ISBA 2026
World Meeting, Nagoya, Japan

Bayesian quantile-based correction and synthesis of climate products

Antonio Aguirre Raquel Prado Bruno Sansó
Department of Statistics, University of California, Santa Cruz • jaguir@ucsc.edu

Code and reproducibility
github.com/AntonioAPDL/Project1

CONTEXT
Archives and ensembles

Recent climate analysis and forecasting rely heavily on climate-center model archives and ensemble prediction systems. **Retrospective products** replay a center's physical–numerical modeling system over past conditions, so simulated flows can be compared with observations and source-specific discrepancies can be estimated. **Forecast ensembles** run the operational system from current initial conditions to issue many plausible futures; their spread is useful, but discrepancies, horizons, and update cycles differ by source.

QUESTION
Correct first, then synthesize

Can quantile-level source discrepancies be estimated from diverse climate-product archives and propagated into newly issued ensemble forecasts on operational update timescales, while preserving an interpretable dynamic decomposition of source corrections, trend, seasonality, and exogenous effects?

Case study: daily San Lorenzo River flow at the USGS Big Trees gauge on the $\log(1 + x)$ scale. Before each forecast origin, observed flow is paired with retrospective ECMWF/GloFAS and NOAA/NWS products to estimate source-specific discrepancy states. After the origin, issued ensembles condition the forecast quantiles, while exogenous covariates from the origin-specific input bundle enter the transfer component. The result is a set of source-adjusted quantile forecasts synthesized into one posterior predictive distribution.

USGS observations
15-min public gauge record aggregated to daily log-flow; state-updating and held-out verification target.

ECMWF / GloFAS
European Centre for Medium-Range Weather Forecasts; retrospectives; daily 51-member ensemble forecasts, leads 1–28.

NOAA / NWS
U.S. National Weather Service under NOAA; retrospectives; hourly 7-member ensemble forecasts, daily leads 1–8.

Exogenous information
Local basin precipitation and soil moisture; the first generalized dynamic principal component (GDPC) summarizes key Pacific, Atlantic, and global climate indices.

VALIDATION
Rolling-origin holdout design

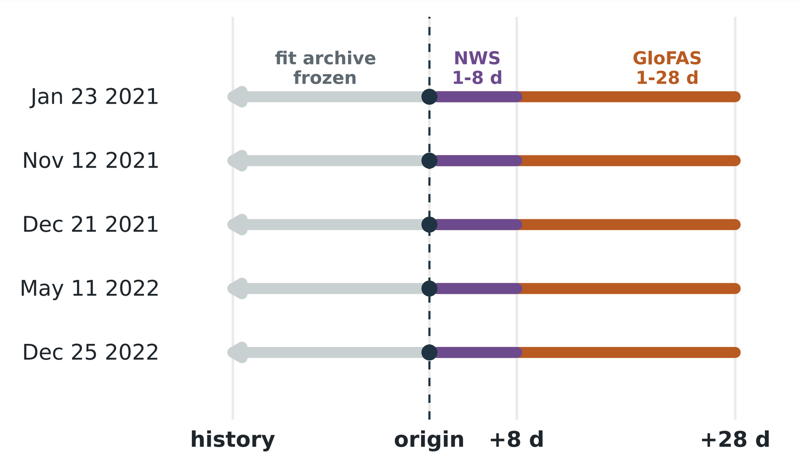

Strict temporal holdout. At each origin, fitting uses only USGS flow and retrospective products available through that date. Issued forecasts and forecast-window exogenous inputs from the origin-specific bundle enter after the origin; future USGS observations are reserved for scoring.

INFERENCE
Scalable Bayesian quantile fitting

Seven quantile-specific extended Dynamic Quantile Linear Models (exDQLMs) are fitted separately at levels 0.05, 0.20, 0.35, 0.50, 0.65, 0.80, and 0.95. Posterior predictive samples are generated for each fitted quantile lane and synthesized into one posterior predictive distribution.

