Escaping the Gravity of the Mean: Quantile Bayesian Predictive Synthesis for Macro-Financial Tail Risk*

Jiaxun Liu†

October, 2026

Click here for latest version

Abstract

While central banks closely monitor downside macro-financial risks, standard dynamic forecast combination methods available to them typically optimize for average conditions rather than distribution tails. This paper introduces Quantile Bayesian Predictive Synthesis (QBPS), a dynamic framework that combines forecasts through time-varying linear weights optimized directly for specific quantiles. I reparameterize the quantile likelihood to separate tail shocks from baseline volatility, preventing observation variance from collapsing and keeping combination weights from overreacting to single observations. I further introduce an observation-driven discounting mechanism that monitors unexpected tail forecast errors, allowing the combination weights to adapt quickly during sudden economic transitions. Applied to quasi-real-time US Growth-at-Risk forecasting using financial indicators across 1- and 4-quarter horizons, QBPS delivers competitive predictive accuracy overall, while outperforming standalone machine learning forecasts, standard mean-based combinations, and established Growth-at-Risk benchmarks during recessions. Finally, dynamic weight decompositions enable real-time risk attribution: at short horizons, the decomposition highlights intermediary leverage as the primary leading indicator for downside risk, while at extended horizons, the predictive contribution of these financial indicators diminishes.

Keywords: Bayesian Predictive Synthesis, Forecast Combination, Growth-at-Risk, Quantile Regression, Macroeconomic Tail Risk, Asymmetric Laplace Distribution
JEL: C11, C22, C53, E37, E44

*I thank Christian Brownlees and David Rossell for their invaluable guidance and feedback. I am equally grateful to Gary Koop for his hospitality and insights during my research visit to the University of Strathclyde. I also thank Lorenzo Cappello and the brown bag seminar participants for the insightful conversations. All remaining errors are my own.

†Faculty of Economics and Business, Universitat Pompeu Fabra. Email: jiaxun.liu@upf.edu

1 Introduction

Central banks and financial regulators actively monitor and forecast macro-financial downside risks to protect economic stability, reflecting a forward-looking risk-management approach to policy design (Greenspan, 2004; Surico, 2007). To assess these prospective vulnerabilities, institutions consult a wide range of models, motivating forecast combination methods that pool information across them (for reviews, see Timmermann, 2006; Martin et al., 2024). However, standard combination frameworks evaluate candidate models by their ability to forecast the conditional mean and project symmetric intervals around this central estimate. This symmetric structure anchors the forecast distribution to the center of the data mass, creating a pull that I term the "*gravity of the mean*". Because the combination is optimized for the conditional mean, its left tail can remain conservative and slow to register early warning signals when downside risk expands.

This paper introduces Quantile Bayesian Predictive Synthesis (QBPS), a dynamic density forecast combination framework designed to target specific distribution quantiles directly. The framework combines individual forecasts through time-varying linear weights optimized directly for a chosen percentile. This design allows practitioners to combine candidate models based on how well they anticipate severe downturns rather than average conditions. These linear weights also provide a real-time diagnostic, showing which models contribute most to downside forecasts as economic conditions change.

The methodology links dynamic forecast combination with Bayesian quantile modeling. McAlinn and West (2019) and McAlinn et al. (2020) introduced Bayesian Predictive Synthesis (BPS) to pool density forecasts using dynamic linear models (West and Harrison, 1997). While BPS evaluates candidate models assuming no single model is the true data-generating process (an $\mathcal{M}$-open framework; see also Chernis et al., 2025; Adrian et al., 2025), its standard Gaussian likelihood optimizes for performance at the conditional mean. To target specific percentiles directly, the econometric literature turns to dynamic quantile regressions (Gonçalves et al., 2020) and dynamic quantile model averaging (Bernardi et al., 2026). These approaches build on the Asymmetric Laplace distribution, which links asymmetric quantile

loss—penalizing over- and under-predictions differently—to likelihood-based estimation (Yu and Moyeed, 2001) through a conditionally Gaussian representation (Kozumi and Kobayashi, 2011). Recent work by Kobayashi et al. (2026) demonstrates the promise of embedding this quantile architecture into dynamic synthesis settings.

The proposed Quantile Bayesian Predictive Synthesis (QBPS) framework incorporates two core methodological developments.

First, I introduce a stabilizing framework that prevents combination weights from overreacting to individual tail events. In standard state-space setups, dynamic quantile models can suffer from numerical instability when the estimated measurement noise collapses toward zero. When this occurs, the filter misinterprets the vanishing noise as a sign of perfect accuracy, causing the weights to overfit a single observation and discard historical information. To prevent this collapse, I propose a scaled parameterization of the Asymmetric Laplace likelihood, adapting non-centered parameterization principles from Papaspiliopoulos et al. (2003) and Frühwirth-Schnatter (2004). This formulation separates tail shocks from baseline volatility, while an exact conditional update enables closed-form posterior simulation.

Second, I introduce an adaptive discounting mechanism that allows combination weights to adjust quickly during sudden economic transitions. Dynamic linear models typically discount past information at a constant rate (West and Harrison, 1997), which can restrict how fast weights adjust to abrupt shifts in macroeconomic conditions. Influential extensions have relaxed this constraint by searching across discrete grids of discount factors (Koop and Korobilis, 2013) or by adapting the discount rate using forecast errors at the conditional mean (Yusupova et al., 2023). I build on this adaptive tradition by extending observation-driven discounting to distribution tails. Using the Generalized Autoregressive Score (GAS) framework (Creal et al., 2013; Harvey, 2013), the mechanism monitors unexpected tail forecast errors to accelerate weight updating during periods of stress, while preserving stability during tranquil times.

Simulation experiments evaluate QBPS in a controlled environment featuring a business cycle interrupted by an abrupt shift toward downside risk. As an operational "plug-in" framework,

QBPS directly combines quantile forecasts from diverse candidate models alongside simple measures of their predictive uncertainty. When forecasting the 5th percentile ($\tau = 0.05$) during the transition to a turbulent regime, QBPS reallocates combination weights quickly and reduces check loss relative to the individual models and standard Gaussian BPS. The method also performs reliably at the distribution center ($\tau = 0.5$), matching the predictive accuracy of standard BPS at the median.

Central banks and international institutions increasingly monitor downside macroeconomic vulnerability through the "Growth-at-Risk" (GaR) framework, which uses financial conditions to forecast the lower tail of GDP growth (e.g., Adrian et al., 2019; Brownlees and Souza, 2021; Amburgey and McCracken, 2023, 2026). This macro-financial focus makes GaR a natural empirical setting to evaluate QBPS. Using quasi-real-time data for the United States, I evaluate downside forecasts of real GDP growth at the 5th percentile across 1- and 4-quarter horizons. The candidate pool pairs a baseline model based on lagged GDP growth alone with models augmented by the aggregate Chicago Fed financial conditions index and its subcomponents. The candidate models rely on Quantile Random Forests (Meinshausen, 2006) to capture non-linear patterns in macro-financial data (Clark et al., 2023, 2024; Delle Monache et al., 2024). However, tree-based models cannot project outcomes beyond historical extremes. The synthesis combines these non-linear forecasts while using a dynamic intercept to adjust the forecast level when unprecedented economic contractions occur.

Out-of-sample evaluations show that QBPS delivers competitive predictive accuracy across the full evaluation sample, while outperforming standalone tree models, standard Gaussian BPS (McAlinn and West, 2019), and the foundational two-step Growth-at-Risk benchmark (Adrian et al., 2019) during recessions. The synthesis also achieves the nominal 5% frequency of extreme tail realizations, correcting the tendency of history-bounded tree models to underpredict severe downturns. Translating the dynamic weights into the added forecast contribution of each financial model beyond the lagged-GDP baseline provides a real-time risk attribution. At the 1-quarter horizon, this decomposition highlights intermediary leverage as the primary leading indicator for downside risk, in line with recent empirical findings (Amburgey and McCracken, 2026). At the 4-quarter horizon, these financial contributions

attenuate, reflecting the natural decay of financial leading signals at extended forecast horizons (Adrian et al., 2019).

The remainder of this paper is organized as follows. Section 2 introduces the QBPS framework, detailing the scaled parameterization and the observation-driven discounting mechanism. Section 3 demonstrates how QBPS adapts and performs across simulated business cycle regimes and abrupt shifts in downside risk. Section 4 reports the empirical application to US Growth-at-Risk, examining out-of-sample accuracy, dynamic synthesis weights, and the real-time decomposition of financial signals. Section 5 concludes.

2 Quantile Bayesian Predictive Synthesis

2.1 Bayesian Predictive Synthesis and the Symmetric Constraint

Motivating the direct evaluation of extreme quantiles begins with examining the structural limitations of state-of-the-art dynamic combination models. Let y_t denote a macroeconomic or financial target variable. Consider a pool of candidate agents $k = 1, \ldots, K$, where each agent provides a latent forecast $x_{k,t}$ for y_t based on the available information set $\mathcal{D}_{t-1}$, characterized by a predictive density $h_{k,t}(x_{k,t})$. The foundational Bayesian Predictive Synthesis (BPS) framework (McAlinn and West, 2019) defines the combined posterior predictive density of y_t as the integral over the joint agent predictive density $H_t(\mathbf{x}_t)$:

$$p(y_t|\mathcal{D}_{t-1}) = \int \alpha_t(y_t|\mathbf{x}_t) H_t(\mathbf{x}_t) d\mathbf{x}_t \tag{1}$$

where $\mathbf{x}_t = (x_{1,t}, \ldots, x_{K,t})'$ is the vector of latent agent forecasts, assuming conditional independence such that $H_t(\mathbf{x}_t) = \prod_{k=1}^{K} h_{k,t}(x_{k,t})$. The core of the BPS framework is the synthesis density, $\alpha_t(y_t|\mathbf{x}_t)$, which calibrates and weights the individual agents.

The integral in equation (1) represents a structural generalization of traditional forecast combination methods. As shown by McAlinn and West (2019), standard Bayesian Model

Averaging (BMA) (Hoeting et al., 1999; Geweke and Amisano, 2011) is mathematically equivalent to a restricted special case of BPS where the synthesis density is constrained to a static probability simplex. By lifting these constraints, BPS flexibly accounts for time-varying agent misspecification and cross-sectional correlation.

To operationalize this integral, the standard BPS literature assumes an underlying dynamic linear structure. By augmenting the latent agent forecasts $\mathbf{x}_t$ with a deterministic intercept to form the state predictor vector $\mathbf{F}_t = (1, \mathbf{x}_t')'$, the synthesis density $\alpha_t(y_t|\mathbf{x}_t)$ is parameterized as a conditionally Gaussian likelihood:

$$\alpha_t(y_t|\mathbf{x}_t, \boldsymbol{\theta}_t, \sigma_t) = \mathcal{N}(y_t|\mathbf{F}_t'\boldsymbol{\theta}_t, \sigma_t^2) \tag{2}$$

where $\boldsymbol{\theta}_t = (\theta_{t0}, \theta_{t1}, \ldots, \theta_{tK})'$ serves as the dynamic weight vector evolving via a latent state transition equation. This Gaussian specification is convenient because it preserves Dynamic Linear Model (DLM) conjugacy, enabling highly efficient, closed-form inference via the Forward Filtering Backward Sampling (FFBS) algorithm (West and Harrison, 1997; Prado and West, 2010).

While computationally elegant, evaluating dynamic weights by maximizing this Gaussian likelihood is mathematically equivalent to minimizing a symmetric, squared-error loss function. This specification is highly efficient during stable expansions, but macro-financial crises typically manifest as asymmetric deviations that skew the risk distribution.

When an extreme tail event occurs, this squared-error equivalence makes the synthesis highly sensitive to the outlier. Yet, constrained by a strictly symmetric shape, the model cannot independently stretch a single tail. To absorb the outlier, it must either shift the entire conditional mean or symmetrically inflate the variance—an adjustment that mathematically projects an equally probable extreme shift in the opposite direction. To balance these competing quadratic penalties, the synthesis function naturally anchors the distribution near the center of the data mass. I term this structural compromise the "gravity of the mean". Consequently, while the framework actively processes the outlier, its symmetric architecture restricts its capacity to quantify asymmetric tail risk.

Explicitly capturing this tail vulnerability motivates a shift from the standard Gaussian likelihood to an objective function optimized directly at the extreme quantile of interest.

2.2 The Zero-Variance Trap and Scaled Stabilization

To directly optimize the dynamic weights for a specific extreme quantile τ, the synthesis framework shifts from a symmetric Gaussian likelihood to the Asymmetric Laplace (AL) distribution[1]. The general Bayesian quantile synthesis equation is thus defined as:

$$\alpha_t(y_t|\mathbf{x}_t, \boldsymbol{\theta}_t^\tau, v_t, \tau) = \mathcal{AL}(y_t|\mathbf{F}_t'\boldsymbol{\theta}_t^\tau, \sqrt{v_t}, \tau) \tag{3}$$

where $\boldsymbol{\theta}_t^\tau$ is the dynamic weight vector optimized for each specific quantile τ, and v_t represents the structural variance.

Because the AL distribution lacks natural conjugacy to the Gaussian priors required for state-space modeling, it is conventionally represented as a conditionally Gaussian normal-exponential mixture[2]. To embed this target likelihood into a dynamic setting while retaining the computational efficiency of the Forward-Filtering Backward-Sampling (FFBS) algorithm, I introduce a *scaled* parameterization that structurally decouples the latent shock from the structural variance. The Quantile BPS (QBPS) hierarchy is fully specified as:

$$\begin{aligned} y_t|\mathbf{F}_t, \boldsymbol{\theta}_t^\tau, e_t, v_t, \tau &\sim \mathcal{N}(\mathbf{F}_t'\boldsymbol{\theta}_t^\tau + a_\tau\sqrt{v_t}e_t,\ b_\tau v_t e_t) \\ \boldsymbol{\theta}_t^\tau|\boldsymbol{\theta}_{t-1}^\tau, \mathbf{W}_t, v_t &\sim \mathcal{N}(\boldsymbol{\theta}_{t-1}^\tau,\ v_t\mathbf{W}_t) \\ e_t &\sim \text{Exp}(1) \\ v_t|\mathcal{D}_{t-1} &\sim \mathcal{IG}(\beta n_{t-1}/2,\ \beta n_{t-1}s_{t-1}/2) \end{aligned} \tag{4}$$

where e_t is a decoupled standard exponential shock, v_t is the structural variance modeled

[1]The AL probability density function is defined as $f(y|\mu, \sigma, \tau) = \frac{\tau(1-\tau)}{\sigma}\exp\left(-\frac{\rho_\tau(y-\mu)}{\sigma}\right)$, where $\rho_\tau(e) = e(\tau - \mathbb{I}(e < 0))$ denotes the asymmetric check-loss function.

[2]By representing the AL distribution as a normal-exponential mixture, the target asymmetric likelihood is recovered marginally, ensuring the conditional space remains strictly Gaussian at each MCMC step to preserve exact FFBS conjugacy.

via an Inverse-Gamma ($\mathcal{IG}$) distribution, $a_\tau = \frac{1-2\tau}{\tau(1-\tau)}$ and $b_\tau = \frac{2}{\tau(1-\tau)}$. To govern the time-varying adaptation of the synthesis weights, the state evolution variance is defined via sequential variance discounting as $\mathbf{W}_t = \frac{1-\delta_t}{\delta_t}\mathbf{C}_{t-1}$, where δ_t is the structural discount factor and $\mathbf{C}_{t-1}$ is the unscaled posterior state covariance matrix from the previous period. Similarly, the evolution of the structural variance v_t is governed by a fixed volatility discount factor β, where n_{t-1} and s_{t-1} denote the prior degrees of freedom and the point estimate of the variance, respectively.

The foundational equivalence between the Asymmetric Laplace distribution and the normal-exponential mixture was originally established by Kozumi and Kobayashi (2011). Proposition 1 confirms that the proposed scaled parameterization mathematically preserves this property. While this result is a direct corollary of the original static mixture theorem, a full analytical integration is provided in Appendix A for completeness.

Proposition 1 (Likelihood Equivalence of the Scaled Hierarchy) *Let the latent shock* $e_t \sim Exp(1)$ *be independent of the structural variance* v_t*. Marginalizing over* e_t *in the scaled hierarchy exactly recovers the Asymmetric Laplace likelihood:*

$$p(y_t|\mathbf{F}_t, \boldsymbol{\theta}_t^\tau, v_t, \tau) = \int_0^\infty \mathcal{N}(y_t|\mathbf{F}_t'\boldsymbol{\theta}_t^\tau + a_\tau\sqrt{v_t}e_t,\ b_\tau v_t e_t) Exp(e_t|1) de_t = \mathcal{AL}(\mathbf{F}_t'\boldsymbol{\theta}_t^\tau, \sqrt{v_t}, \tau) \quad (5)$$

Proof. See Appendix A.

Because the target AL likelihood is fully recovered marginally at every time step, the framework gains the mathematical flexibility to restructure the underlying variance parameters. This serves as a deliberate functional departure from the standard static literature. In traditional static Bayesian quantile regression (Kozumi and Kobayashi, 2011), the model is constructed using a *coupled* parameterization where the distribution of the latent mixing shock u_t depends explicitly on a fixed scale parameter σ:

$$\begin{aligned} y_t|\mathbf{F}_t, \boldsymbol{\theta}^\tau, u_t, \sigma, \tau &\sim \mathcal{N}(\mathbf{F}_t'\boldsymbol{\theta}^\tau + a_\tau u_t,\ b_\tau\sigma u_t) \\ u_t|\sigma &\sim \text{Exp}(1/\sigma) \\ \sigma &\sim \mathcal{IG}(n_0/2,\ s_0/2) \end{aligned} \quad (6)$$

where the latent mixing shock u_t is mathematically tied to the static scale σ, $a_\tau = \frac{1-2\tau}{\tau(1-\tau)}$, $b_\tau = \frac{2}{\tau(1-\tau)}$, and the hyperparameters n_0 and s_0 denote the fixed prior degrees of freedom and prior scale, respectively.

Transitioning to a dynamic synthesis engine requires the static parameters to evolve over time as $\boldsymbol{\theta}_t^\tau$ and σ_t. However, directly embedding the coupled architecture from (6) into a sequential filter triggers a "zero-variance trap". The dynamic framework relies on a Gibbs sampler that continuously alternates between drawing latent shocks and updating the time-varying weights via the Kalman filter. The conditional observation variance $b_\tau\sigma_t u_t$ is deterministically tied to the mixing shock via $u_t|\sigma_t \sim Exp(1/\sigma_t)$. Consequently, a near-zero MCMC draw for u_t mechanically collapses the observation variance to zero[3]. The Kalman filter misinterprets

[3]In the sequential update, the Kalman gain is defined as $\mathbf{A}_t = \mathbf{R}_t\mathbf{F}_t(\mathbf{F}_t'\mathbf{R}_t\mathbf{F}_t + V_{obs})^{-1}$, where $\mathbf{R}_t$ is

this collapsed variance as a perfect measurement. Stripped of its historical anchor, the filter forces the dynamic weights to completely absorb the empirical outlier, causing the model to diverge.