7	$\leq$12,995	~2–4 h
quantile lanes	USGS archive days	fit + synthesis

Parallel quantile lanes
Each quantile level is fitted as its own state-space model; the lanes can run concurrently before forecast quantiles are synthesized.

Scalable VB updates
Within the mean-field variational updates, Gaussian state blocks are updated by Kalman filtering and backward smoothing. Laplace–Delta approximations update the non-conjugate exAL scale and skewness parameters.

MODEL
Main Workflow

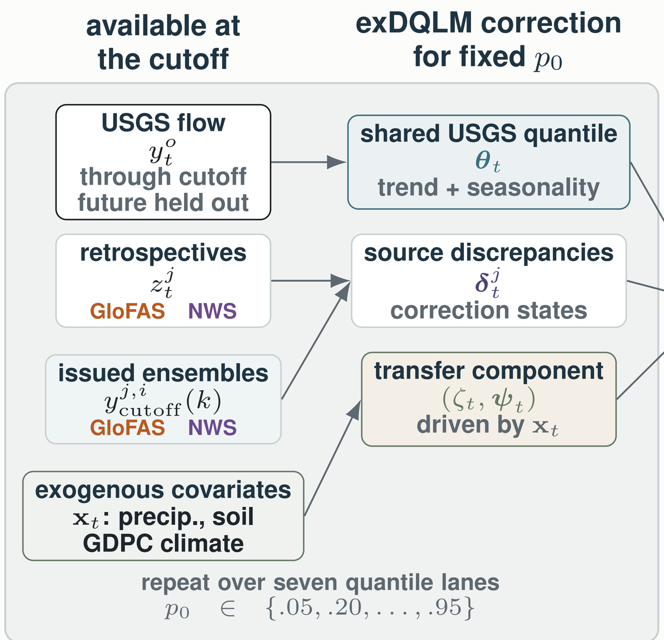

Source-aware exDQLM workflow. For each target quantile p_0, pre-origin USGS observations and retrospective product histories estimate a latent river-flow quantile and source-specific discrepancy states. After the cutoff, issued ensemble forecasts and origin-bundle covariates propagate corrected forecast-window quantiles. The seven fitted quantile lanes are then synthesized into one posterior predictive distribution using isotonic adjustment and monotone rearranged draws, then scored only with held-out USGS observations.

extended Dynamic Quantile Linear Model (exDQLM)
For each target quantile p_0, let T denote the forecast origin/cutoff. The exDQLM uses an extended asymmetric Laplace (exAL) likelihood whose location is the p_0 quantile. The selected model links observed flow, retrospective product histories, and issued forecast members through one latent river-flow quantile. Source discrepancies estimate pre-origin corrections; forecast ensembles and origin-bundle covariates then propagate corrected quantiles after the origin.

$$\begin{aligned}
&\textbf{Observed flow:} && y_t^o \sim \mathrm{exAL}_{p_0}(\mathbf{F}_t'\boldsymbol{\theta}_t + \zeta_t, \sigma^o, \gamma^o), && 1 \le t \le T,\\
&\textbf{Retrospectives:} && z_t^j \sim \mathrm{exAL}_{p_0}\left(\mathbf{F}_t'(\boldsymbol{\theta}_t + \boldsymbol{\delta}_t^j) + \zeta_t, \sigma^j, \gamma^j\right), && j \in \mathcal{J},\ 1 \le t \le T,\\
&\textbf{Forecast ensembles:} && y_T^{j,i}(k) \sim \mathrm{exAL}_{p_0}\left(\mathbf{F}_{T+k}'(\boldsymbol{\theta}_{T+k} + \boldsymbol{\delta}_{T+k}^j) + \zeta_{T+k}, \sigma^j, \gamma^j\right), && j \in \mathcal{J},\ i \in \mathcal{I}_j,\ k \in \mathcal{K}_j.
\end{aligned}$$