To resolve this instability, the scaled hierarchy in (4) restructures the states by applying the transformations $e_t = u_t/\sigma_t$ and $v_t = \sigma_t^2$. This algebraically maps the original coupled parameters (σ_t, u_t) from (6) directly into the decoupled states (v_t, e_t) of (4). By transforming the latent mixing shock into a standard exponential variable $e_t \sim \text{Exp}(1)$, the framework breaks the deterministic link between the AL shock and the structural volatility, improving the stability of the filter during the Gibbs sampling loop.

This decoupling requires a corresponding modification to the prior. Sequential variance discounting demands a dynamic parameterization, but assuming a time-varying scale $\sigma_t|\mathcal{D}_{t-1} \sim \mathcal{IG}(\beta n_{t-1}/2,\ \beta n_{t-1}s_{t-1}/2)$ forces the structural variance $v_t = \sigma_t^2$ into a squared Inverse-Gamma distribution. This breaks the Normal-Gamma conjugacy required for analytical filtering. The scaled hierarchy bypasses this constraint by placing the $\mathcal{IG}$ prior directly on v_t, restoring the convenient DLM conjugacy.

Finally, stabilizing the Kalman filter through decoupling introduces a distinct sampling challenge. During the MCMC evaluation, if the decoupled shock e_t were drawn purely from its static prior Exp(1), the draws would fail to condition on the realized observation. Consequently, when an extreme macroeconomic outlier occurs, this prior-drawn shock would fail to stretch and accommodate the error, allowing the transient noise to destabilize the core synthesis weights ($\boldsymbol{\theta}_t^\tau$).

To prevent the weights from absorbing extreme tail events, e_t is sampled conditionally on the realized forecast error. As established in Proposition 2, the full conditional posterior of the latent shock resolves exactly to a Generalized Inverse Gaussian (GIG) distribution.

the prior state covariance and V_{obs} is the conditional observation variance. If $V_{obs} \to 0$, the Kalman gain simplifies such that the scalar projection $\mathbf{F}_t'\mathbf{A}_t \to 1$. Because the posterior state mean (of $\boldsymbol{\theta}_t^\tau$) is updated via $\boldsymbol{m}_t = \boldsymbol{m}_{t-1} + \mathbf{A}_t(y_t - \mathbf{F}_t'\boldsymbol{m}_{t-1})$, this limit forces the newly updated forecast $\mathbf{F}_t'\boldsymbol{m}_t$ to equal y_t exactly. The dynamic weights then completely abandon their historical memory to over-fit the single observation.

Proposition 2 (Exact Conditional Posterior of the Latent AL Shock) *Let* $\mu_t = \mathbf{F}_t'\boldsymbol{\theta}_t^{\tau}$. *Under the scaled hierarchy, the full conditional posterior distribution of the latent standard exponential shock* e_t *is exactly a Generalized Inverse Gaussian distribution:*

$$e_t | y_t, \mu_t, v_t, \tau \sim \mathcal{GIG}\left(\lambda = 0.5,\ \chi = \frac{(y_t - \mu_t)^2}{b_\tau v_t},\ \psi = \frac{a_\tau^2}{b_\tau} + 2\right) \tag{7}$$

Proof. See Appendix B.

By resampling e_t conditionally via this exact GIG posterior update, the latent shock acts as a localized shock absorber. It inflates to process extreme macro-financial outliers, preserving the stability of the dynamic agent weights and ensuring robust MCMC convergence across structural regime breaks. The complete MCMC block sampling procedure is detailed in the Internet Appendix.

2.3 Observation-Driven Asymmetric Discounting

To process the large data deviations inherent in tail-risk forecasting, predictive synthesis models navigate a stability-plasticity dilemma through sequential variance discounting. Leveraging convenient conjugacy for analytical updating within the Forward-Filtering Backward-Sampling (FFBS) algorithm (West and Harrison, 1997; Prado and West, 2010), the transition from the time $t-1$ posterior to the time t prior dynamically inflates the state covariance matrix via $\mathbf{R}_t = \mathbf{C}_{t-1}/\delta_t$. Here, the discount factor δ_t acts to mathematically regularize the structural variance injected into the filter. Existing frameworks typically fix this injection rate with a static discount factor δ. The foundational Bayesian Predictive Synthesis (BPS) framework (McAlinn and West, 2019) suggests a range (e.g., $\delta \in [0.9, 0.99]$) for baseline stability, whereas recent dynamic quantile extensions (Kobayashi et al., 2026) advocate for a smaller static discount factor (e.g., $\delta = 0.85$) to accelerate adaptation.

However, a static variance regulator introduces a structural trade-off under the scaled Asymmetric Laplace (AL) hierarchy. Targeting extreme quantiles (e.g., $\tau = 0.05$) substantially

inflates the conditional observation variance via the multiplier b_τ[4]. This high variance renders the tail likelihood structurally flat, leaving the Kalman filter heavily dependent on the injected prior state covariance $\mathbf{R}_t$ to anchor the dynamic weights. Consequently, employing a small discount factor to encourage rapid adaptation becomes mathematically problematic: inflating the prior variance against a flat likelihood compromises the filter's regularizing anchor, causing the dynamic weights to become overly sensitive to standard macroeconomic noise. To prevent this instability, the baseline discount factor demands a tight initialization (e.g., $\delta_{max} = 0.995$).

Yet, maintaining this tight baseline creates a fixed-memory inertia. While a static $\delta = 0.995$ ensures the stability of tail forecasts during normal business cycle fluctuations, it strictly limits the amount of new variance injected into the system. When a structural break or regime shift strikes, the model adapts too slowly to the new data-generating process, resulting in prolonged forecast deviations and cumulative check-loss penalties.

To resolve this, I introduce an observation-driven asymmetric discounting module grounded in the Generalized Autoregressive Score (GAS) framework (Creal et al., 2013; Harvey, 2013). Building on recent gradient-based advances in adaptive discounting (Yusupova et al., 2023), this module endogenizes the filter's plasticity, offering a continuous updating mechanism that extends traditional discrete grid-search methodologies (e.g., Koop and Korobilis, 2013).

The mechanism operates through a *two-step* architecture: first, it constructs an observation-driven state variable by extracting structural stress signals directly from the gradient of the AL log-likelihood; second, it maps this state variable through a calibrated update rule to execute the asymmetric discounting. The filter links δ_t directly to the likelihood. When macro-financial regime shifts occur, the filter dynamically drops the discount factor to inject a targeted surge of variance.

Starting with the observation-driven component, the GAS framework extracts its driving

[4]Because the variance multiplier is $b_\tau = \frac{2}{\tau(1-\tau)}$, forecasting extreme tail quantiles dramatically increases the observation variance relative to the median ($b_{0.05} \approx 42$ vs. $b_{0.5} = 8$). This mathematically diffuses the likelihood as we target extreme quantiles, meaning the probability density is widely dispersed and individual observations provide weak information regarding the exact location of the parameters.

signal by differentiating the AL log-likelihood with respect to the structural scale parameter, $\sqrt{v_t}$. Scaling this resultant gradient yields a dimensionless gauge of asymmetric predictive surprise.

Proposition 3 (The AL Score as the Normalized Adaptation Gradient) *In asymmetric forecasting frameworks, the scaled Asymmetric Laplace score isolates the core dimensionless innovation. Let Z_t denote this "asymmetric surprise" utilized by the asymmetric discounting mechanism:*

$$Z_t = \frac{\rho_\tau(y_t - \mu_t)}{\sqrt{v_t}} \tag{8}$$

where $\mu_t = \mathbf{F}_t'\boldsymbol{\theta}_t^\tau$ is the unshifted location forecast. By dividing the realized directional tail shock $\rho_\tau(y_t - \mu_t)$ by the model's localized structural capacity $\sqrt{v_t}$, the statistic Z_t acts as the normalized gradient driving local parameter adaptation, independent of the baseline variance.

Proof. See Appendix C.

To utilize Z_t as a dynamic structural test, its exact distribution is first defined under normal market conditions, where "stability" in a quantile forecasting framework implies that the true data-generating process matches the model's predictive distribution.

Proposition 4 (Marginal Distribution of the Scaled Score) *Let the null hypothesis of structural stability (H_0) denote a state where the true data-generating process matches the model's predictive density. Under H_0, the observations y_t are conditionally distributed as an Asymmetric Laplace distribution, $\mathcal{AL}(\mu_t, \sqrt{v_t}, \tau)$. Consequently, the sequence of scaled check-losses constitutes independent and identically distributed (i.i.d.) Standard Exponential variables:*

$$Z_t \sim Exp(1) \tag{9}$$

Proof. See Appendix D.

A statistical rejection of this null hypothesis statistically indicates a structural regime shift. In macro-financial applications, this instability is frequently associated with the sudden

onset of binding constraints, such as intermediary capital limits or margin spirals (He and Krishnamurthy, 2013; Brunnermeier and Sannikov, 2014). Regardless of the specific trigger, when the baseline data-generating process fractures, the current structural variance becomes inadequate. To dynamically track this breakdown, the discounting mechanism smooths the stationary innovations via an autoregressive process, generating an internal state variable S_t that measures sustained structural stress:

$$S_t = \rho S_{t-1} + (1 - \rho) Z_t \tag{10}$$

Under the null hypothesis of economic stability (H_0), the expected value of the smoothed state is $\mathbb{E}[S_t] = 1$. The activation threshold is calibrated to function as an early-warning signal. By setting the autoregressive persistence to $\rho = 0.3$, the smoothed state can be approximated by a Gamma distribution via the method of moments. Because the internal state S_t expands as an infinite geometrically weighted sum of standard exponential innovations, the Gamma distribution provides a highly principled approximation. Unlike a standard normal approximation, it strictly preserves the $[0, \infty)$ support and the heavy right-skewness inherent to these underlying exponentials. This mapping allows us to establish a theoretical False Discovery Rate (FDR). By setting the activation threshold at $\overline{S} = 1.75$, the derived FDR is approximately 11.2%[5].

While this roughly 1σ threshold appears statistically lenient relative to standard econometric conventions, it represents a deliberate design choice for tail-risk forecasting: the asymmetric cost of missed crises is highly consequential. Waiting for traditional 3σ statistical certainty during a rapidly compounding liquidity freeze constrains the filter's performance during periods that require rapid adaptation.

To properly capture the macro-financial reality that crises materialize abruptly but dissipate gradually—a temporal asymmetry central to the literature on volatility clustering (e.g.,

[5]Assuming stationarity under H_0, the AR(1) process yields a variance of $\text{Var}(S_t) = \frac{1-\rho}{1+\rho} \approx 0.54$. By mapping the theoretical moments $\mathbb{E}[S_t] = 1$ and $\text{Var}(S_t) = 0.54$ to a Gamma distribution $\mathcal{G}(k, \theta)$, the shape parameter resolves to $k = \mathbb{E}[S_t]^2/\text{Var}(S_t) \approx 1.86$ and the scale parameter to $\theta = \text{Var}(S_t)/\mathbb{E}[S_t] \approx 0.54$. Evaluating the general survival function $P(S_t \geq \overline{S}) = 1 - \int_0^{\overline{S}} \frac{1}{\Gamma(k)\theta^k} x^{k-1} e^{-x/\theta} dx$ at the threshold $\overline{S} = 1.75$ yields an upper-tail probability (FDR) of approximately 11.2%.

Engle, 1982; Bollerslev, 1986; Nelson, 1991; Glosten et al., 1993)—the mechanism utilizes a path-dependent, asymmetric update rule. 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}:

$$\delta_t^* = \begin{cases} \max\left(\delta_{min},\ \delta_{max} - \kappa_{drop}(S_{t-1} - \overline{S})\right) & \text{if } S_{t-1} \geq \overline{S} \\ \delta_{max} & \text{otherwise} \end{cases} \tag{11}$$

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:

$$\delta_t = \min\left(\delta_t^*,\ \delta_{t-1} + \kappa_{recover}\right) \tag{12}$$

When the score rejects the null hypothesis ($S_{t-1} \geq \overline{S}$), the memory target drops proportionally to the excess gradient, calibrated to $\kappa_{drop} = 0.15$, subject to a structural ceiling at $\delta_{max} = 0.995$ and a floor at $\delta_{min} = 0.95$. This floor ensures that during episodes of severe predictive incoherence, the model amplifies its baseline rate of variance injection approximately tenfold. When an acute shock materializes, the update rule instantly decreases the discount factor to meet this target, breaking the fixed-memory inertia with a targeted surge of plasticity. Conversely, when the acute shock dissipates, the recovery micro-step ($\kappa_{recover} = 0.01$) acts as a speed limit, leading a graduated variance decay back to the baseline ceiling. This paired asymmetry—an instantaneous variance injection followed by a graduated recovery—allows QBPS to cross regime gaps abruptly while stably locking the newly discovered synthesis weights into place.

2.4 Posterior Inference and Quantile Forecasting

To navigate the high-dimensional parameter space of the non-centered hierarchy, posterior inference is conducted via an efficient, three-block Markov Chain Monte Carlo (MCMC) Gibbs sampler.

In the first block, conditioned on the latent AL shocks and agent states, the dynamic synthesis weights $\boldsymbol{\theta}^{\tau}_{1:T}$ and structural variances $v_{1:T}$ are jointly sampled using a Forward-Filtering Backward-Sampling (FFBS) algorithm with the observation-driven asymmetric discounting module detailed in section 2.3. The second block conditionally resamples the latent predictive agent states $\mathbf{X}_{1:T}$ and their associated precisions, incorporating the inherent predictive uncertainty of the individual forecast models. Finally, the third block processes extreme outliers by exactly resampling the decoupled Asymmetric Laplace shocks $e_{1:T}$ from the Generalized Inverse Gaussian (GIG) posterior detailed in section 2.2. Complete sampling algorithms are provided in the Internet Appendix.

Once the posterior distributions of the dynamic states are obtained, QBPS then generates out-of-sample forecast distributions. The synthesis weights $\boldsymbol{\theta}^{\tau}_{T+k}$ are anchored strictly to the terminal in-sample posterior mean, which naturally serves as the out-of-sample prior[6]. Simultaneously, consistent with the forecasting architecture of McAlinn and West (2019), the k-step ahead candidate vector $\mathbf{F}_{T+k}$ is populated using independent location-scale Student-t draws derived from the individual agents' submitted forecast moments. Under the model's structural hierarchy, projecting these sampled candidate forecasts against the synthesis weights directly defines the predictive Asymmetric Laplace location parameter:

$$\hat{y}^{\tau}_{T+k} = \mathbf{F}'_{T+k}\boldsymbol{\theta}^{\tau}_{T+k} \tag{13}$$

By evaluating this projection iteratively across thousands of MCMC draws, QBPS constructs the empirical posterior predictive distribution of the targeted extreme quantile. This provides a structural advantage for tail-risk analysis: the model not only dynamically tracks the τ-th

[6]Under the structural random walk prior detailed in section 2.2, the terminal posterior mean mathematically defines the expected out-of-sample trajectory.

percentile of the risk distribution but natively outputs corresponding credible intervals to quantify the conditional parameter uncertainty under the assumed AL likelihood.

2.5 Computational Stabilization Toolkit

Beyond the non-centered parameterization and the GAS module, optimizing dynamic linear models directly at extreme quantiles requires robust numerical architecture. As the target approaches the extreme tails (e.g., $\tau = 0.05$), the AL conditional variance multiplier b_τ sharply inflates, exposing the FFBS algorithm to numerical instability. To maintain filter cohesion, QBPS employs two structural stabilizers during estimation.

First, the algorithm replaces the standard Kalman filter covariance update with the Joseph stabilized form (Bucy and Joseph, 2005; Simon, 2006). Rather than utilizing standard matrix subtraction to update the state covariance, the posterior matrix is computed via addition:

$$\mathbf{C}_t = (\mathbf{I} - \mathbf{A}_t\mathbf{F}_t')\mathbf{R}_t(\mathbf{I} - \mathbf{A}_t\mathbf{F}_t')' + \mathbf{A}_t V_{obs}\mathbf{A}_t' \tag{14}$$

Because it is the sum of two positive semi-definite matrices, this formulation mathematically guarantees that the state covariance matrices remain positive-definite, even when processing the highly inflated observation variances characteristic of AL tail shocks.

Second, to ensure that the inflated tail variance does not induce spurious parameter drift during the MCMC burn-in phase, the prior covariance for the candidate agents is anchored using a tight precision variance (e.g., $\mathbf{C}_0 = 0.01 \cdot \mathbf{I}_K$). This concentrated initialization suppresses the initial Kalman gain, constraining the dynamic weights to navigate the volatile parameter space rather than overfitting the extreme noise inherent to the AL formulation.

3 Simulation Studies: Structural Shifts and Estimation Inertia

To evaluate the predictive tracking and dynamic model synthesis of the proposed econometric agents, this simulation specifies a quasi-empirical Data Generating Process (DGP) featuring a smooth macroeconomic cycle, time-varying conditional volatility, and a structural shift in skewness. The full sample length is $T = 400$.

3.1 The Quasi-Empirical Data Generating Process

The target scalar series y_t is generated conditionally on the lag of underlying latent state variables via a skew-normal distribution:

$$y_t \mid \mathcal{F}_{t-1} \sim \mathcal{SN}\big(\xi = X_{t-1}^{\text{macro}},\, \omega = X_{t-1}^{\text{vix}},\, \alpha = X_{t-1}^{\text{skew}}\big), \quad t = 2, \ldots, T \tag{15}$$

which explicitly decomposes macroeconomic risk into location ($\xi \in \mathbb{R}$), scale ($\omega > 0$), and shape/slant ($\alpha \in \mathbb{R}$). The true conditional quantile targeted by the agents for a given coverage level $\tau \in (0, 1)$ is evaluated directly via the inverse cumulative distribution function of the skew-normal distribution. To simulate a business cycle punctuated by a sudden crisis, the underlying states evolve dynamically. The macroeconomic baseline follows a slow-moving, low-frequency cyclical path $X_t^{\text{macro}} = 5 + \sin(0.25\theta_t)$ over an equidistant time grid $\theta_t \in [0, 10\pi]$, calibrating the underlying mean to complete approximately one full macroeconomic cycle over the sample period. At index $t^* = 300$, a structural transition occurs, governed by a piecewise-linear weight $w_t = \min\{\max(\frac{t-300}{5}, 0), 1\}$. This transition triggers a simultaneous shift in both volatility and tail asymmetry: the conditional scale expands via $X_t^{\text{vix}} = 1.5 + 1.5w_t + \varepsilon_t^{\text{vix}}$ (where $\varepsilon_t^{\text{vix}} \sim \mathcal{N}(0, 0.3^2)$), and the shape parameter reverses via $X_t^{\text{skew}} = 2(1 - w_t) - 2w_t + \varepsilon_t^{\text{skew}}$ (where $\varepsilon_t^{\text{skew}} \sim \mathcal{N}(0, 0.5^2)$). Consequently, prior to the structural shift ($w_t = 0$), the baseline scale fluctuates around 1.5 with positive skewness centered at +2. In the post-transition regime ($w_t = 1$), conditional dispersion doubles to

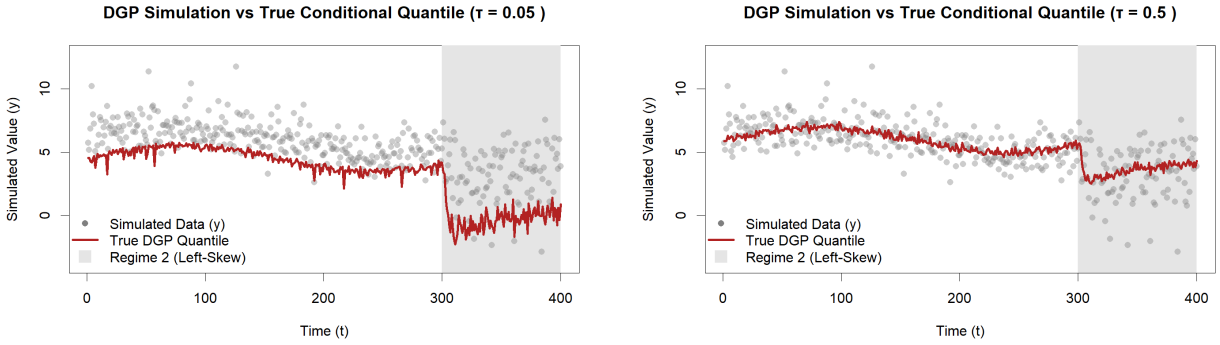
an average level of 3.0 and the slant abruptly reverses to negative skewness centered at -2, inducing a sharp expansion in downside risk.

To emulate empirical tracking under imperfect information, the candidate econometric agents observe noisy proxies of the underlying latent drivers rather than the true parameter trajectories. While the macroeconomic baseline is observed perfectly ($\widetilde{X}_t^{\text{macro}} = X_t^{\text{macro}}$), the volatility and skewness states are contaminated by independent Gaussian noise, such that $\widetilde{X}_t^{\text{vix}} = X_t^{\text{vix}} + \eta_t^{\text{vix}}$ and $\widetilde{X}_t^{\text{skew}} = X_t^{\text{skew}} + \eta_t^{\text{skew}}$, where $\eta_t^{\text{vix}} \sim \mathcal{N}(0, 1.0^2)$ and $\eta_t^{\text{skew}} \sim \mathcal{N}(0, 2.0^2)$.

Because y_t depends on latent variables dated $t-1$, candidate agents form 1-step-ahead quantile and mean forecasts for period t strictly utilizing information available up to period $t-1$, ensuring an out-of-sample expanding-window evaluation free of look-ahead bias.

Figure 2 illustrates this symmetric behavior within the simulated environment. Panels (a) and (b) plot the true DGP alongside the theoretical 5th and 50th percentiles, respectively. At the onset of the structural transition ($t = 300$), the left tail expands significantly downward, whereas the median exhibits a structurally smoother and more bounded adjustment.

Figure 1: Simulated DGP & Theoretical Quantiles



(a) DGP and Theoretical Quantiles ($\tau = 0.05$)

(b) DGP and Theoretical Quantiles ($\tau = 0.5$)

Notes: **(a) & (b):** The plots show the raw simulated DGP according to (15) and the theoretical quantiles implied by the parameters of the skewed normal distribution at (a) $\tau = 0.05$ and (b) $\tau = 0.5$.

3.2 Candidate Agents and Variance Calibration

To evaluate dynamic synthesis performance during non-linear regime shifts, both the Gaussian BPS and QBPS engines require a vector of candidate point forecasts, $\boldsymbol{a}_t = (a_{t,1}, \ldots, a_{t,J})'$, and an associated predictive uncertainty matrix, $A_t = \text{diag}(A_{t,1}, \ldots, A_{t,J})$, to scale the latent agent precision draws. Because standard Gaussian BPS optimizes for the conditional mean under a symmetric squared-error loss, it explicitly requires the agents' mean forecasts ($\boldsymbol{a}_t^{\text{mean}}$). In contrast, because QBPS directly targets a specific tail via the Asymmetric Laplace likelihood, it receives the agents' quantile forecasts ($\boldsymbol{a}_t^{\text{quant}}$). I deploy $J = 4$ distinct parametric and semi-parametric econometric agents to generate these dual inputs.

A distinct advantage of this framework is its $\mathcal{M}$-open flexibility. Because QBPS does not assume the true data-generating process exists within the candidate pool, it potentially synthesizes forecasts across disparate model classes. By deploying standard frequentist models as base agents, this evaluation highlights the engine's "plug-and-play" utility. While sequential filtering formally requires full predictive densities to scale state precision, practitioners can feed off-the-shelf point forecasts directly into the dynamic state vector. Pairing these forecasts with a calibrated variance proxy satisfies this requirement—ensuring a reasonable level playing field.

Agents $j \in \{1, 3, 4\}$ constitute a unified class of linear regression experts capturing the macroeconomic baseline, volatility, and asymmetric tail, respectively. Let $x_{t-1}^{(j)}$ denote the specific covariate for agent j, where $x^{(1)} = \widetilde{X}_t^{\text{macro}}$, $x^{(3)} = \widetilde{X}_t^{\text{vix}}$ and $x^{(4)} = \widetilde{X}_t^{\text{skew}}$. These agents produce conditional mean forecasts via Ordinary Least Squares (OLS) and direct quantile forecasts via linear Quantile Regression (QR):

$$a_{t,j}^{\text{mean}} = \widehat{\beta}_{0,t}^{\text{OLS}} + \widehat{\beta}_{1,t}^{\text{OLS}} x_{t-1}^{(j)}, \tag{16}$$

$$a_{t,j}^{\text{quant}} = \widehat{\beta}_{0,t}^{\text{QR}}(\tau) + \widehat{\beta}_{1,t}^{\text{QR}}(\tau) x_{t-1}^{(j)}. \tag{17}$$

Agent 2 is a univariate AR(1)–GARCH(1,1) model[7] capturing autoregressive return dynam-

[7]The complete AR(1)-GARCH(1,1) specification is defined by the mean equation $y_t = c + \phi y_{t-1} + \varepsilon_t$,

ics and dynamic conditional volatility. Letting $\widehat{\mu}_{t|t-1}$ denote the AR(1) conditional mean and $\sigma^2_{t|t-1}$ the 1-step-ahead GARCH conditional variance, its forecasts are:

$$a^{\text{mean}}_{t,2} = \widehat{\mu}_{t|t-1}, \tag{18}$$

$$a^{\text{quant}}_{t,2} = \widehat{\mu}_{t|t-1} + \Phi^{-1}(\tau)\,\sigma_{t|t-1}, \tag{19}$$

where $\Phi^{-1}(\cdot)$ is the standard Gaussian quantile function.

To accommodate the lag structure required by the AR-GARCH model, the candidate agents are initially estimated over a 100-period historical window from $t = 2$ to $t = 101$. Following this initialization, all agents generate 1-step-ahead forecasts recursively from $t = 102$ to T using an expanding historical window. Next, the Bayesian synthesis engines evaluate these candidate forecasts over a 100-period rolling window, yielding a final out-of-sample evaluation period that spans $t = 202$ to T. This setup is designed to evaluate the mechanics of estimation inertia. While the expanding-window agents naturally remain anchored to pre-transition data, this structure allows us to test whether the rolling-window synthesis maintains the capacity to adapt to sudden regime shifts.

Because BPS and QBPS optimize for different statistical objectives, they process distinct candidate inputs and extract their final out-of-sample quantile predictions differently:

- **Gaussian BPS (Conditional Mean Synthesis):** Utilizing the mean inputs ($\boldsymbol{a}^{\text{mean}}_t$), BPS generates a synthesized Gaussian predictive density $\mathcal{N}(f^{\text{BPS}}_t, q^{\text{BPS}}_t)$. The out-of-sample τ-th quantile prediction is then extracted analytically post-estimation as $\hat{y}^{(\tau,\text{BPS})}_t = f^{\text{BPS}}_t + \Phi^{-1}(\tau)\sqrt{q^{\text{BPS}}_t}$.

- **Quantile BPS (Direct Tail Synthesis):** Utilizing the direct quantile forecasts ($\boldsymbol{a}^{\text{quant}}_t$), QBPS generates the out-of-sample τ-th quantile prediction natively from the posterior predictive distribution. Detailed in section 2.4, the final point forecast is extracted by averaging the projected Asymmetric Laplace location parameter across

where the innovation $\varepsilon_t = \sigma_t z_t$ with $z_t \sim \mathcal{N}(0,1)$, and the conditional variance equation $\sigma^2_t = \omega + \alpha\varepsilon^2_{t-1} + \beta\sigma^2_{t-1}$.

the M effective MCMC draws: $\hat{y}_t^{(\tau,\text{QBPS})} = \frac{1}{M}\sum_{i=1}^{M} \boldsymbol{F}_t^{(i)\prime}\boldsymbol{\theta}_t^{\tau(i)}$.

Despite these disparate objective functions, both frameworks require the variance matrix A_t for sequential filtering. Because standard linear Quantile Regressions estimate a specific conditional quantile without assuming a full parametric distribution, they do not natively output a predictive density variance. Furthermore, while the GARCH model provides a dynamic variance, it characterizes the symmetric conditional distribution rather than a quantile-specific dispersion. As a practical approximation, this design utilizes the asymptotic predictive variance of the conditional mean to proxy the uncertainty of the quantile forecasts derived from the identical predictor set. Applying these variance proxies allows the synthesis engines to allocate dynamic weights based primarily on structural forecast accuracy rather than differences in variance scaling methodology:

$$A_{t,j}^{\text{quant}} = A_{t,j}^{\text{mean}} = \begin{cases} \widehat{\sigma}^2_{\text{GARCH},t|t-1}, & \text{for } j = 2, \\ \widehat{\text{Var}}\big(\widehat{\varepsilon}_{t,j}^{\text{OLS}}\big), & \text{for } j \in \{1,3,4\}, \end{cases} \tag{20}$$

where $\widehat{\text{Var}}(\widehat{\varepsilon}_{t,j}^{\text{OLS}})$ denotes the in-sample residual variance of the corresponding OLS mean regression evaluated on the agent's exact information set, and $\widehat{\sigma}^2_{\text{GARCH},t|t-1}$ is the 1-step-ahead conditional variance.

To initialize the dynamic synthesis engines for the out-of-sample evaluation ($t = 202$ to 400), prior hyperparameters calibrate baseline memory and state uncertainty. Following McAlinn and West (2019), the Gaussian BPS benchmark initializes with an equal-weight prior mean $m_0 = (0, 1/J, \ldots, 1/J)'$, a diffuse state covariance $C_0 = 1.0\mathbf{I}$, an initial scale $s_0 = 0.002$, and discount factors $(\delta, \beta) = (0.95, 0.99)$, which fixes the initial degrees of freedom to $n_0 = 1/(1-\beta)$. Both DRQS and QBPS share the same m_0, β, and n_0 with BPS. The Dynamic Regression Quantile Synthesis (DRQS) benchmark of Kobayashi et al. (2026) also maintains the diffuse covariance $C_0 = 1.0\mathbf{I}$, paired with their recommended static discount factor of $\delta = 0.85$ for the tail ($\tau = 0.05$) and $\delta = 0.95$ at the median ($\tau = 0.50$). In contrast, for tail synthesis ($\tau = 0.05$), QBPS initializes with a concentrated prior covariance $C_0 = 0.01\mathbf{I}$ and operates within an adaptive discount band $[\delta_{min}, \delta_{max}] = [0.95, 0.995]$. At

the median ($\tau = 0.50$), score-driven discounting is deactivated and QBPS fixes $\delta = 0.95$, matching the standard benchmark memory. Both quantile synthesis frameworks anchor their initial scale s_0 to the empirical check loss evaluated over the initial training window[8].

For above frameworks, sequential MCMC inference utilizes a strict baseline of 200 burn-in iterations and 1,000 post-burn-in draws per time step. Aggregated convergence diagnostics confirm high Effective Sample Sizes and robust parameter stationarity of QBPS for these simulated environments (see Appendix H, Panel A).

3.3 Symmetric Constraints and the "Gravity of the Mean"

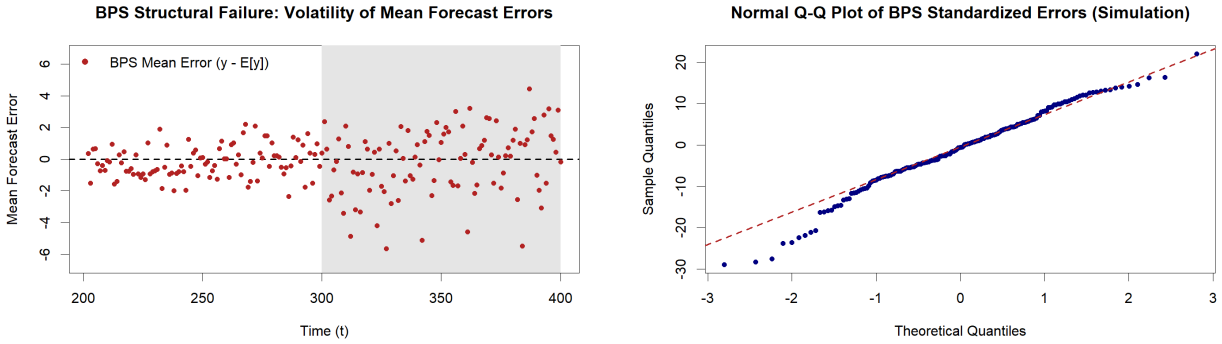
Standard Gaussian BPS maximizes a symmetric likelihood, inherently optimizing for the conditional mean. While effective during periods of symmetric volatility, this architecture becomes a constraint during asymmetric structural shifts. Bound by a symmetric loss function, the model suffers from a "gravity of the mean" that anchors the synthesized forecast to the center of the data mass, structurally preventing it from fully projecting into skewed tails. Consequently, as shown in Panel (a) of Figure 2, the BPS mean forecast residuals exhibit increased dispersion and pronounced downward deviations during the post-transition regime.

Panel (b) of Figure 2 isolates this distributional limit by mapping the standardized predictive residuals against theoretical standard normal quantiles. While perfectly Gaussian errors would align strictly along the diagonal, the empirical residuals deviate sharply downward at the lower extremes. This deviation mathematically confirms that the asymmetric left-tail realizations generated by the skewness shift fall outside the boundaries of the symmetric predictive distribution. This structural behavior establishes the necessity of transitioning from a symmetric Gaussian likelihood to an asymmetric framework when the explicit objective is

[8]Unlike standard Gaussian BPS, accurate variance initialization is critical for QBPS at extreme quantiles. Because the Asymmetric Laplace observation variance is scaled by $b_\tau = \frac{2}{\tau(1-\tau)}$, targeting $\tau = 0.05$ inflates this multiplier to approximately 42. If s_0 is arbitrarily initialized, this massive multiplier mathematically flattens the likelihood, collapsing the Kalman gain and artificially freezing the dynamic weights. Anchoring s_0 directly to the empirical ML estimator helps avoid this computational trap.

tail-risk evaluation.

Figure 2: Gaussian Benchmark (BPS) Mean Forecast Residuals



(a) BPS Forecast Residual

(b) Standardized Residual Q-Q Plot

Notes: The plotted residuals in (a) represent the forecast errors of the 1-quarter ahead BPS predictive mean: $y_t - \hat{y}_t^{BPS}$. (b) displays the normal Q-Q plots of BPS' 1-quarter ahead standardized mean predictive residuals' correspondence with the theoretical quantiles of a standard normal distribution.

3.4 QBPS Dynamics: Adaptation and the Intercept Mechanism

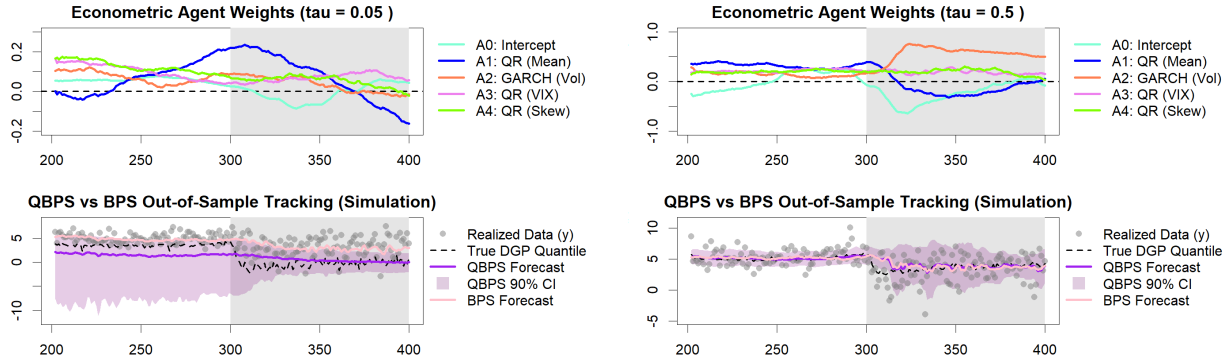
To rigorously evaluate the framework's capacity to navigate asymmetric macro-financial shocks, the primary simulation exercise targets the extreme left tail ($\tau = 0.05$). During the structural transition to asymmetric left-skewness, the dynamic posterior weights capture the shifting signal-to-noise ratio by progressively reallocating weight among the parametric experts (see Panel (a) of Figure 3).