$$\begin{aligned}
\textbf{Trend/seasonal:}\quad & \boldsymbol{\theta}_t \mid \boldsymbol{\theta}_{t-1} \sim \mathcal{N}(\mathbf{G}_t\boldsymbol{\theta}_{t-1}, \mathbf{W}_t^{\theta}),\\
\textbf{Discrepancies:}\quad & \boldsymbol{\delta}_t^j \mid \boldsymbol{\delta}_{t-1}^j \sim \mathcal{N}(\mathbf{G}_t\boldsymbol{\delta}_{t-1}^j, \mathbf{W}_t^{\delta^j}),\\
\textbf{Transfer:}\quad & \zeta_t \mid \zeta_{t-1}, \boldsymbol{\psi}_{t-1} \sim \mathcal{N}(\lambda\zeta_{t-1} + \mathbf{x}_t'\boldsymbol{\psi}_{t-1}, w_t^{\zeta}),\\
\textbf{Covariate effects:}\quad & \boldsymbol{\psi}_t \mid \boldsymbol{\psi}_{t-1} \sim \mathcal{N}(\boldsymbol{\psi}_{t-1}, \mathbf{W}_t^{\psi}).
\end{aligned}$$

Notation: y_t^o is USGS flow, z_t^j retrospective product j, and $y_T^{j,i}(k)$ forecast member i from source j at lead k; $\mathcal{J}, \mathcal{I}_j, \mathcal{K}_j$ index sources, members, and leads. $\mathbf{F}_t$ loads the state, $\mathbf{G}_t$ evolves trend–seasonal blocks, and $\mathbf{W}_t, w_t^{\zeta}$ are evolution covariances. $\boldsymbol{\theta}_t$ is the shared river quantile, $\boldsymbol{\delta}_t^j$ the source correction, and $(\zeta_t, \boldsymbol{\psi}_t, \lambda)$ the transfer block driven by $\mathbf{x}_t$ (precipitation, soil moisture, GDPC). σ^s, γ^s are exAL channel scale and skewness. **Priors:** $\boldsymbol{\theta}_0 \sim \mathcal{N}(\mathbf{m}_0, \mathbf{C}_0)$, $\sigma^s \sim IG(a_s, b_s)$, and $\gamma^s \sim t_{(L,U)}(0, \phi^s, \nu^s)$; component discounting controls pre-cutoff state covariances $\mathbf{W}_t$. In the forecast window, the active state-block covariance receives an Inverse Wishart prior centered on a scaled version of the final pre-cutoff covariance restricted to active forecast states.