Because the individual candidate agents are estimated over expanding historical windows, their predictions heavily reflect pre-transition data, which induces structural inertia during the abrupt shift. To address this slow adaptation, the unconstrained QBPS intercept acts as a structural shock absorber (Panel (a) of Figure 3). Rather than waiting for the parametric agents to fully update their boundaries, the intercept adjusts into negative territory, serving as a dynamic bias-corrector that pulls the synthesized forecast downward to track the shifting tail. Simultaneously, the filter rapidly reallocates the dynamic posterior weights, dynamically shifting mass toward the candidate agents that currently provide the most accurate predictive signal for the newly skewed regime. This two-pronged mobilization of the state vector is

governed by the observation-driven asymmetric discounting module, which actively drops the variance discount factor (Panel (a) of Figure 4) to inject the structural variance necessary to execute both the intercept correction and the weight reallocation.

Panel A of Table 1 summarizes out-of-sample performance at the 5th percentile ($\tau = 0.05$). While Gaussian BPS targets the conditional mean, projecting symmetric intervals yields an empirical hit rate of 46.73%, reflecting the pull of the distribution center. The DRQS framework of Kobayashi et al. (2026) targets quantile loss directly, improving upon mean combinations. However, under a constant discount factor ($\delta = 0.85$), it records a 10.05% hit rate and statistically significant higher check loss during the crisis regime (relative loss of 0.760, DM statistic of -2.61). In comparison, the proposed QBPS records a 4.52% hit rate, achieves the lowest check loss across both periods, and satisfies the Dynamic Quantile test.

Figure 3: QBPS: Dynamic Weighting and Tail Forecasts (Parametric Agents)



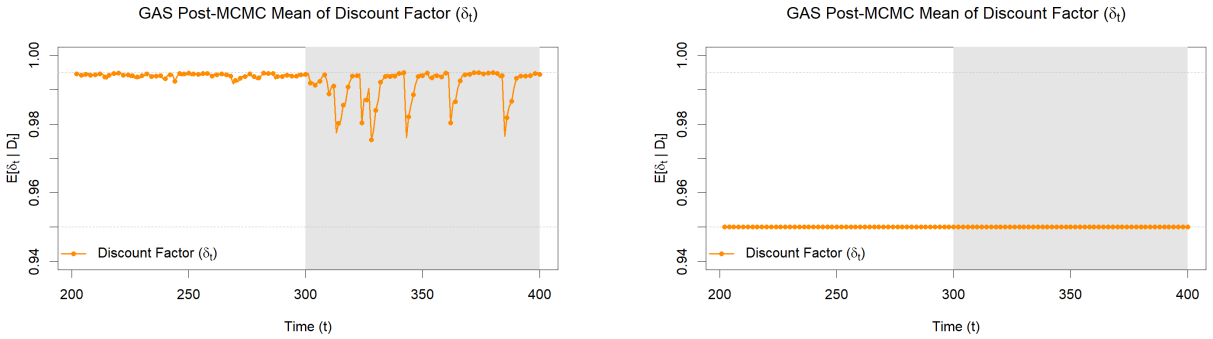
(a) $\tau = 0.05$ (b) $\tau = 0.5$

Notes: Target is the 1-period ahead 5th and 50th percentile ($\tau = 0.05$ and 0.5) of simulated times series. **Top Sub-panels:** Posterior weight evolution for the candidate parametric agents across the synthesis for $\tau = 0.05$ and 0.5. **Bottom Sub-panels:** Out-of-sample forecast tracking: the forecast of the Gaussian BPS benchmark (pink) and QBPS (purple), and the asymmetric QBPS 90% credible interval (shaded purple).

3.5 Performance at the Median

To evaluate the framework's baseline efficiency and ensure it does not overreact in stable environments, a secondary simulation exercise shifts the target quantile to the median ($\tau = 0.5$). For this symmetric evaluation, the asymmetric discounting module is deactivated, and the filter relies on a constant variance discount factor of $\delta = 0.95$ (Panel (b) of Figure

Figure 4: Score-Driven Dynamic Discounting (Parametric Agents)



(a) $\tau = 0.05$ (b) $\tau = 0.5$

Notes: The left panel displays the mean of the dynamic discount factor (δ_t) extracted from the left-tail QBPS predictive MCMC draws ($\tau = 0.05$) while the right panel displays the constant ($\delta_t = 0.95$) initialized for the median QBPS forecast ($\tau = 0.5$). The dynamics follows from the observation-driven asymmetric discounting module. The left panel illustrates the structural adaptation for the left-tail synthesis, while the right panel displays the constant discount factor. Crises periods are marked gray.

Table 1: Out-of-Sample Simulation Performance (Parametric Agents)

Model	**Relative Check Loss**		**Hit Rate (%)**	**DM Statistic**		**DQ Test**
	Overall	**Crisis**		**Overall**	**Crisis**	
Panel A: Target $\tau = 0.05$						
True DGP (Oracle)	1.333	1.191	5.03	3.92***	1.66*	✓
QBPS (Baseline)	**1.000**	**1.000**	**4.52**	–	–	✓
DRQS (Kobayashi et al.)	0.960	0.760	10.05	-0.42	-2.61***	×
BPS (Gaussian)	0.379	0.342	46.23	-5.93***	-5.70***	×
Agent: QR (Mean)	0.574	0.414	20.60	-1.97**	-3.19***	×
Agent: GARCH (Vol)	0.851	0.684	13.57	-1.32	-2.85***	×
Agent: QR (VIX)	0.744	0.642	27.64	-2.16**	-2.56**	×
Agent: QR (Skew)	0.799	0.696	23.62	-2.12**	-2.85***	×
Panel B: Target $\tau = 0.50$						
True DGP (Oracle)	1.120	1.146	49.75	3.27***	2.82***	✓
QBPS (Baseline)	**1.000**	**1.000**	**51.26**	–	–	✓
DRQS (Kobayashi et al.)	1.015	1.033	50.25	0.82	1.82*	×
BPS (Gaussian)	0.978	0.973	53.27	-1.25	-1.23	×
Agent: QR (Mean)	0.750	0.656	67.84	-3.97***	-6.18***	×
Agent: GARCH (Vol)	0.852	0.884	73.37	-5.16***	-4.42***	×
Agent: QR (VIX)	0.771	0.800	80.90	-4.95***	-3.37***	×
Agent: QR (Skew)	0.769	0.800	80.40	-5.01***	-3.36***	×

Notes: Relative losses are expressed as the ratio of the QBPS baseline check loss to the competitor's out-of-sample mean check loss ($Loss_{QBPS}/Loss_{Model}$). Values < 1 indicate that the competitor underperforms the QBPS baseline. Diebold-Mariano (DM) statistics are calculated using Newey-West HAC standard errors; negative statistics indicate that the competitor yields significantly worse (higher) predictive loss than QBPS. The True DGP (Oracle) is the only benchmark mathematically capable of outperforming the synthesis (> 1 relative loss, positive DM stat). Asterisks denote the statistical significance of the DM test: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The Dynamic Quantile (DQ) Test evaluates the correct specification of the out-of-sample hit sequence using the Engle-Manganelli joint Wald test (✓ denotes a passing p-value > 0.05, indicating the sequence is unpredictable; × denotes failure).

4). Importantly, this evaluation at the median strictly utilizes the check loss—while MSE inherently elicits the conditional mean, the check loss remains the strictly consistent scoring function for evaluating conditional median forecasts (Koenker and Bassett Jr, 1978; Gneiting, 2011).

When targeting the median ($\tau = 0.5$), the Asymmetric Laplace likelihood naturally collapses into a symmetric form. As reported in Panel B of Table 1, the predictive accuracy of QBPS aligns closely with both standard Gaussian BPS and DRQS. All three synthesis frameworks generate empirical hit rates centered near 50% and exhibit comparable check-loss metrics. This convergence confirms that the scaled hierarchy and score-driven discounting mechanisms introduced to stabilize extreme tail estimation operate cleanly without degrading central predictive accuracy.

4 Empirical Application: Synthesizing Growth-at-Risk

4.1 The Economic Stakes and Candidate Agents

Growth-at-Risk (GaR) has emerged as the modern standard for evaluating macroeconomic vulnerability, shifting the focus of policymakers away from average fluctuations and toward the extreme left tail of the growth distribution as business cycle skewness fluctuates (Adrian et al., 2019; Salgado et al., 2019; Delle Monache et al., 2024). While monitoring the 5th percentile of GDP growth is now central to institutional risk assessment, standard rolling-window econometric frameworks are designed to smooth over historical data. While this equal weighting of past observations provides stable estimates during normal expansions, it limits the model's speed of adaptation, causing it to react slowly to rapid regime shifts and sudden market crashes.

In the following quasi-real-time downside risk forecasting exercise, the primary focus is the 1-quarter ahead horizon to capture immediate macro-financial vulnerability. A supplementary 4-quarter ahead evaluation then tests long-horizon stability and the structural decay

of predictive signals. Across both horizons, the target variable is annualized US Real GDP Growth[9], and the primary predictors feature the Chicago Fed National Financial Conditions Index (NFCI) alongside its distinct subcomponents: Risk, Credit, and Leverage[10]. The dataset spans from 1971Q1 to 2026Q2. Following Amburgey and McCracken (2023), forecast origins are aligned to the second month of the next quarter to match the actual release schedule of advance GDP estimates (e.g., utilizing Q1 data on May 1 to forecast Q2). To prevent look-ahead bias, financial conditions are aggregated via a backward-looking 13-week rolling average. The model then extracts the specific NFCI value from the averaged observation nearest to the forecast origin. The candidate agents are estimated via an expanding historical window initialized with 80 quarters of data. The synthesis engines then evaluate and dynamically weight these agent forecasts utilizing a 40-quarter rolling window, yielding a final out-of-sample evaluation period from 2001Q2 through 2026Q2.

Forecasting extreme macroeconomic tail risk during major disruptions highlights a practical modeling trade-off. Standard linear models assume constant marginal relationships; lacking structural boundaries, an unprecedented anomaly like the 2020 COVID shock forces these models to mechanically over-extrapolate into unrealistically deep negative territory. Conversely, tree-based architectures like Quantile Random Forests (QRFs) (Meinshausen, 2006) can discover complex, non-linear predictive interactions while providing convenient "bounded stability" against such explosive overshooting (Clark et al., 2023; Chernis et al., 2025). However, this bounding mathematically floors predictions at the edge of historical training support, causing standalone QRFs to under-react to unprecedented tail realizations.

To navigate this trade-off, the empirical design pairs the bounded, non-linear predictive capacity of QRF agents with the dynamic linear interpretability of the QBPS framework. By mapping complex non-linear forecasts into a time-varying linear combination, the synthesis enables structural evaluation of the underlying non-linear signals. I deploy these QRF agents across two distinct synthesis exercises. The first evaluates aggregate financial conditions by synthesizing a purely macroeconomic baseline agent (trained exclusively on lagged GDP)

[9]Federal Reserve Economic Data (FRED): `https://fred.stlouisfed.org/series/A191RL1Q225SBEA`.
[10]Federal Reserve Bank of Chicago: `https://www.chicagofed.org/research/data/nfci/current-data`.

with a macro-financial agent augmented by the aggregate NFCI. The second decomposes this vulnerability by synthesizing the baseline agent alongside three distinct experts augmented with the Risk, Credit, and Leverage subindices. Depending on the specific synthesis engine, standard Gaussian BPS synthesizes the conditional mean of the QRF distributions, while QBPS directly processes their 5th percentile predictions. Echoing the variance design of the earlier simulation, agent j's forecast uncertainty $A_{j,t}$ is approximated via the interquartile range implied by its empirical QRF distribution, signifying the confidence of agent predictions:

$$A_{j,t} = \left(\frac{Q_{0.75,t} - Q_{0.25,t}}{1.349} \right)^2 \tag{21}$$

where $Q_{0.75,t}$ and $Q_{0.25,t}$ denote the 75th and 25th predictive quantiles of the agent's ensemble, respectively[11].

To mathematically compensate for the QRFs' historical boundaries during crises, QBPS structurally adapts via a bifurcated prior state covariance initialization ($\mathbf{C}_0$). The dynamic weights of the candidate agents are tightly anchored at 0.01 to stabilize the Kalman gain against extreme tail noise. In contrast, the unconstrained baseline intercept receives a significantly looser prior variance of 0.5. This experimental asymmetry designates the intercept as a rapid shock absorber: it remains free to plunge and capture deep level deficits that fall entirely outside the agents' historical bounds. As the intercept accommodates this structural gap, the filter smoothly reallocates dynamic weights toward the non-linear signals providing the highest marginal predictive value[12]. To ensure a fair baseline comparison in a highly volatile empirical environment, the initial scale s_0 for the Gaussian BPS benchmark is increased from 0.002 to 0.01, granting it a more flexible variance base. Beyond these targeted adjustments to prior initialization, all remaining hyperparameter specifications match the simulation settings—including the data-driven rule anchoring the QBPS initial scale s_0 to the empirical check-loss of the training window. The empirical framework also maintains

[11] For a standard Gaussian distribution, the interquartile range (IQR) evaluates to approximately 1.349σ. Dividing the empirical IQR by 1.349 yields a robust, non-parametric variance proxy that accurately scales predictive uncertainty while resisting extreme tree ensemble outliers.

[12] Initial evaluations utilizing parametric agents demonstrated severe unbounded extrapolation during the 2020 COVID shock. Consequently, the candidate pool is restricted to non-parametric learners. While QRFs serve as the primary representative ensemble, alternative bounded architectures (e.g., Gradient Boosted Machines) yield qualitatively similar baseline shapes and synthesis dynamics.

the baseline configuration of 200 burn-in and 1,000 post-burn-in MCMC iterations. Continuous out-of-sample diagnostics confirm stable parameter convergence across the recursive evaluation window, indicating that the baseline algorithm remains robust across shifting macro-financial regimes without requiring dataset-specific tuning (see Appendix H, Panels B and C).

To evaluate this dynamic synthesis against standard GaR methodologies, the out-of-sample performance of QBPS is also compared against the foundational two-step parametric benchmark of Adrian et al. (2019) (hereafter ABG Skew-t). This framework first estimates linear quantile regressions—incorporating the aggregate NFCI alongside lagged GDP—and subsequently fits a conditional Skew-t distribution to the predicted quantiles via numerical optimization. The procedure features an expanding estimation window, keeping the out-of-sample strictly matched with QBPS (2001Q1 to 2026Q2).

4.2 The Symmetric Constraints of Gaussian Synthesis

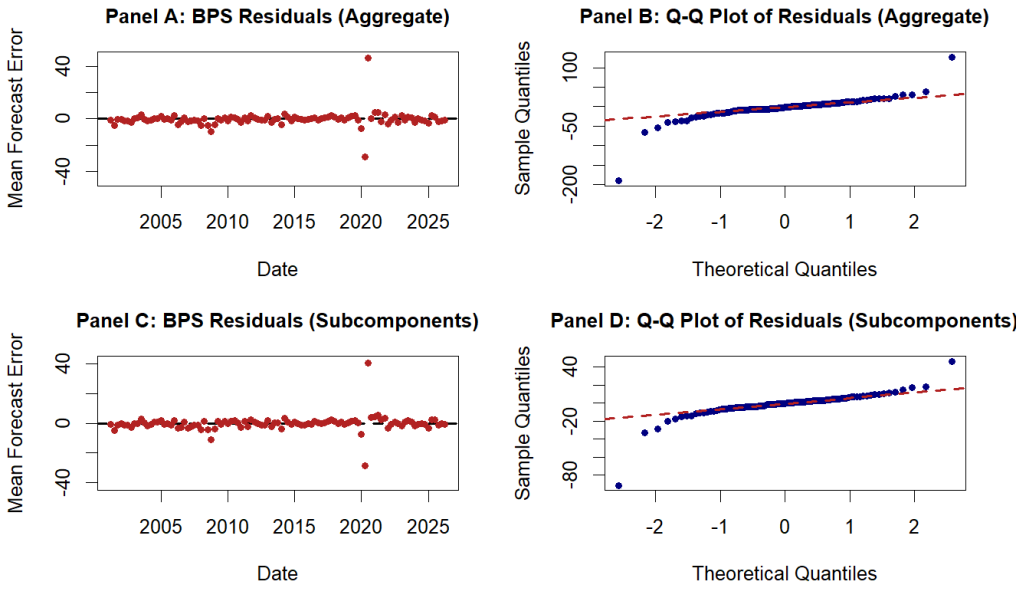
During tranquil, well-behaved economic periods, Standard Gaussian BPS' symmetric penalty is a highly desirable feature that provides baseline stability and prevents the forecast from overreacting to minor noise. However, during a highly skewed macro-financial crisis, this exact stabilizing feature acts as a rigid constraint.

This mechanism is clearly visible in the BPS out-of-sample forecast errors. Figure 5 plots the 1-quarter ahead predictive mean residuals and their corresponding Q-Q plots for the Gaussian BPS benchmark. In Panels A and C, the residuals exhibit large, abrupt deviations from zero specifically during the 2008 Great Recession and the 2020 COVID shock. This behavior reflects the natural consequence of the model remaining securely anchored to the mean while the realized economic distribution sharply shifts downward.

The Q-Q plots in Panels B and D further illuminate this structural constraint by mapping the sample quantiles of the residuals against the theoretical quantiles of a standard normal distribution. While perfectly Gaussian errors would align strictly along the diagonal, the

plotted residuals deviate sharply downward at the lower extremes. Across both the aggregate and subcomponent synthesis exercises, this deviation confirms a pronounced, heavy left tail compared to the right tail, aligning with established empirical evidence regarding the asymmetric, downside-skewed behavior of GDP growth. Applying a perfectly symmetric likelihood to an asymmetric reality naturally leaves extreme tail realizations outside the modeled distribution, rationalizing the transition to the Asymmetric Laplace forecast synthesis framework for institutions explicitly mandated to evaluate and manage downside risk.

Figure 5: Gaussian Benchmark (BPS) During Structural Breaks
1-Quarter Ahead



Notes: **Panels A & C:** The plotted residuals represent the forecast errors of the 1-quarter ahead BPS predictive mean: $y_t - \hat{y}_t^{BPS}$. **Panels B & D:** The normal Q-Q plots display BPS' 1-quarter ahead standardized mean predictive residuals' correspondence with the theoretical quantiles of a standard normal distribution.

4.3 Out-of-Sample Tail Forecasting Performance

Table 2 evaluates the out-of-sample 1-quarter-ahead downside risk forecasts across overall check loss, recession-specific accuracy, empirical hit rates and quantile specification tests. By leveraging a granular financial information set, the subcomponent synthesis dynamically captures the shifting risks of Credit, Leverage, and Volatility, yielding the most competitive

Table 2: 1-Quarter Ahead Downside Risk Evaluation ($\tau = 0.05$)

Model	Relative Check Loss		Hit Rate (%)	DM Statistic		DQ Test
	Overall	Recession		Overall	Recession	
I. Universal Benchmarks						
Empirical (Unconditional)	0.835	0.536	6.93	-0.93	-2.44**	×
Agent: Baseline (Lagged GDP)	0.946	0.827	7.92	-0.92	-1.96*	✓
ABG Skew-t	0.902	0.726	9.90	-0.67	-1.75*	✓
II. Application 1: Aggregate NFCI						
QBPS (Aggregate)	**0.976**	**0.980**	**3.96**	**-1.68***	**-0.71**	✓
BPS (Aggregate)	0.479	0.399	45.54	-2.88***	-3.55***	×
Agent: Aggregate NFCI	1.000	0.771	9.90	0.00	-1.74*	✓
III. Application 2: NFCI Subcomponents						
QBPS (Subcomponents)	**1.000**	**1.000**	**4.95**	–	–	✓
BPS (Subcomponents)	0.500	0.412	37.62	-2.70***	-3.48***	×
Agent: Risk	0.970	0.750	10.89	-0.30	-2.05**	✓
Agent: Credit	0.968	0.768	8.91	-0.35	-2.06**	✓
Agent: Leverage	0.988	0.831	9.90	-0.13	-1.31	✓

Notes: Relative losses are standardized against the QBPS (Subcomponents) model; values < 1 indicate higher loss (worse performance) than the global benchmark. Underlying base agents are stabilized via an ensemble average of 50 independent Quantile Random Forest (QRF) repetitions (number of trees = 500, node size = 40). Hit rate indicates the out-of-sample empirical percentage of model forecasts below the true series. Diebold-Mariano (DM) statistics are calculated using Newey-West HAC standard errors; negative statistics indicate underperformance relative to the benchmark. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The Dynamic Quantile (DQ) Test evaluates correct specification of the hit sequence at the 5% significance level, where (✓) denotes a passing p-value > 0.05 and (×) denotes a failure. Raw statistics can be found in Appendix E.

overall predictive accuracy. Consequently, QBPS (Subcomponents) is selected as the global benchmark against which the Diebold-Mariano tests evaluate all alternative frameworks.

The hit rate measures the out-of-sample percentage of times realized GDP growth falls below the forecasted 5th percentile. While the Gaussian BPS benchmark is structurally pulled toward the median (yielding hit rates near 40%), the parametric ABG Skew-t benchmark and the standalone QRF agents yield 9.90%. For the QRF agents, this higher hit rate naturally reflects the historical bounds of tree-based models when encountering unprecedented shocks. For the ABG framework, balancing the potential overextrapolation of underlying linear regressions with the shape constraints of a parametric distribution can result in limited flexibility when tracking extreme empirical tails. In contrast, the QBPS (Subcomponents) model leverages its dynamic intercept to navigate these structural trade-offs, achieving a 4.95% hit rate that closely aligns with the nominal target.

In terms of asymmetric check loss, the candidate agents and the ABG Skew-t framework provide highly competitive baseline forecasts during normal expansions. However, QBPS

yields lower predictive loss than the empirical baseline, the individual QRF experts, and the parametric ABG framework during macroeconomic recessions. Because expanding-window frameworks inherently adapt gradually to sudden crashes, the Diebold-Mariano (DM) statistics (Diebold and Mariano, 2002) confirm that QBPS offers a statistically significant improvement in predictive accuracy during these macro-financial stress periods.

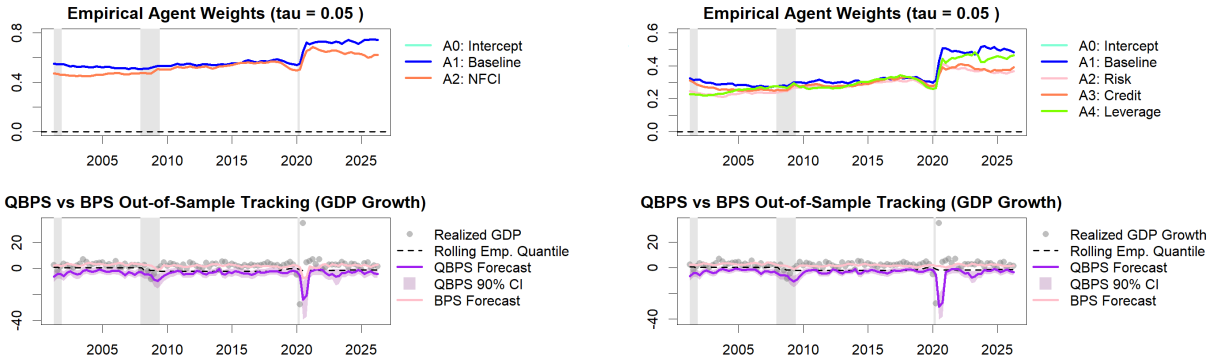
Finally, the Engle-Manganelli Dynamic Quantile (DQ) test (Engle and Manganelli, 2004) evaluates the correct dynamic specification of the hit sequence. While the unconditional empirical quantile fails the test due to autocorrelated misses, the ABG Skew-t, QRF, and QBPS models all pass, confirming that their errors do not strictly cluster chronologically. However, while the ABG and QRF frameworks are dynamically well-specified, QBPS pairs this dynamic specification with a hit rate calibrated to the nominal 5% target and a lower asymmetric check loss during crises. Full unscaled original results of this evaluation are provided in Appendix E. To ensure that this out-of-sample performance is structurally driven by the synthesis engine rather than specific ensemble or sampler configurations, I conduct computational robustness evaluations utilizing alternative base-agent capacities (expanding the QRF ensemble to 100 independent repetitions) and extended MCMC sampling chains (2,000 iterations). Across these variations, the framework's predictive accuracy, hit rates, and dynamic specification remain both quantitatively and qualitatively stable. Full tabular results are provided in Appendix F.

4.4 The Mechanics of Crisis Adaptation

The outperformance of the QBPS framework stems from its dynamic weighting mechanism and its capacity to adjust for the boundary limits of its constituent agents. Figure 6 illustrates this dynamic weight evolution at the 1-quarter ahead horizon. In the subcomponent synthesis displayed in Panel (b), the model continuously reallocates posterior weight among the Risk, Credit, and Leverage experts, providing a shifting consensus as business cycle conditions evolve.

Given that the tree-based QRF agents are mathematically floored by their expanding historical training windows, they are structurally prevented from predicting crashes deeper than those previously observed in the sample. As the candidate agents hit this historical floor during the unprecedented 2020 COVID shock, the filter deploys a two-pronged structural adaptation. First, the unconstrained QBPS intercept actively diagnoses this collective boundary, plunging into negative territory to act as a dynamic bias-corrector that pulls the synthesized forecast's level downward. Second, following a brief initial dip, the dynamic weights allocated to the QRF agents hike to a permanently elevated level. Because the agents' signals correctly identify the direction of the crisis but are bounded in magnitude, the synthesis engine scales up these multiplier weights to mathematically amplify the constrained signals, aligning them to match the true depth and variance of the new economic regime. Full-scale plots detailing the complete intercept trajectories alongside the synthesis weights are available in Appendix E.

Figure 6: QBPS: Dynamic Weighting and Tail Forecasts (1-Quarter Ahead)



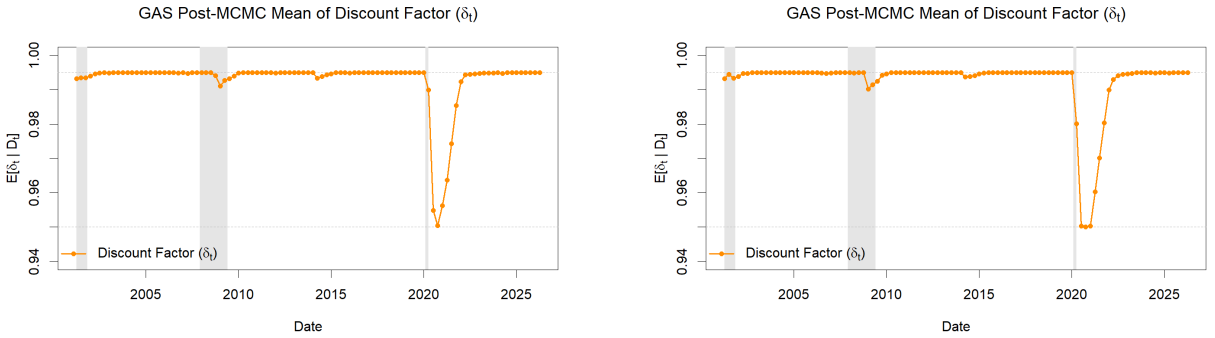
(a) Application 1: Aggregate NFCI

(b) Application 2: NFCI Subcomponents

Notes: Target is the 1-quarter ahead 5th percentile ($\tau = 0.05$) of GDP growth. **Top Sub-panels:** Posterior weight evolution for the candidate QRF agents across the Aggregate NFCI synthesis and the Subcomponent synthesis. Full-scale plots with intercepts are in Appendix E. **Bottom Sub-panels:** Out-of-sample forecast tracking. The Gaussian BPS benchmark (pink) optimizing for MSE is structurally pulled toward the mean, systematically failing to cover severe downside crashes. In contrast, QBPS (purple) correctly tracks the empirical tail, with the BPS forecast frequently falling completely outside the asymmetric QBPS 90% credible interval (shaded purple).

At the same time, this rapid parameter adaptation is governed by the asymmetric discounting module, illustrated in Figure 7. The dynamic discount factor drops at the onset of NBER recessions to inject necessary structural variance into the Kalman filter. Notably, the framework distinguishes between the magnitudes of crises: while the discount factor exhibits a

Figure 7: Score-Driven Dynamic Discounting (1-Quarter Ahead)



(a) Application 1: Aggregate NFCI

(b) Application 2: NFCI Subcomponents

Notes: Both panels display the mean of the dynamic discount factor (δ_t) extracted post-MCMC as the mean of effective MCMC draws. The dynamics follows from the observation-driven asymmetric discounting module. The left panel illustrates the structural adaptation for the aggregate NFCI synthesis, while the right panel displays the subcomponent synthesis. NBER recessions are marked gray.

distinct but moderate drop during the 2008 Great Recession, it plunges dramatically during the unprecedented 2020 COVID shock, reflecting the massive scale of the variance injection required to mobilize the state vector and execute synthesis weight adjustment.

4.5 Real-Time Attribution: Decomposing Macro-Financial Vulnerability

While the dynamic weights indicate which candidate forecast exerts the greatest influence on the synthesis, they operate as relative predictive allocations rather than absolute macroeconomic magnitudes. To isolate the additive predictive value of individual financial channels, we can leverage the shared macroeconomic anchor across the candidate agents. Recall that each QRF agent forecast ($\hat{y}_{j,t}^{(\tau)}$) is constructed as:

$$\hat{y}_{1,t}^{(\tau)} = \text{QRF}_{\text{Baseline}}\left(y_{t-1}\right), \ \hat{y}_{j,t}^{(\tau)} = \text{QRF}_{j}\left(y_{t-1}, x_{j,t-1}\right) \quad \text{for } j \in \{\text{Risk}, \text{Credit}, \text{Leverage}\} \tag{22}$$

where y_{t-1} is lagged GDP growth, $x_{j,t-1}$ denotes the respective NFCI subindex, and $\tau = 0.05$ targets Growth-at-Risk.

Because this baseline macroeconomic signal is common across all models, the synthesized tail quantile under QBPS can be algebraically decomposed into macroeconomic momentum and isolated financial components:

$$\hat{y}_t^{(\tau)} = \theta_{0,t}^{(\tau)} + \left(\sum_{k=1}^{4} \theta_{k,t}^{(\tau)} \right) \hat{y}_{1,t}^{(\tau)} + \sum_{j \in \{\text{Risk,Credit,Leverage}\}} \mathcal{M}_{j,t}^{(\tau)}, \tag{23}$$

where the dynamic marginal tail-risk contribution $\mathcal{M}_{j,t}^{(\tau)}$ of financial channel j is defined as:

$$\mathcal{M}_{j,t}^{(\tau)} \equiv \theta_{j,t}^{(\tau)} \left(\hat{y}_{j,t}^{(\tau)} - \hat{y}_{1,t}^{(\tau)} \right). \tag{24}$$

Here, $\mathcal{M}_{j,t}^{(\tau)}$ isolates the agent's marginal additive contribution to the 5th percentile forecast, expressed in exact percentage points relative to the macroeconomic baseline. When $\mathcal{M}_{j,t}^{(\tau)} < 0$, the financial agent projects a steeper contraction than macroeconomic inertia alone, quantifying the additive magnitude by which financial frictions signal elevated downside risk. Conversely, when $\mathcal{M}_{j,t}^{(\tau)} > 0$, benign financial signals elevate the tail forecast above the baseline trajectory, indicating a reduced exposure to downside tail-events.

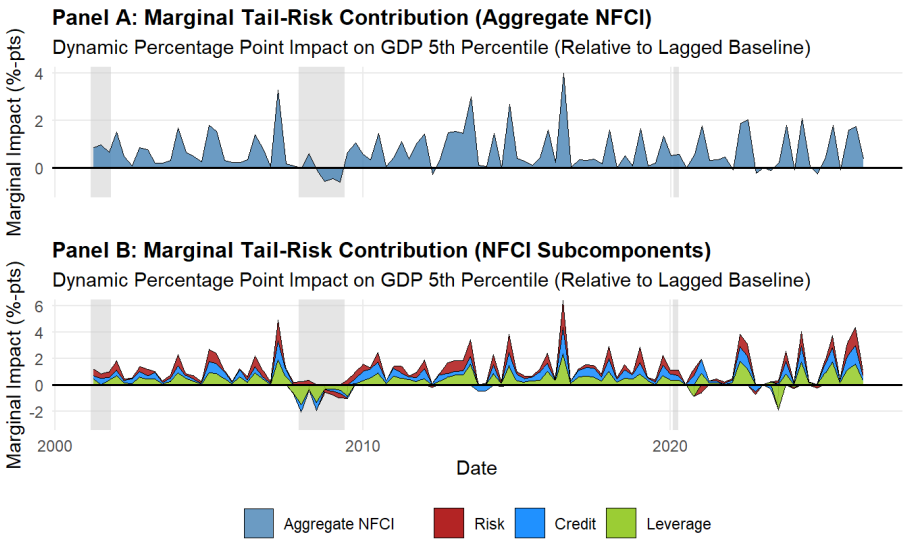
Figure 8 contrasts the marginal tail-risk contribution of the aggregate NFCI[13] (Panel A) against that of its subcomponents (Panel B). Tracking these dynamic contributions reveals a distinct structural asymmetry in how specific financial indicators forecast real economic vulnerability across the business cycle.

4.5.1 Proportional Cushioning During Economic Expansions

During tranquil economic periods and prolonged expansions, the marginal predictive contributions ($\mathcal{M}_{j,t}^{(\tau)}$) consistently operate above zero. In these regimes, loose financial conditions systematically signal reduced downside macroeconomic vulnerability, projecting a GaR trajectory safely above the autoregressive baseline. This synchronous signaling aligns with the

[13]For the aggregate financial conditions exercise, this marginal contribution is derived using the exact same algebraic decomposition applied to the two-agent synthesis: $\mathcal{M}_{\text{NFCI},t}^{(\tau)} \equiv \theta_{\text{NFCI},t}^{(\tau)} (\hat{y}_{\text{NFCI},t}^{(\tau)} - \hat{y}_{1,t}^{(\tau)})$, where $\hat{y}_{\text{NFCI},t}^{(\tau)}$ represents the forecast of the QRF agent using lagged GDP and aggregate NFCI index.

Figure 8: Dynamic Growth-at-Risk (GaR) Marginal Attribution
1-Quarter Ahead



Notes: Both panels display the percentage point contribution to 5th percentile ($\tau = 0.05$) GDP growth relative to a purely autoregressive lagged-GDP baseline ($\mathcal{M}_{j,t}^{(0.05)}$). **Panel A** displays the marginal downside impact evaluated using the aggregate Chicago Fed NFCI. **Panel B** utilizes the QBPS subcomponent synthesis to dynamically decompose the index, isolating the distinct, real-time tail-risk contributions of Risk, Credit, and Leverage. NBER recessions are shaded in gray.

procyclicality of the financial system documented by Adrian and Shin (2010), who note that during low-volatility regimes, intermediaries symmetrically expand balance sheet leverage and credit supply. As formally demonstrated by Adrian et al. (2019), this collective easing strongly predicts a compression of the left tail of GDP growth, indicating a temporary reduction in downside exposure. As observed in Panel B throughout the mid-2000s and the post-2010 expansion, this positive predictive contribution is highly proportional. Favorable risk appetite, credit availability, and leverage expansion move synchronously, providing relatively balanced, symmetric signals that collectively elevate the overall tail-risk forecast.

4.5.2 The Asymmetric Signal of Leverage in Amplifying Crises

During structural downturns and market shocks, this proportionality fractures. As the economy transitions into a financially stressed period, the marginal contributions sharply invert below zero. However, rather than all three subcomponents degrading equally, the decom-

position in Panel B of Figure 8 reveals that the *Leverage* channel systematically provides the primary predictive signal for elevated downside risk during episodes of financial stress. During the acute phase of the 2008 Great Recession, the Leverage component serves as the leading indicator of the contraction's severity, projecting a powerful downward revision to the GDP tail, while the Risk and Credit channels exhibit comparatively muted predictive signals.

This diagnostic capability is equally evident during the 2022 global bond market crash. Driven by rapid monetary tightening, fixed-income intermediaries and duration-heavy portfolios faced severe margin pressures. Panel B of Figure 8 captures this friction: the Leverage subcomponent plunges into negative territory, serving as the primary driver for the downward revision of the GaR forecast, even as the Risk and Credit channels remain relatively benign.

Importantly, the dynamic synthesis distinguishes these endogenous market stresses from exogenous real shocks. During the 2020 COVID-19 pandemic—a severe but non-financially driven macroeconomic contraction—the aggregate GaR forecast plunged, yet the Leverage subcomponent did not exhibit the extreme negative predictive contribution seen in 2008 or 2022. This empirical divergence confirms that the NFCI subcomponent synthesis isolates specifically financially driven downside risk, rather than mechanically extracting negative predictive signals during every economic contraction.

These empirical findings structurally align with the macro-finance literature regarding intermediary asset pricing and financial frictions. Theoretical models by Adrian and Shin (2010) and He and Krishnamurthy (2013) demonstrate how intermediary leverage constraints bind exclusively during downturns, forcing deleveraging spirals and fire sales that non-linearly amplify shocks. Furthermore, Brunnermeier and Sannikov (2014) explicitly show how this leverage generates endogenous, non-linear macroeconomic tail risk, a dynamic empirically connected to Growth-at-Risk by Adrian et al. (2019). Building on this foundation, recent empirical work by Amburgey and McCracken (2026) corroborates these dynamics, finding that leverage serves as the leading component driving the predictive ability of the NFCI

over extreme macroeconomic tail risk. While the theoretical literature identifies leverage as a structural causal mechanism during liquidity shocks, the QBPS framework paired with NFCI subcomponent synthesis provides the corresponding empirical forecasting evidence: financially driven macroeconomic tail events are most strongly signaled not by broad financial deterioration, but specifically by intermediary leverage.