RESULTS
1–28-step-ahead forecast CRPS comparison

Selected exDQLM has the lowest 1–28-step-ahead CRPS at all five origins CRPS ↓

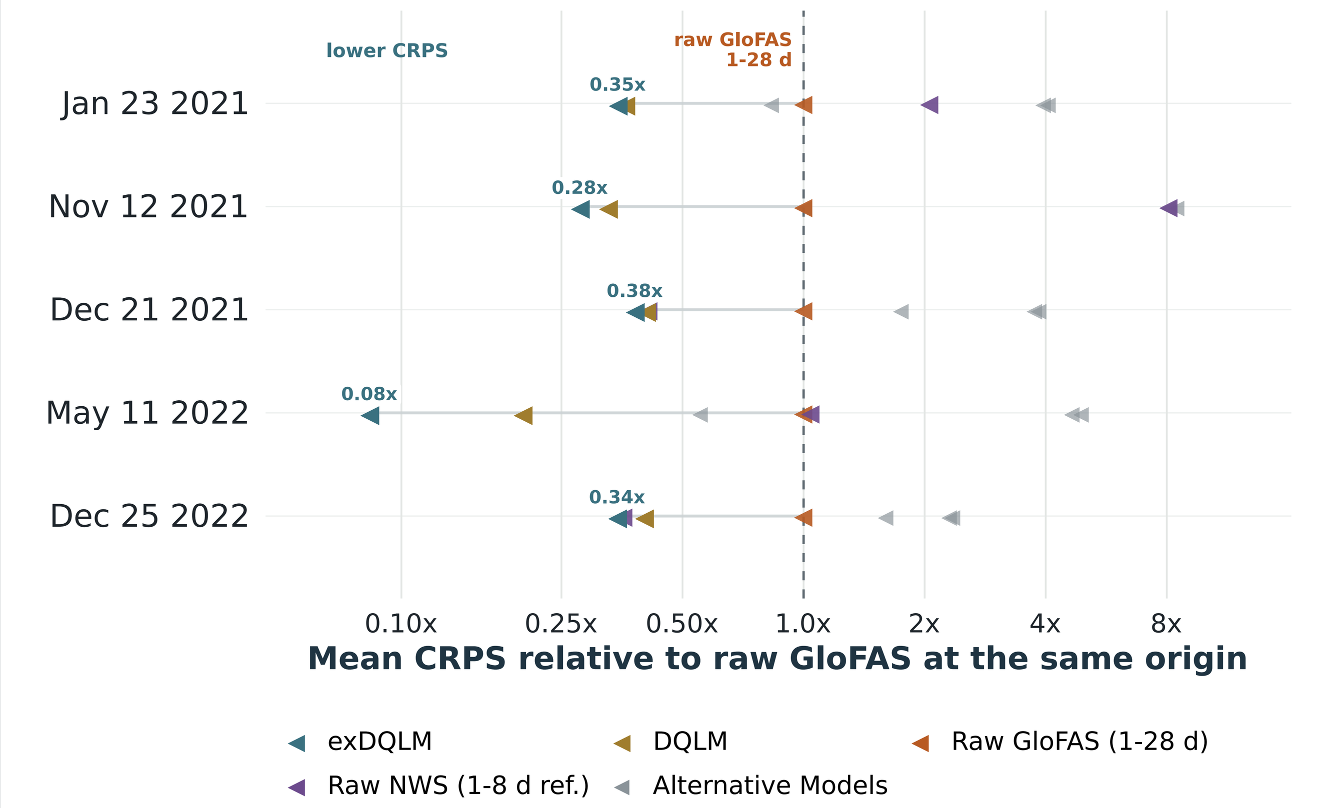

CRPS is the continuous ranked probability score; lower values are better. Left-pointing markers emphasize lower CRPS on the horizontal scale. Colors compare the multivariate keep exDQLM and DQLM variants, raw forecast products, and gray alternative models. Raw GloFAS is the 1–28-day reference, while raw NWS is shown as a 1–8-day horizon reference on the same raw-GloFAS normalization.

FIT
Wet-period quantile fit

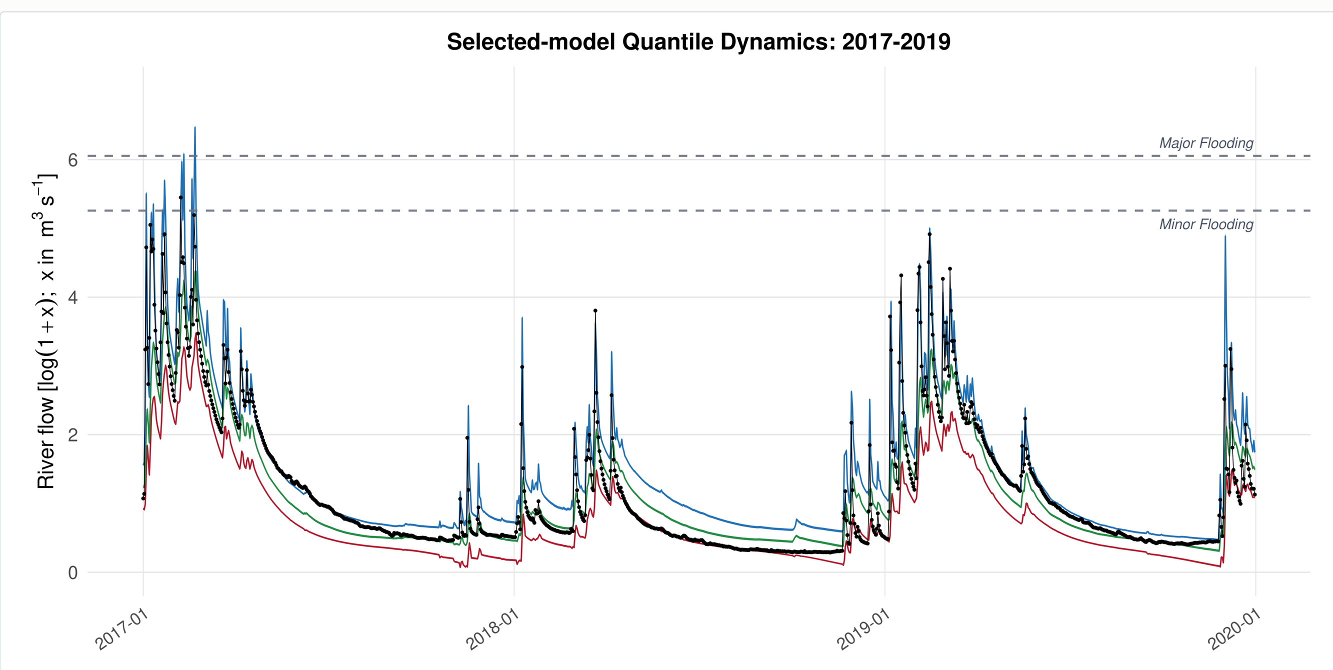

Historical fit diagnostic. Posterior mean fitted conditional quantiles and 95% credible intervals from the selected exDQLM during the 2017–2019 wet period. The 5th, 50th, and 95th quantile fits are shown with observed USGS flow; wider upper-tail bands make winter high-flow behavior visible.