4.6 Long-Horizon Forecasting and the Post-Crisis Recovery

To evaluate long-horizon stability, candidate agents construct direct $H = 4$ quarter-ahead forecasts, which the framework then dynamically synthesizes using the k-step ahead projection detailed in Section 2.4. Beyond the extended horizon, the empirical design remains strictly identical to the 1-quarter ahead evaluation, utilizing the same QRF ensembles across both the Aggregate NFCI and NFCI Subcomponent applications. To ensure the final out-of-sample evaluation covers the same period, the initial agent expanding window is reduced by four quarters, offsetting the target shift required for direct forecasting. Full out-of-sample results and figures for this extended evaluation are provided in Appendix G.

The ABG Skew-t model remains a Universal Benchmark for this evaluation. As documented by Adrian et al. (2019), the predictive signal of current financial variables attenuates at longer horizons. As the underlying linear conditional relationships weaken, the ABG parametric fit converges toward the unconditional historical distribution. This convergence provides a highly robust baseline for long-horizon forecasting.

The Symmetric Constraints of Gaussian Synthesis

Figure 10 illustrates a distinct shift in the residual distribution of the Gaussian BPS benchmark. While the heavy left tail associated with unpredicted sudden contractions persists from the 1-quarter horizon, a pronounced heavy right tail now emerges. At 4 quarters ahead, the initial shock is already embedded within the agents' lagged information sets, shifting the primary forecasting challenge from predicting sudden economic contractions to

tracking rapid macroeconomic rebounds. Because symmetric BPS models generally project a gradual reversion to the mean, the framework tends to under-predict the magnitude of the subsequent recovery. This structural characteristic generates large positive forecast errors, adding a heavy right tail to the residual distribution.

Out-of-Sample Tail Forecasting Performance

Tables 5 and 6 detail the out-of-sample performance. While the ABG benchmark yields the lowest overall check loss at this extended horizon, the performance difference between QBPS (Subcomponents) and ABG (overall and during crisis regimes) is statistically insignificant. In contrast, standard symmetric models (BPS) and standalone baseline agents exhibit statistically significant underperformance during these critical periods. Overall, the QBPS framework remains highly competitive: it achieves a 4.95% empirical hit rate—closely matching the 5% nominal target—and satisfies the DQ test for dynamic specification, confirming its structural reliability at this extended forecast horizon.

Estimation Inertia and Crisis Adaptation

The weight dynamics in Figure 11 highlight the extrapolation challenge posed by the 2020 COVID shock. Observing the severe contraction in the second quarter of 2020, the tree-based QRF agents—which are inherently constrained by the boundaries of historical training data—extrapolate a prolonged downturn four quarters ahead, persistently under-predicting the subsequent rapid expansion. To compensate for this estimation inertia, the unconstrained QBPS intercept provides a critical structural correction. In direct contrast to its behavior at $H = 1$—where it adjusted downward to capture the initial contraction—the intercept at $H = 4$ shifts sharply positive. This dynamic upward bias-correction bridges the gap between the agents' constrained downside predictions and the realized macroeconomic rebound.

Signal Attenuation and Macroeconomic Inertia

Finally, Figure 13 reveals a distinct shift in the marginal predictive contributions of Risk, Credit, and Leverage. This contrasts with the 1-quarter horizon, where financial conditions provided strong leading signals for imminent contractions. At a 4-quarter horizon, the negative predictive contributions of these variables attenuate significantly during the post-2020 recovery period. Conversely, the positive marginal impacts (the spikes above zero) are objectively larger than at $H = 1$. This pattern indicates that these financial condition indicators offer limited predictive content for Growth-at-Risk at extended horizons. This aligns with the findings of Adrian et al. (2019) that the predictive signal of financial conditions for downside risk attenuates as the forecasting horizon lengthens. Consequently, the observation-driven QBPS filter dynamically accommodates this signal decay by down-weighting the financial subcomponents, relying instead on the unconstrained intercept to correct the baseline estimation inertia.

5 Conclusion

While central banks and regulatory authorities actively monitor and forecast downside macro-financial risks, standard forecast combination methods typically optimize for average conditions rather than distribution tails. In particular, existing Gaussian synthesis frameworks focus on the conditional mean, leaving their tail forecasts vulnerable to the "gravity of the mean"—an anchoring that can bypass early warning signals and limit a model's capacity to track abrupt economic transitions.

This paper proposes Quantile Bayesian Predictive Synthesis (QBPS) to escape this gravitational pull. The framework incorporates two core methodological developments. First, a scaled parameterization of the Asymmetric Laplace likelihood separates quantile shocks from baseline volatility. This formulation prevents observation noise from collapsing, ensuring that combination weights do not overreact to single observations or discard historical

information. An exact Generalized Inverse Gaussian (GIG) conditional update further ensures closed-form posterior simulation. Second, an observation-driven Generalized Autoregressive Score (GAS) mechanism monitors tail forecast errors to regulate parameter memory. This mechanism accelerates weight adaptation during sudden economic shifts while preserving parameter stability during tranquil periods. It achieves this responsiveness by tracking unexpected tail surprises through the score of the Asymmetric Laplace likelihood.

Simulation experiments evaluate QBPS in a controlled environment featuring a business cycle interrupted by an abrupt shift toward downside risk. As an operational "plug-in" framework, QBPS directly combines quantile forecasts from diverse models alongside simple measures of their predictive uncertainty. When forecasting the 5th percentile ($\tau = 0.05$) during the transition to a turbulent regime, QBPS reallocates combination weights quickly and reduces check loss relative to standalone candidate models and standard Gaussian BPS. The method also performs reliably at the distribution center ($\tau = 0.5$), matching the predictive accuracy of standard BPS at the median.

Applied to quasi-real-time US Growth-at-Risk forecasting across 1- and 4-quarter horizons, QBPS translates these theoretical and simulation gains into tangible empirical performance. When forecasting the extreme 5th percentile of GDP growth, the synthesis yields competitive out-of-sample check loss overall and outperforms standalone tree models, standard Gaussian BPS, and the foundational two-step Growth-at-Risk benchmark during recessions. Beyond predictive accuracy, translating dynamic weights into forecast contributions provides real-time risk attribution: at short horizons, the decomposition highlights intermediary leverage as the primary leading indicator of downside risk, while at extended horizons, it captures the diminishing predictive contribution of financial indicators as their signals diminish.

Ultimately, while standard synthesis frameworks remain well suited for forecasting the conditional mean, evaluating macro-financial vulnerability benefits from complementary tools explicitly designed for distribution tails. By combining computational stability with real-time risk attribution, QBPS offers a practical approach for tracking both the severity and the shifting composition of downside risk alongside baseline forecasts.

References

Adrian, M. T., D. Giannone, M. Luciani, and M. West (2025). *Scenario synthesis and macroeconomic risk*. International Monetary Fund.

Adrian, T., N. Boyarchenko, and D. Giannone (2019). Vulnerable growth. *American Economic Review 109*(4), 1263–1289.

Adrian, T. and H. S. Shin (2010). Liquidity and leverage. *Journal of financial intermediation 19*(3), 418–437.

Amburgey, A. J. and M. W. McCracken (2023). On the real-time predictive content of financial condition indices for growth. *Journal of Applied Econometrics 38*(2), 137–163.

Amburgey, A. J. and M. W. McCracken (2026). Growth-at-risk is investment-at-risk. *Review of Economics and Statistics*, 1–27.

Bernardi, M., R. Casarin, B. B. Maillet, and L. Petrella (2026). Bayesian dynamic quantile model averaging. *Annals of Operations Research 363*(2), 1355–1383.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of econometrics 31*(3), 307–327.

Brownlees, C. and A. B. Souza (2021). Backtesting global growth-at-risk. *Journal of Monetary Economics 118*, 312–330.

Brunnermeier, M. K. and Y. Sannikov (2014). A macroeconomic model with a financial sector. *American Economic Review 104*(2), 379–421.

Bucy, R. S. and P. D. Joseph (2005). *Filtering for stochastic processes with applications to guidance*, Volume 326. American Mathematical Soc.

Chernis, T., N. Hauzenberger, F. Huber, G. Koop, and J. Mitchell (2025). Predictive density combination using bayesian machine learning. *International Economic Review 66*(3), 1287–1315.

Clark, T. E., F. Huber, G. Koop, M. Marcellino, and M. Pfarrhofer (2023). Tail forecasting with multivariate bayesian additive regression trees. *International Economic Review 64*(3), 979–1022.

Clark, T. E., F. Huber, G. Koop, M. Marcellino, and M. Pfarrhofer (2024). Investigating growth-at-risk using a multicountry nonparametric quantile factor model. *Journal of Business & Economic Statistics 42*(4), 1302–1317.

Creal, D., S. J. Koopman, and A. Lucas (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics 28*(5), 777–795.

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

Diebold, F. X. and R. S. Mariano (2002). Comparing predictive accuracy. *Journal of Business & economic statistics 20*(1), 134–144.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica: Journal of the econometric society*, 987–1007.

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.

Frühwirth-Schnatter, S. (2004). Efficient bayesian parameter estimation. *State Space and Unobserved Component Models: Theory and Applications*, 123–151.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculations of posterior moments. *Bayesian statistics 4*, 641–649.

Geweke, J. and G. Amisano (2011). Optimal prediction pools. *Journal of Econometrics 164*(1), 130–141.

Glosten, L. R., R. Jagannathan, and D. E. Runkle (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The journal of finance 48*(5), 1779–1801.

Gneiting, T. (2011). Making and evaluating point forecasts. *Journal of the American Statistical Association 106*(494), 746–762.

Gonçalves, K. C., H. S. Migon, and L. S. Bastos (2020). Dynamic quantile linear models: A bayesian approach. *Bayesian Analysis*, 335–362.

Greenspan, A. (2004). Risk and uncertainty in monetary policy. *American Economic Review 94*(2), 33–40.

Harvey, A. C. (2013). *Dynamic Models for Volatility and Heavy Tails: With Applications to Financial and Economic Time Series*. Cambridge University Press.

He, Z. and A. Krishnamurthy (2013). Intermediary asset pricing. *American Economic Review 103*(2), 732–770.

Hoeting, J. A., D. Madigan, A. E. Raftery, and C. T. Volinsky (1999). Bayesian model averaging: a tutorial. *Statistical science 14*(4), 382–401.

Kobayashi, G., S. Sugasawa, Y. Yamauchi, and D. Han (2026). Dynamic bayesian regression quantile synthesis for forecasting outlook-at-risk. *arXiv preprint arXiv:2603.11474*.

Koenker, R. and G. Bassett Jr (1978). Regression quantiles. *Econometrica: journal of the Econometric Society*, 33–50.

Koop, G. and D. Korobilis (2013). Large time-varying parameter vars. *Journal of Econometrics 177*(2), 185–198.

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

Martin, G. M., D. T. Frazier, W. Maneesoonthorn, R. Loaiza-Maya, F. Huber, G. Koop, J. Maheu, D. Nibbering, and A. Panagiotelis (2024). Bayesian forecasting in economics and finance: A modern review. *International Journal of Forecasting 40*(2), 811–839.

McAlinn, K., K. A. Aastveit, J. Nakajima, and M. West (2020). Multivariate bayesian predictive synthesis in macroeconomic forecasting. *Journal of the American Statistical Association 115*(531), 1092–1110.

McAlinn, K. and M. West (2019). Dynamic bayesian predictive synthesis in time series forecasting. *Journal of Econometrics 210*(1), 155–169.

Meinshausen, N. (2006). Quantile regression forests. *Journal of Machine Learning Research 7*(6).

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the econometric society*, 347–370.

Papaspiliopoulos, O., G. O. Roberts, and M. Sköld (2003). Non-centered parameterisations for hierarchical models and data augmentation. In *Bayesian Statistics 7: Proceedings of the Seventh Valencia International Meeting*, Volume 307. Oxford University Press, USA.

Prado, R. and M. West (2010). *Time series: modeling, computation, and inference*. Chapman and Hall/CRC.

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

Salgado, S., F. Guvenen, and N. Bloom (2019). Skewed business cycles. Technical report, National Bureau of Economic Research.

Simon, D. (2006). *Optimal state estimation: Kalman, H infinity, and nonlinear approaches*. John Wiley & Sons.

Surico, P. (2007). The fed's monetary policy rule and us inflation: The case of asymmetric preferences. *Journal of Economic Dynamics and Control 31*(1), 305–324.

Timmermann, A. (2006). Forecast combinations. *Handbook of economic forecasting 1*, 135–196.

West, M. and J. Harrison (1997). *Bayesian forecasting and dynamic models*. Springer.

Yu, K. and R. A. Moyeed (2001). Bayesian quantile regression. *Statistics & Probability Letters 54*(4), 437–447.

Yusupova, A., N. G. Pavlidis, and E. G. Pavlidis (2023). Dynamic linear models with adaptive discounting. *International Journal of Forecasting 39*(4), 1925–1944.

A Proof of Proposition 1: AL Likelihood Equivalence

This appendix provides the formal proof for Proposition 1, demonstrating that the non-centered parameterization strictly preserves the Asymmetric Laplace (AL) likelihood.

The non-centered conditional observation density and the prior for the latent shock e_t are defined as:

$$p(y_t|\boldsymbol{\theta}_t^\tau, v_t, e_t) = \mathcal{N}(y_t|\mathbf{F}_t'\boldsymbol{\theta}_t^\tau + a_\tau\sqrt{v_t}e_t,\ b_\tau v_t e_t)$$
$$p(e_t) = \text{Exp}(e_t|1)$$

To find the marginalized likelihood $p(y_t|\boldsymbol{\theta}_t^\tau, v_t)$, I integrate out the independent standard exponential shock e_t:

$$p(y_t|\boldsymbol{\theta}_t^\tau, v_t) = \int_0^\infty \frac{1}{\sqrt{2\pi b_\tau v_t e_t}} \exp\left(-\frac{(y_t - \mathbf{F}_t'\boldsymbol{\theta}_t^\tau - a_\tau\sqrt{v_t}e_t)^2}{2b_\tau v_t e_t}\right)\exp(-e_t)de_t \tag{25}$$

I introduce the direct scaling transformation $u_t = \sqrt{v_t}e_t$. This implies that $e_t = u_t/\sqrt{v_t}$ and the differential scales as $de_t = \frac{1}{\sqrt{v_t}}du_t$. Substituting these into the integral yields:

$$= \int_0^\infty \frac{1}{\sqrt{2\pi b_\tau v_t(u_t/\sqrt{v_t})}} \exp\left(-\frac{(y_t - \mathbf{F}_t'\boldsymbol{\theta}_t^\tau - a_\tau u_t)^2}{2b_\tau v_t(u_t/\sqrt{v_t})}\right)\exp\left(-\frac{u_t}{\sqrt{v_t}}\right)\frac{1}{\sqrt{v_t}}du_t \tag{26}$$

Simplifying the variance term in the denominator of the Gaussian constant and the exponent ($v_t/\sqrt{v_t} = \sqrt{v_t}$) results in:

$$= \int_0^\infty \frac{1}{\sqrt{2\pi b_\tau\sqrt{v_t}u_t}} \exp\left(-\frac{(y_t - \mathbf{F}_t'\boldsymbol{\theta}_t^\tau - a_\tau u_t)^2}{2b_\tau\sqrt{v_t}u_t}\right)\left[\frac{1}{\sqrt{v_t}}\exp\left(-\frac{u_t}{\sqrt{v_t}}\right)\right]du_t \tag{27}$$

Replacing $\sqrt{v_t}$ by σ_t (a simple notation transformation for visual equivalence, can be kept the same otherwise):

$$= \int_0^\infty \frac{1}{\sqrt{2\pi b_\tau \sigma_t u_t}} \exp\left(-\frac{(y_t - \mathbf{F}_t'\boldsymbol{\theta}_t^\tau - a_\tau u_t)^2}{2b_\tau \sigma_t u_t}\right) \left[\frac{1}{\sigma_t} \exp\left(-\frac{u_t}{\sigma_t}\right)\right] du_t \tag{28}$$

This equation is mathematically identical to the integral of the standard centered parameterization (Kozumi and Kobayashi, 2011):

$$= \int_0^\infty \mathcal{N}(y_t | \mathbf{F}_t'\boldsymbol{\theta}_t^\tau + a_\tau u_t,\ b_\tau \sigma_t u_t) \text{Exp}(u_t | 1/\sigma_t) du_t \tag{29}$$

By definition, this normal-exponential scale-location mixture resolves exactly to the Asymmetric Laplace distribution, $\mathcal{AL}(\mathbf{F}_t'\boldsymbol{\theta}_t^\tau, \sigma_t, \tau)$, completing the proof. ■

B Proof of Proposition 2: Exact GIG Posterior Update

This appendix provides the formal proof for Proposition 2, demonstrating that the full conditional posterior of the latent shock e_t under the non-centered parameterization resolves exactly to the Generalized Inverse Gaussian (GIG) distribution.

Proof. Let $\mu_t = \mathbf{F}_t'\boldsymbol{\theta}_t^\tau$ denote the provisional predictive mean. By Bayes' Theorem, the full conditional posterior for e_t is proportional to the product of the Gaussian observation likelihood and the Standard Exponential prior:

$$p(e_t | y_t, \mu_t, v_t) \propto p(y_t | \mu_t, v_t, e_t) p(e_t) \tag{30}$$

Substituting the known density functions of the non-centered hierarchy yields:

$$p(e_t | y_t, \mu_t, v_t) \propto \left[\frac{1}{\sqrt{e_t}} \exp\left(-\frac{(y_t - \mu_t - a_\tau \sqrt{v_t} e_t)^2}{2b_\tau v_t e_t}\right)\right] \exp(-e_t) \tag{31}$$

To isolate the kernel of this distribution with respect to e_t, I expand the squared term in the numerator of the Gaussian exponent:

$$(y_t - \mu_t - a_\tau\sqrt{v_t}e_t)^2 = (y_t - \mu_t)^2 - 2(y_t - \mu_t)a_\tau\sqrt{v_t}e_t + a_\tau^2 v_t e_t^2 \tag{32}$$

Dividing this expanded polynomial by the variance denominator ($2b_\tau v_t e_t$) separates the exponent into three distinct terms:

$$\frac{(y_t - \mu_t)^2}{2b_\tau v_t e_t} - \frac{2(y_t - \mu_t)a_\tau\sqrt{v_t}e_t}{2b_\tau v_t e_t} + \frac{a_\tau^2 v_t e_t^2}{2b_\tau v_t e_t} \quad = \quad \frac{(y_t - \mu_t)^2}{2b_\tau v_t}e_t^{-1} \; - \; \frac{(y_t - \mu_t)a_\tau}{b_\tau\sqrt{v_t}} \; + \; \frac{a_\tau^2}{2b_\tau}e_t \tag{33}$$

Observe that the middle cross-term, $\frac{(y_t-\mu_t)a_\tau}{b_\tau\sqrt{v_t}}$, is strictly independent of the random variable e_t. Consequently, $\exp\left(\frac{(y_t-\mu_t)a_\tau}{b_\tau\sqrt{v_t}}\right)$ acts merely as a scaling constant and is completely absorbed into the proportionality sign.