EXAMPLE
Three forecast-window examples

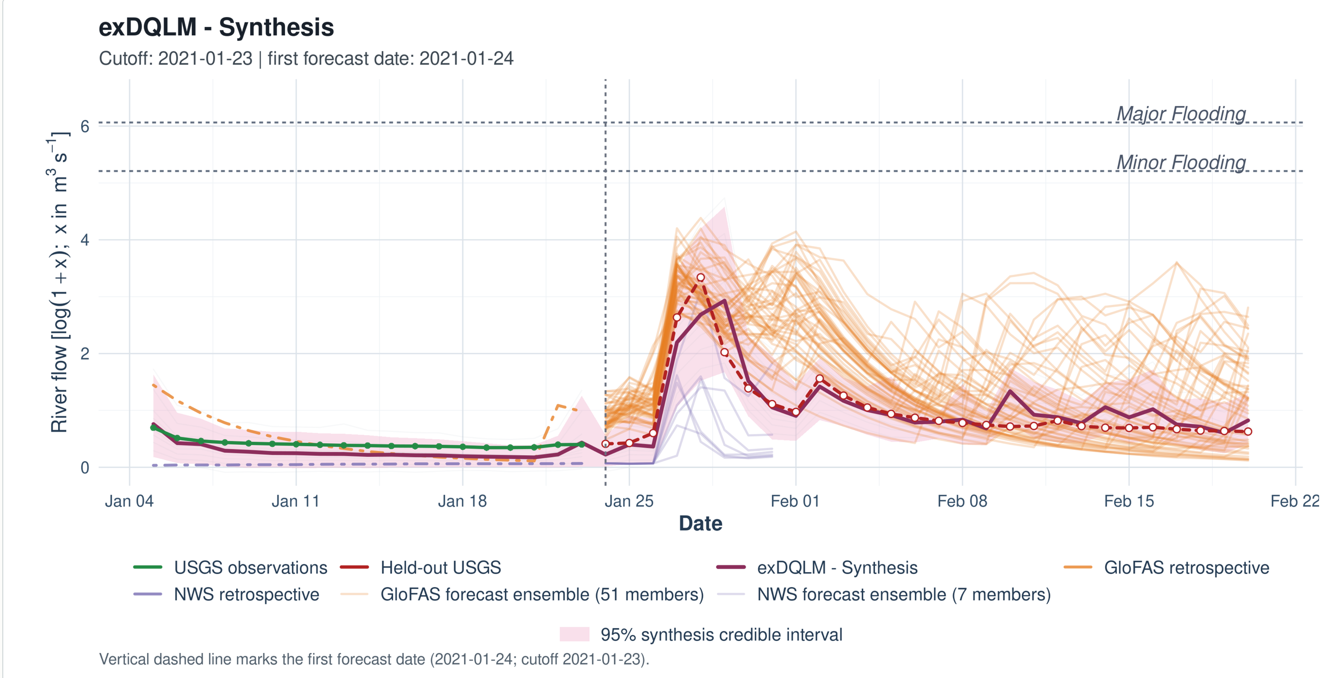

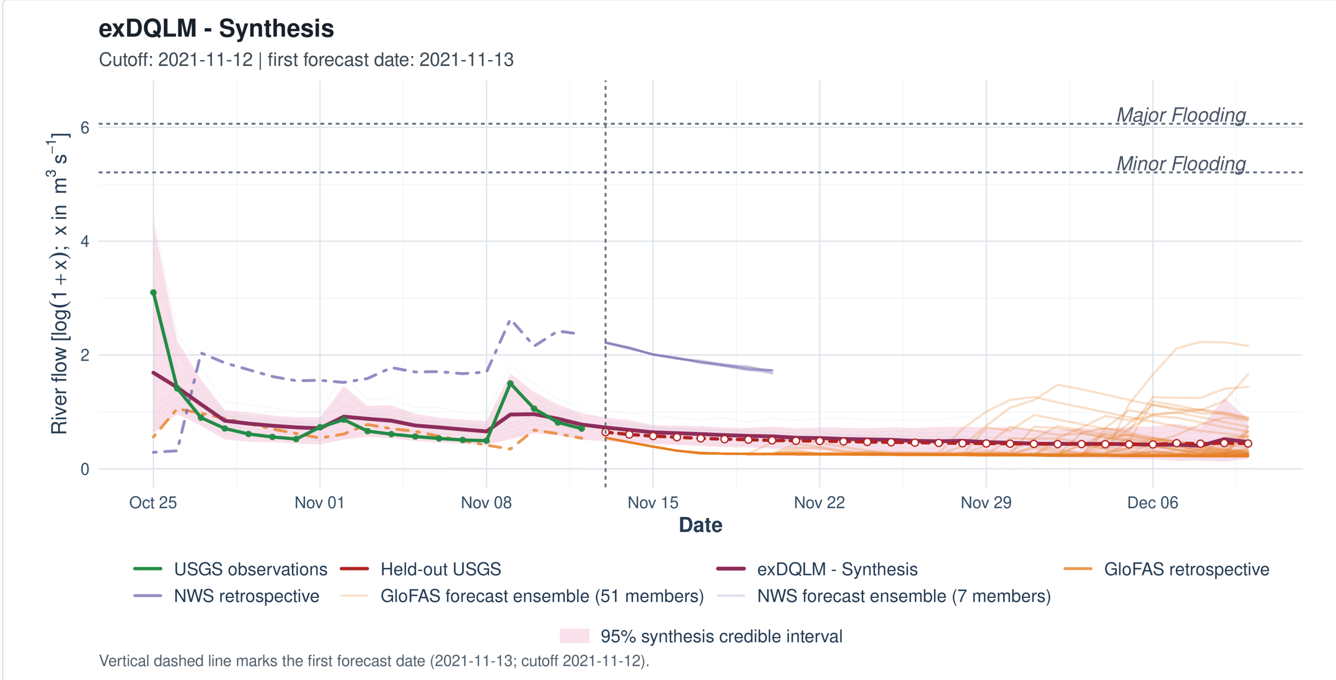

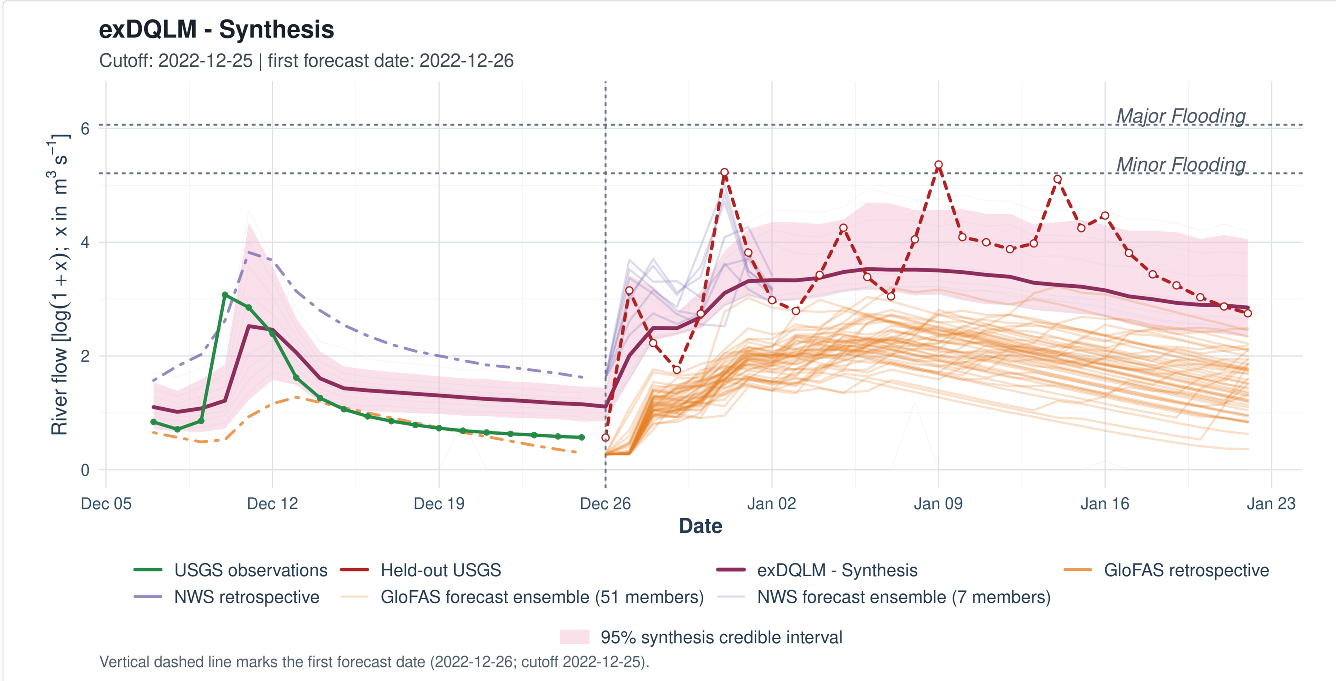

Origin-level illustrations. Three selected exDQLM forecast windows show how corrected quantile forecasts combine into a posterior predictive envelope after the cutoff. The panels are illustrative checks of behavior at individual origins; the aggregate comparison remains the five-origin CRPS benchmark.