Discarding this constant term and combining the remaining terms with the Exponential prior ($\exp(-e_t) = \exp(-\frac{2}{2}e_t)$), the conditional posterior kernel simplifies to:

$$\begin{aligned} p(e_t|y_t, \mu_t, v_t) &\propto e_t^{-1/2}\exp\left(-\frac{1}{2}\left[\frac{(y_t - \mu_t)^2}{b_\tau v_t}e_t^{-1} + \frac{a_\tau^2}{b_\tau}e_t\right]\right)\exp\left(-\frac{2}{2}e_t\right) \\ &\propto e_t^{-1/2}\exp\left(-\frac{1}{2}\left[\frac{(y_t - \mu_t)^2}{b_\tau v_t}e_t^{-1} + \left(\frac{a_\tau^2}{b_\tau} + 2\right)e_t\right]\right) \end{aligned} \tag{34}$$

This derived kernel structurally matches the probability density function of the Generalized Inverse Gaussian (GIG) distribution. The standard GIG density for a random variable $x \sim \mathcal{GIG}(\lambda, \chi, \psi)$ is defined as:

$$f(x) \propto x^{\lambda-1}\exp\left(-\frac{1}{2}(\chi x^{-1} + \psi x)\right) \tag{35}$$

By direct pattern matching between the derived kernel and the standard GIG density, the

parameters characterizing the exact posterior distribution are identified as:

$$\lambda - 1 = -1/2 \implies \lambda = 0.5 \tag{36}$$

$$\chi = \frac{(y_t - \mu_t)^2}{b_\tau v_t} \tag{37}$$

$$\psi = \frac{a_\tau^2}{b_\tau} + 2 \tag{38}$$

Therefore, the full conditional posterior is exactly $e_t \sim \mathcal{GIG}(\lambda, \chi, \psi)$, allowing for a direct, rejection-free Gibbs block update. ■

C Proof of Proposition 3: Score Derivation as a Sufficient Statistic

This appendix provides the formal proof for Proposition 3, demonstrating that the asymmetric surprise variable Z_t isolates the core innovation of the marginal Asymmetric Laplace log-likelihood.

Proof. In the Generalized Autoregressive Score (GAS) framework, time-varying parameters are updated via the score—the first derivative of the log-likelihood with respect to the targeted parameter. Because the discount factor δ_t dictates the scale of structural variance injected into the filter, we derive the gradient of the predictive likelihood with respect to the scale parameter $\sqrt{v_t}$.

The likelihood of an observation y_t under the Asymmetric Laplace distribution, given the unshifted location forecast μ_t, structural volatility $\sqrt{v_t}$, and target quantile τ, is defined as:

$$p(y_t|\mu_t, \sqrt{v_t}, \tau) = \frac{\tau(1-\tau)}{\sqrt{v_t}} \exp\left(-\frac{\rho_\tau(y_t - \mu_t)}{\sqrt{v_t}}\right) \tag{39}$$

where $\rho_\tau(e_t) = e_t(\tau - \mathbb{I}(e_t < 0))$ is the asymmetric check-loss function. The conditional

log-likelihood of y_t can be expressed as:

$$\ln p(y_t|\mu_t, \sqrt{v_t}, \tau) \propto -\ln(\sqrt{v_t}) - \frac{\rho_\tau(y_t - \mu_t)}{\sqrt{v_t}} \tag{40}$$

Taking the partial derivative of the log-likelihood with respect to the structural scale parameter yields the theoretical score (∇_t):

$$\nabla_t = \frac{\partial \ln p(y_t)}{\partial \sqrt{v_t}} = -\frac{1}{\sqrt{v_t}} + \frac{\rho_\tau(y_t - \mu_t)}{v_t} \tag{41}$$

To stabilize the variance of the raw score across time-varying volatility regimes, it is optimal to scale the raw score by a time-varying factor $\mathcal{S}_t = \sqrt{v_t}$. Multiplying the raw score by this factor yields the dimensionless *scaled score*, s_t:

$$s_t = \sqrt{v_t}\nabla_t = \frac{\rho_\tau(y_t - \mu_t)}{\sqrt{v_t}} - 1 \tag{42}$$

Setting $Z_t = \frac{\rho_\tau(y_t - \mu_t)}{\sqrt{v_t}}$, the scaled score collapses to a direct equivalence: $s_t = Z_t - 1$. Because s_t represents the optimal, variance-stabilized gradient of the objective function, Z_t (a simple linear transformation of s_t) acts as a sufficient statistic driving the structural update. ■

D Proof of Proposition 4: Marginal Distribution of the Scaled Score

This appendix provides the formal proof for Proposition 4, demonstrating that under the null hypothesis of structural stability, the scaled check-loss Z_t follows a Standard Exponential distribution.

Proof. To find the marginal distribution of $Z_t = \rho_\tau(y_t - \mu_t)/\sqrt{v_t}$, I calculate its survival function $P(Z_t > z)$. For notational simplicity, let the scale parameter be $\sigma_t = \sqrt{v_t}$. The

integration is split into the right tail ($y_t > \mu_t$) and the left tail ($y_t < \mu_t$). Under the null hypothesis, the conditional probability density function of y_t is the AL distribution: $f(y_t) = \frac{\tau(1-\tau)}{\sigma_t} \exp\left(-\frac{\rho_\tau(y_t-\mu_t)}{\sigma_t}\right)$.

Right Tail ($y_t > \mu_t$): Here, $\rho_\tau(y_t - \mu_t) = (y_t - \mu_t)\tau$. The condition $Z_t > z$ implies $y_t > \mu_t + \frac{\sigma_t z}{\tau}$.

$$I_{right} = \int_{\mu_t + \frac{\sigma_t z}{\tau}}^{\infty} \frac{\tau(1-\tau)}{\sigma_t} \exp\left(-\frac{(y_t - \mu_t)\tau}{\sigma_t}\right) dy_t \tag{43}$$

Using the substitution $u = \frac{(y_t-\mu_t)\tau}{\sigma_t}$ (where $dy_t = \frac{\sigma_t}{\tau}du$), the integration bounds map from z to ∞:

$$I_{right} = (1-\tau)\int_z^{\infty} \exp(-u)du = (1-\tau)\exp(-z) \tag{44}$$

Left Tail ($y_t < \mu_t$): Here, $\rho_\tau(y_t - \mu_t) = (y_t - \mu_t)(\tau - 1)$. The condition $Z_t > z$ implies $y_t < \mu_t + \frac{\sigma_t z}{\tau-1}$.

$$I_{left} = \int_{-\infty}^{\mu_t + \frac{\sigma_t z}{\tau-1}} \frac{\tau(1-\tau)}{\sigma_t} \exp\left(-\frac{(y_t - \mu_t)(\tau-1)}{\sigma_t}\right) dy_t \tag{45}$$

Using the substitution $u = \frac{(y_t-\mu_t)(\tau-1)}{\sigma_t}$ (where $dy_t = \frac{\sigma_t}{\tau-1}du$), the limits map from ∞ down to z. This flips the sign of the integral to standard order:

$$I_{left} = \tau\int_z^{\infty} \exp(-u)du = \tau\exp(-z) \tag{46}$$

Summing the two tail probabilities yields the exact survival function for Z_t:

$$P(Z_t > z) = I_{right} + I_{left} = (1-\tau)\exp(-z) + \tau\exp(-z) = \exp(-z) \tag{47}$$

Because $S(z) = P(Z_t > z) = \exp(-z)$ uniquely defines the survival function of a Standard Exponential distribution, $Z_t \sim \text{Exp}(1)$ under the null hypothesis of structural stability. ■

E Growth-at-Risk Supplementary Results

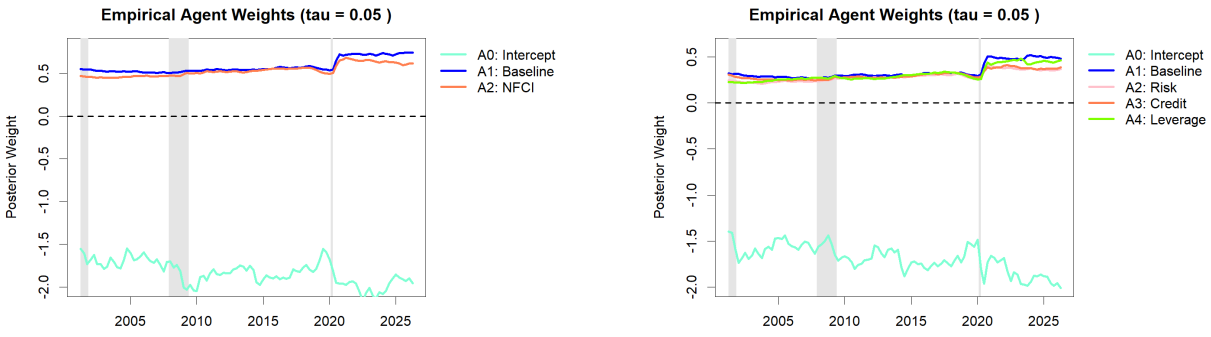
This appendix displays the raw statistics and dynamic synthesis weights in full scale.

Table 3: 1-Quarter Ahead Downside Risk Evaluation ($\tau = 0.05$) in Raw Terms

Model	Raw Check Loss		Hit Rate (%)	DM Statistic		DQ p-value
	Overall	Recession		Overall	Recession	
I. Universal Benchmarks						
Empirical (Unconditional)	0.637	4.013	6.93	-0.93	-2.44**	0.000
Agent: Baseline (Lagged GDP)	0.563	2.602	7.92	-0.92	-1.96*	0.067
ABG Skew-t	0.590	2.966	9.90	-0.67	-1.75*	0.225
II. Application 1: Aggregate NFCI						
QBPS (Aggregate)	**0.545**	**2.197**	**3.96**	**-1.68***	**-0.71**	**0.085**
BPS (Aggregate)	1.112	5.394	45.54	-2.88***	-3.55***	0.000
Agent: Aggregate NFCI	0.532	2.791	9.90	0.00	-1.74*	0.244
III. Application 2: NFCI Subcomponents						
QBPS (Subcomponents)	**0.532**	**2.153**	**4.95**	–	–	**0.358**
BPS (Subcomponents)	1.064	5.230	37.62	-2.70***	-3.48***	0.000
Agent: Risk	0.549	2.870	10.89	-0.30	-2.05**	0.171
Agent: Credit	0.550	2.804	8.91	-0.35	-2.06**	0.561
Agent: Leverage	0.539	2.591	9.90	-0.13	-1.31	0.435

Notes: The table reports raw asymmetric check losses; lower values indicate better out-of-sample predictive accuracy. Underlying base agents are stabilized via an ensemble average of 50 independent Quantile Random Forest (QRF) repetitions (number of trees = 500, node size = 40). Hit rate indicates the out-of-sample empirical percentage of model forecasts below the true series. Diebold-Mariano (DM) statistics are calculated using Newey-West HAC standard errors against the global benchmark (QBPS Subcomponents); negative statistics indicate underperformance relative to the benchmark. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The Dynamic Quantile (DQ) Test evaluates correct specification of the hit sequence at the 5% significance level, reporting the p-value where a value > 0.05 indicates a correctly specified model.

Figure 9: QBPS: Dynamic Weighting (Full, 1-Quarter Ahead)



(a) Application 1: Aggregate NFCI

(b) Application 2: NFCI Subcomponents

Notes: Target is the 1-quarter ahead 5th percentile ($\tau = 0.05$) of GDP growth. The two plots display posterior weight evolution for the candidate QRF agents across the Aggregate NFCI synthesis and the Subcomponent synthesis.

F Computational Robustness Checks

Table 4: Computational Robustness: 1-Quarter Ahead Downside Risk Evaluation ($\tau = 0.05$) in Raw Terms

Model	Raw Check Loss		Hit Rate (%)	DM Statistic		DQ p-value
	Overall	Recession		Overall	Recession	
Panel A: Extended Base-Agent Capacity (100 QRF Repetitions, Baseline 1000 MCMC)						
I. Universal Benchmarks						
Empirical (Unconditional)	0.637	4.013	6.93	-0.86	-2.53**	0.000
Agent: Baseline (Lagged GDP)	0.564	2.610	7.92	-0.77	-1.93*	0.066
ABG Skew-t	0.590	2.966	9.90	-0.58	-1.73*	0.225
II. Application 1: Aggregate NFCI						
QBPS (Aggregate)	**0.543**	**2.176**	**3.96**	**-0.14**	**0.34**	**0.087**
BPS (Aggregate)	1.097	5.317	45.54	-2.90***	-3.79***	0.000
Agent: Aggregate NFCI	0.530	2.796	9.90	0.24	-1.75*	0.246
III. Application 2: NFCI Subcomponents						
QBPS (Subcomponents)	**0.541**	**2.188**	**4.95**	–	–	**0.338**
BPS (Subcomponents)	1.048	5.226	39.60	-2.71***	-3.67***	0.000
Agent: Risk	0.553	2.899	9.90	-0.22	-2.09**	0.193
Agent: Credit	0.550	2.814	8.91	-0.17	-2.04**	0.561
Agent: Leverage	0.541	2.586	9.90	0.01	-1.26	0.434
Panel B: Extended MCMC Sampling (Baseline 50 QRF Repetitions, 2000 MCMC)						
I. Universal Benchmarks						
Empirical (Unconditional)	0.637	4.013	6.93	-0.90	-2.46**	0.000
Agent: Baseline (Lagged GDP)	0.562	2.588	7.92	-0.83	-1.99**	0.068
ABG Skew-t	0.590	2.966	9.90	-0.62	-1.72*	0.225
II. Application 1: Aggregate NFCI						
QBPS (Aggregate)	**0.548**	**2.193**	**3.96**	**-1.09**	**-0.13**	**0.081**
BPS (Aggregate)	1.100	5.282	44.55	-2.98***	-3.66***	0.000
Agent: Aggregate NFCI	0.534	2.806	9.90	0.07	-1.73*	0.241
III. Application 2: NFCI Subcomponents						
QBPS (Subcomponents)	**0.538**	**2.187**	**4.95**	–	–	**0.343**
BPS (Subcomponents)	1.080	5.198	38.61	-2.85***	-3.59***	0.000
Agent: Risk	0.549	2.866	9.90	-0.21	-2.17**	0.197
Agent: Credit	0.550	2.815	8.91	-0.25	-2.04**	0.560
Agent: Leverage	0.541	2.594	9.90	-0.07	-1.26	0.434

Notes: The table reports raw asymmetric check losses evaluated to confirm the computational robustness of the framework; lower values indicate better out-of-sample predictive accuracy. **Panel A** extends the base-agent capacity utilizing 100 independent Quantile Random Forest (QRF) repetitions under the baseline 1,000 MCMC iterations. **Panel B** extends the synthesis estimation utilizing 2,000 MCMC iterations under the baseline 50 QRF repetitions. In both panels, QRF parameters are fixed at 500 trees and a node size of 40. Hit rate indicates the out-of-sample empirical percentage of model forecasts below the true series. Diebold-Mariano (DM) statistics are calculated using Newey-West HAC standard errors against the global benchmark (QBPS Subcomponents within the respective panel). Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The Dynamic Quantile (DQ) Test evaluates correct specification of the hit sequence at the 5% significance level, reporting the p-value where a value > 0.05 indicates a correctly specified model.

G GaR 4-Quarter Ahead Forecast Results

Table 5: 4-Quarter Ahead Downside Risk Evaluation ($\tau = 0.05$)

Model	Relative Check Loss		Hit Rate (%)	DM Statistic		DQ Test
	Overall	Recession		Overall	Recession	
I. Universal Benchmarks						
Empirical (Unconditional)	0.948	0.788	6.93	-0.50	-3.40***	×
Agent: Baseline (Lagged GDP)	0.947	0.817	5.94	-0.54	-2.17**	×
ABG Skew-t	1.028	0.935	7.92	0.41	-1.59	×
II. Application 1: Aggregate NFCI						
QBPS (Aggregate)	**0.981**	**0.961**	**3.96**	**-0.59**	**-1.41**	✓
BPS (Aggregate)	0.479	0.492	49.50	-3.00***	-6.28***	×
Agent: Aggregate NFCI	0.965	0.885	7.92	-0.45	-2.00**	✓
III. Application 2: NFCI Subcomponents						
QBPS (Subcomponents)	**1.000**	**1.000**	**4.95**	–	–	✓
BPS (Subcomponents)	0.509	0.479	38.61	-2.40**	-5.04***	×
Agent: Risk	0.973	0.895	7.92	-0.29	-2.03**	✓
Agent: Credit	0.972	0.945	7.92	-0.44	-1.48	✓
Agent: Leverage	0.972	0.929	6.93	-0.38	-1.58	✓

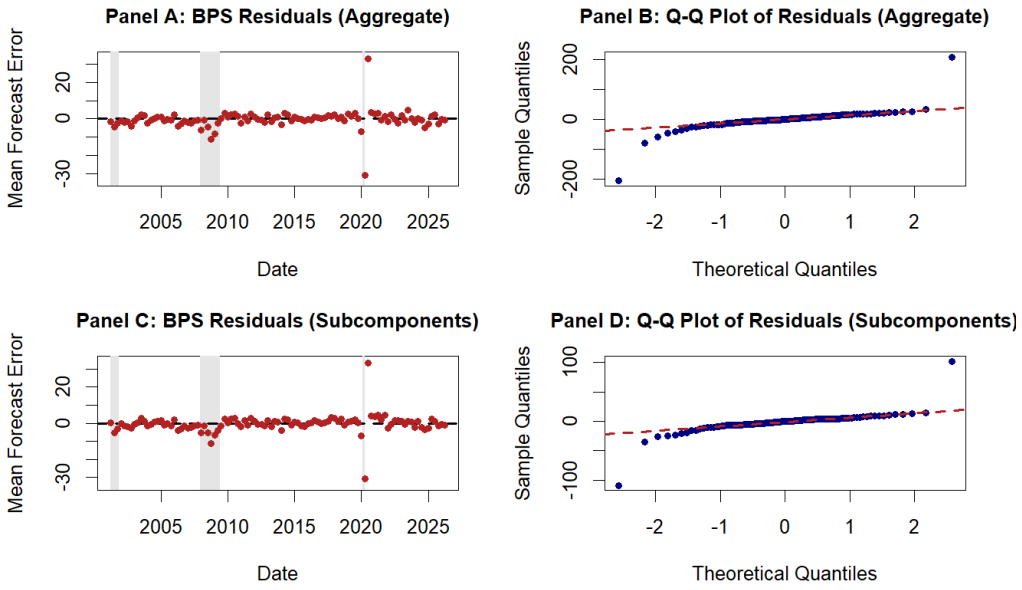
Notes: Relative losses are standardized against the QBPS (Subcomponents) model; values < 1 indicate higher loss (worse performance) than the global benchmark. Underlying base agents are stabilized via an ensemble average of 50 independent Quantile Random Forest (QRF) repetitions (number of trees = 500, node size = 40). Hit rate indicates the out-of-sample empirical percentage of model forecasts below the true series. Diebold-Mariano (DM) statistics are calculated using Newey-West HAC standard errors; negative statistics indicate underperformance relative to the benchmark. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The Dynamic Quantile (DQ) Test evaluates correct specification of the hit sequence at the 5% significance level, where (✓) denotes a passing p-value > 0.05 and (×) denotes a failure. Raw statistics can be found in Appendix E.