DYNAMICS
80-month harmonic diagnostic

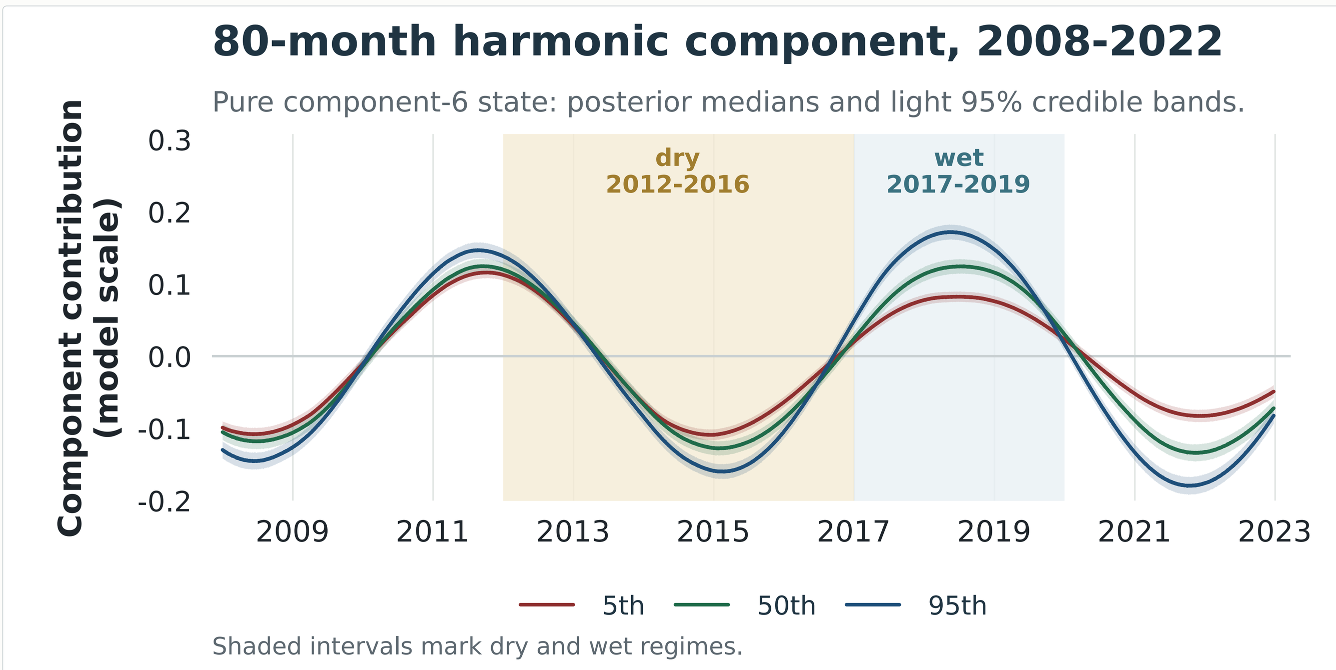

Retrospective diagnostic. One component of the interpretable state decomposition. From the selected exDQLM at the 2022-12-25 origin, lines show the 80-month harmonic state for the 5th, 50th, and 95th target-quantile lanes over 2008–2022; ribbons show 95% credible intervals. Shaded dry (2012–2016) and wet (2017–2019) regimes mark Santa Cruz-relevant water-supply contrasts, and the recovered long-cycle state summarizes multi-year variation rather than additional forecast-validation evidence.

TAKEAWAYS
Main takeaways

- **Source correction.** Quantile-level source discrepancies are estimated from retrospective archives and propagated into newly issued ensemble forecasts.
- **Operational synthesis.** Seven corrected quantile lanes are fitted separately, run in parallel, and synthesized into one posterior predictive distribution on update timescales.
- **Interpretable dynamics.** The state decomposition separates source corrections from trend, seasonality, and exogenous effects, linking product discrepancies to river-flow dynamics.

Limitations and next steps

- **Scope.** Evidence is from one gauge and five held-out origins; broader basin and regime testing is needed before operational generalization.
- **Quantile coherence.** Crossing was not material here; joint quantile estimation could improve coherence and possibly forecast skill.
- **Covariate uncertainty.** Forecast-window precipitation and soil moisture enter as staged ensemble summaries; a fuller model would propagate covariate-ensemble uncertainty through the transfer component.
- **Model selection.** Linear dynamics aid interpretation, while discount factors and transfer rate λ still require tuning before operational use.