Table 6: 4-Quarter Ahead Downside Risk Evaluation ($\tau = 0.05$) in Raw Terms

Model	Raw Check Loss		Hit Rate (%)	DM Statistic		DQ p-value
	Overall	Recession		Overall	Recession	
I. Universal Benchmarks						
Empirical (Unconditional)	0.637	4.013	6.93	-0.50	-3.40***	0.000
Agent: Baseline (Lagged GDP)	0.638	3.874	5.94	-0.54	-2.17**	0.000
ABG Skew-t	0.588	3.384	7.92	0.41	-1.59	0.006
II. Application 1: Aggregate NFCI						
QBPS (Aggregate)	**0.616**	**3.294**	**3.96**	**-0.59**	**-1.41**	**0.127**
BPS (Aggregate)	1.262	6.432	49.50	-3.00***	-6.28***	0.000
Agent: Aggregate NFCI	0.626	3.575	7.92	-0.45	-2.00**	0.149
III. Application 2: NFCI Subcomponents						
QBPS (Subcomponents)	**0.604**	**3.164**	**4.95**	–	–	**0.306**
BPS (Subcomponents)	1.188	6.610	38.61	-2.40**	-5.04***	0.000
Agent: Risk	0.621	3.534	7.92	-0.29	-2.03**	0.205
Agent: Credit	0.622	3.348	7.92	-0.44	-1.48	0.077
Agent: Leverage	0.622	3.407	6.93	-0.38	-1.58	0.067

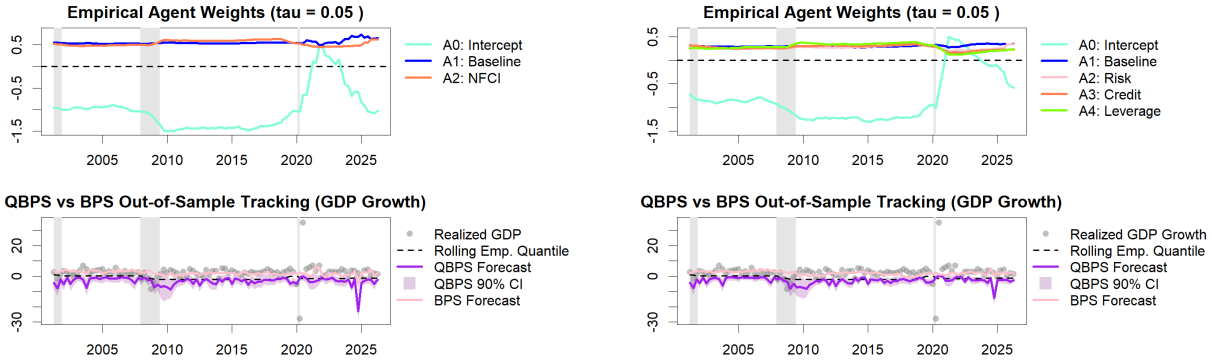
Notes: The table reports raw asymmetric check losses; lower values indicate better out-of-sample predictive accuracy. Underlying base agents are stabilized via an ensemble average of 50 independent Quantile Random Forest (QRF) repetitions (number of trees = 500, node size = 40). Hit rate indicates the out-of-sample empirical percentage of model forecasts below the true series. Diebold-Mariano (DM) statistics are calculated using Newey-West HAC standard errors against the global benchmark (QBPS Subcomponents); negative statistics indicate underperformance relative to the benchmark. Significance levels: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The Dynamic Quantile (DQ) Test evaluates correct specification of the hit sequence at the 5% significance level, reporting the p-value where a value > 0.05 indicates a correctly specified model.

Figure 10: Gaussian Benchmark (BPS) During Structural Breaks (4-Quarter Ahead)



Notes: **Panels A & C:** The plotted residuals represent the forecast errors of the 4-quarter ahead BPS predictive mean: $y_t - \hat{y}_t^{BPS}$. **Panels B & D:** The normal Q-Q plots display BPS' 4-quarter ahead standardized mean predictive residuals' correspondence with the theoretical quantiles of a standard normal distribution.

Figure 11: QBPS: Dynamic Weighting and Tail Forecasts (4-Quarter Ahead)

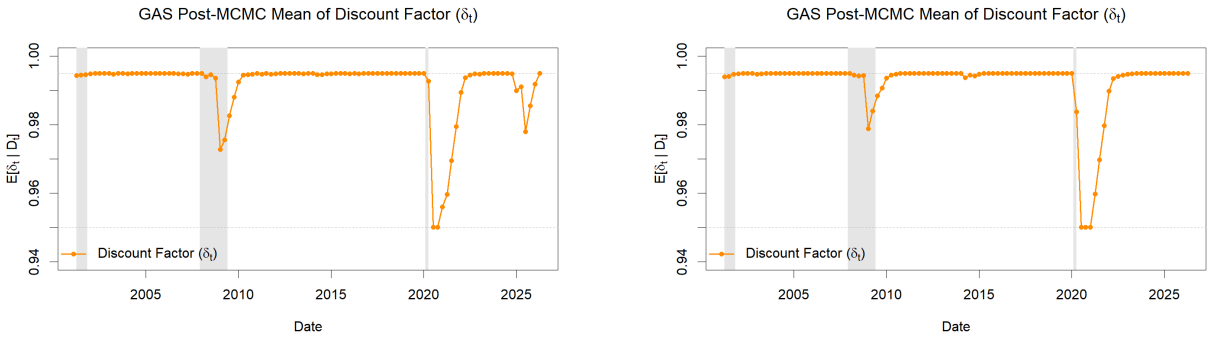


(a) Application 1: Aggregate NFCI

(b) Application 2: NFCI Subcomponents

Notes: Target is the 4-quarter ahead 5th percentile ($\tau = 0.05$) of GDP growth. **Top Sub-panels:** Posterior weight evolution for the candidate QRF agents across the Aggregate NFCI synthesis and the Subcomponent synthesis. Full-scale plots with intercepts are in Appendix E. **Bottom Sub-panels:** Out-of-sample forecast tracking. The Gaussian BPS benchmark (pink) optimizing for MSE is structurally pulled toward the mean, systematically failing to cover severe downside crashes. In contrast, QBPS (purple) correctly tracks the empirical tail, with the BPS forecast frequently falling completely outside the asymmetric QBPS 90% credible interval (shaded purple).

Figure 12: Score-Driven Dynamic Discounting (4-Quarter Ahead)

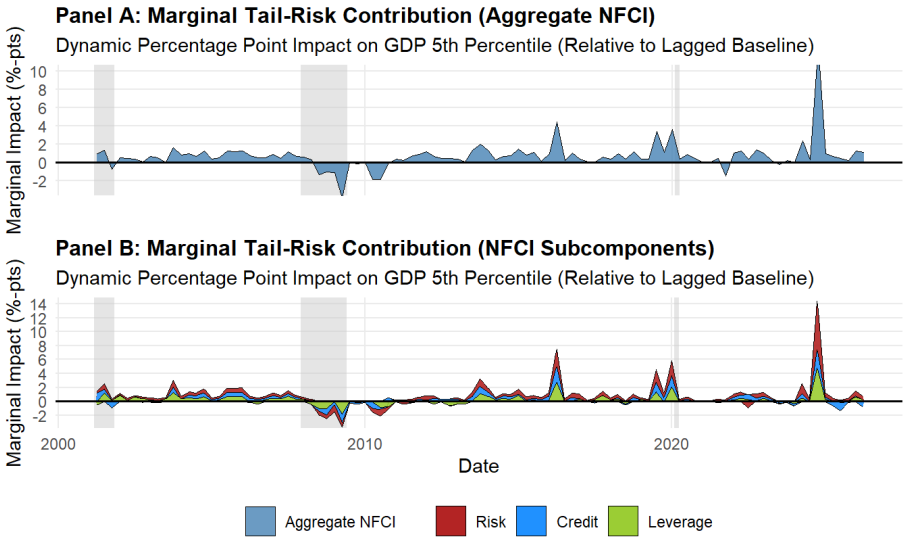


(a) Application 1: Aggregate NFCI

(b) Application 2: NFCI Subcomponents

Notes: Both panels display the mean of the dynamic discount factor (δ_t) extracted post-MCMC as the mean of effective MCMC draws. The dynamics follows from the observation-driven asymmetric discounting module. The left panel illustrates the structural adaptation for the aggregate NFCI synthesis, while the right panel displays the subcomponent synthesis. NBER recessions are marked gray.

Figure 13: Dynamic Growth-at-Risk (GaR) Marginal Attribution (4-Quarter Ahead)



Notes: Both panels display the percentage point contribution to 5th percentile ($\tau = 0.05$) GDP growth relative to a purely autoregressive lagged-GDP baseline ($\mathcal{M}_{j,t}^{(0.05)}$). **Panel A** displays the marginal downside impact evaluated using the aggregate Chicago Fed NFCI. **Panel B** utilizes the QBPS subcomponent synthesis to dynamically decompose the index, isolating the distinct, real-time tail-risk contributions of Risk, Credit, and Leverage. NBER recessions are shaded in gray.

H MCMC Convergence Diagnostics

Table 7 and Figure 14 report aggregated Markov Chain Monte Carlo (MCMC) convergence diagnostics evaluated continuously across the recursive out-of-sample windows. To demonstrate that the Quantile Bayesian Predictive Synthesis (QBPS) filter operates robustly without ad-hoc, dataset-specific tuning, I set a baseline of 200 burn-in and 1,000 post-burn-in draws across all recursive forecast origins. In dynamic state-space and time-varying parameter settings, maintaining an effective sampling efficiency of 10% to 20% (or 100 to 200 effective draws per 1,000 iterations)—is the standard benchmark for chain mixing (e.g., Primiceri, 2005; McAlinn and West, 2019). Convergence rests on two primary metrics: the average Effective Sample Size (ESS) across all dynamic states at each time step, and the empirical failure rate of the Geweke diagnostic (Geweke, 1992). Evaluated at a 95% confidence level ($|Z| > 1.96$), an ideally stationary Markov chain theoretically yields a 5.00% baseline failure rate due to Type I error.

In the controlled simulations (Table 7 Panel A), QBPS demonstrates satisfying computational stability at the targeted tail ($\tau = 0.05$). The sampler yields an average ESS of 424.01 and a Geweke failure rate of 5.95%—functionally matching the theoretical noise floor. This confirms the Generalized Inverse Gaussian (GIG) block updates run stably at extreme quantiles. At the median ($\tau = 0.50$), ESS drops to 270.58 and the failure rate rises to 15.75%. This divergence is an expected consequence of the likelihood specification. At the median, the Asymmetric Laplace distribution becomes symmetric. Because the parametric agents predict highly correlated central tendencies, the likelihood surface flattens into a multidimensional ridge. The MCMC sampler is forced to take smaller, highly correlated steps along this ridge—driving down ESS—yet it still extracts a robust posterior mean that captures the central predictive distribution.

The empirical applications replicate these mechanics across forecasting horizons. For the 1-quarter ahead evaluations (Table 7 Panel B), the predictive signals of financial variables remain strong, creating a well-defined likelihood surface. The filter maintains average ESS values above 370 and anchors Geweke failure rates between 7.92% and 8.25%. At the 4-

quarter ahead horizon (Table 7 Panel C), predictive signals attenuate. This weaker signal flattens the likelihood surface, granting the sampler freedom to take larger jumps across the posterior space. This flexibility mathematically reduces autocorrelation—stabilizing the average ESS around 389—while allowing minor drift in the dynamic weights, which marginally elevates the Geweke failure rates.

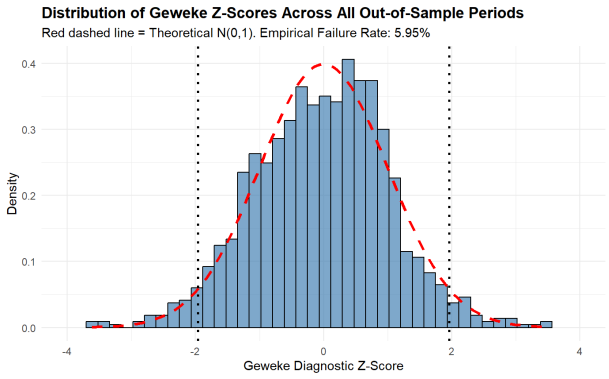
Figure 14 provides visual confirmation of the Geweke failure rate metrics. The empirical histograms of the aggregated Geweke Z-scores symmetrically trace the theoretical $\mathcal{N}(0,1)$ curve. The distributions for the tail simulation (a), the $H = 1$ empirical model (c) and and $H = 4$ empirical model (d) exhibit strict standard normal behavior, while the median simulation (b) displays the expected tail thickening associated with their flatter likelihood surfaces. Overall, these diagnostics supports that a lean 1,000-draw MCMC baseline gives stationary parameter estimation across the dynamic evaluation framework.

Table 7: Aggregated MCMC Convergence Diagnostics

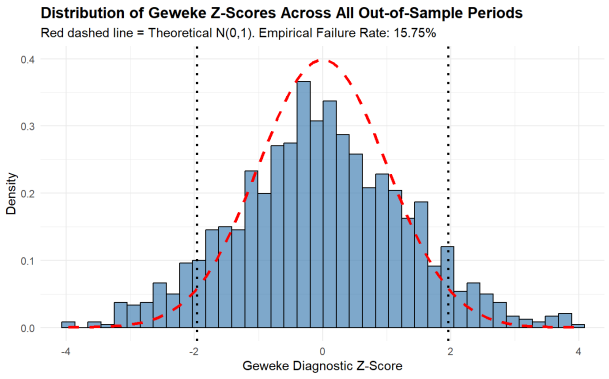
Synthesis Model	**Target (τ)**	**Horizon (H)**	**Avg. ESS**	**Geweke Failure Rate (%)**
Panel A: Simulation Studies				
Parametric Agents	0.05	1	424.01	5.95%
Parametric Agents	0.50	1	270.58	15.75%
Panel B: Empirical Application (1-Quarter Ahead)				
Aggregate NFCI	0.05	1	390.23	7.92%
NFCI Subcomponents	0.05	1	372.21	8.25%
Panel C: Empirical Application (4-Quarter Ahead)				
Aggregate NFCI	0.05	4	389.39	9.90%
NFCI Subcomponents	0.05	4	389.31	7.10%

Notes: The table reports aggregated MCMC convergence metrics across the recursive out-of-sample evaluation windows (using a strictly consistent 200 burn-in and 1,000 post-burn-in draws across all specifications). **Avg. ESS** denotes the average of the Effective Sample Size across all dynamic parameters at each time step. **Geweke Failure Rate** represents the empirical percentage of dynamic parameters yielding a Geweke diagnostic Z-score exceeding the 95% confidence threshold ($|Z| > 1.96$). A perfectly stationary Markov chain mathematically expects a 5.00% baseline failure rate due to standard Type I error.

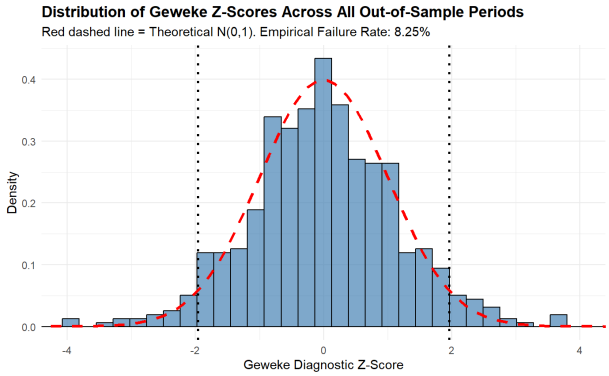
Figure 14: Empirical Distributions of Geweke Convergence Diagnostics



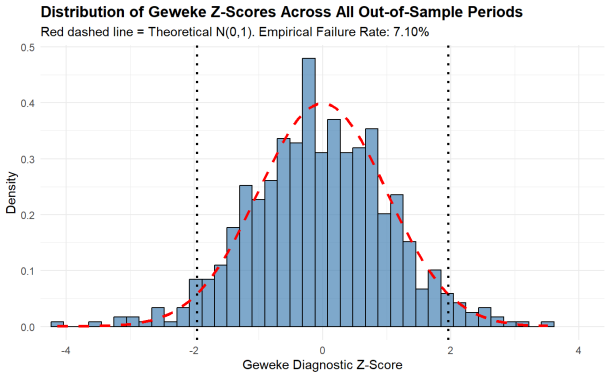
(a) Simulation ($\tau = 0.05$)



(b) Simulation ($\tau = 0.50$)



(c) Empirical Subcomponents ($H = 1$)



(d) Empirical Subcomponents ($H = 4$)

Notes: **(a)–(d):** The panels display the aggregated Geweke diagnostic Z-scores of the dynamic parameters, evaluated continuously across the recursive out-of-sample window for both the simulated and empirical (Application 2: Subcomponent Synthesis) frameworks. The red dashed line represents the theoretical standard normal distribution $\mathcal{N}(0, 1)$ expected under perfect Markov chain stationarity.