Applied Economics and Finance; Vol. 13, No. 1; 2026
ISSN 2332-7294 E-ISSN 2332-7308
Published by Redfame Publishing

Modeling Reserve-Currency Attrition: A Bayesian Neural Framework for U.S. Dollar Reserve Holdings

Anna Knezevic[1]

[1] School of Electrical Engineering, Computing and Mathematical Sciences, Curtin University of Technology, Australia

Correspondence: Anna Knezevic, School of Electrical Engineering, Computing and Mathematical Sciences, Curtin University of Technology, Australia.

Received: November 1, 2025 Revised: December 22, 2025 Accepted: January 23, 2026

Available online: February 12, 2026 URL: https://doi.org/10.11114/aef.v13i1.8117

Abstract

This study develops a hybrid Bayesian modeling framework to estimate the probability of large-scale foreign divestment from U.S. Treasury securities and to assess the stability of the U.S. dollar's reserve-currency status. The model integrates a Mixture Density Network (MDN) with Hamiltonian Monte Carlo (HMC) sampling to approximate the conditional distribution of reserve holdings given observed macroeconomic fundamentals and an inferred latent variable representing geopolitical alignment. Principal component analysis and hierarchical regularization are applied to ensure parsimony and mitigate overfitting in a small-sample environment. The empirical hazard rate is defined as the annual probability of a 10 percent contraction in foreign U.S. Treasury holdings—a threshold consistent with historical reserve reallocations during the sterling's decline and central-bank portfolio adjustment behavior. Results indicate an estimated 6–12 percent annual probability at a 95 percent confidence interval, implying that a 10 percent drawdown could occur roughly once every 8–16 years under current macro-financial conditions. While Gaussian mixture components likely understate extreme tail risks, the findings highlight the resilience of U.S. reserve status and suggest that any future transition would be gradual, multi-decade, and contingent on structural geopolitical realignments rather than cyclical shocks.

Keywords: Reserve Currency, Bayesian Econometrics, Hamiltonian Monte Carlo, Financial Stability, U.S. Treasury, Tail Risk

1. Introduction

1.1 Economic Framework and Equilibrium Interpretation

This paper treats reserve reallocations as forward-looking portfolio choices under uncertainty. The empirical strategy uses Bayesian latent-variable inference as a computational device to approximate the equilibrium allocation outcomes implied by rational expectations and global asset-substitution logic.

Global reserve-currency transitions are gradual, path-dependent processes shaped by network externalities, safe asset supply constraints, and geopolitical alignments (Eichengreen, 2017; Caballero & Farhi, 2017). Historical precedent—most notably the drawn-out replacement of the British pound sterling by the U.S. dollar—illustrates decade-scale adjustment dynamics rather than abrupt regime shifts (Cirje, 2021; Horii, 1986). Consequently, the "loss of reserve status" is conceptualized not as a one-time event but as a sequence of persistent adverse movements to foreign U.S. dollar holdings that gradually undermine network dominance. These movements stem from both economic and political sources.

The 10 percent contraction threshold serves as a policy-relevant benchmark rather than a purely statistical cutoff. Empirically, variations of this magnitude correspond to significant but manageable shifts in official reserve composition, typically triggering liquidity responses or communication interventions by major central banks (Sarr & Lybek, 2002). Historically, during the sterling's gradual displacement by the U.S. dollar between the 1920s and 1950s, cumulative reserve reallocations occurred in multi-year increments of roughly 8–12 percent before broader regime changes became visible (Cirje, 2021; Eichengreen, 2017). From a modeling standpoint, adopting a 10 percent threshold balances signal sensitivity and probabilistic realism under a Gaussian mixture framework: it defines an event large enough to represent systemic movement while avoiding over-weighting of improbable tail outcomes given the light-tailed normal assumption. This level thus captures economically meaningful stress consistent with both historical precedent and statistical tractability within the latent-variable Bayesian structure.

Multi-year transition probabilities are obtained through compounding these annual hazards. Core explanatory variables follow the theoretical transmission channels through which global reserve adjustments occur—trade intensity (network and real-activity linkage), the term-structure level (safe-asset demand and macro stance), domestic inflation (price stability credibility), public debt ratios (safe-asset supply), and foreign reserve composition (external portfolio allocation). A latent variable captures alignment and institutional trust as a non-economic geopolitical factor (Eichengreen, 2017; Obstfeld & Rogoff, 1996).

Unlike recent work that quantifies the macroeconomic impacts of reserve-currency decline (CFA Institute, 2024, Xing, 2025), this study inverts the perspective by estimating the probability of foreign divestment from U.S. dollar assets. This reframing moves the discussion from ex post measurement to ex ante vulnerability assessment, focusing on structural preconditions for a transition rather than its consequences.

Three theoretical frameworks underpin the analysis:

1. Network Externality Theory, explaining persistent use of the incumbent reserve currency due to coordination and liquidity benefits (Obstfeld & Rogoff, 1996).
2. The Triffin Dilemma, which posits that global demand for reserve assets forces the issuer into external imbalances that ultimately erode confidence (Ferrara, 2023).
3. The Safe Asset Framework, linking global demand for highly rated, liquid debt to perceived fiscal capacity and institutional credibility (Caballero & Farhi, 2017).

Empirically, these frameworks converge through observable macro-financial channels—trade flows, inflation, foreign reserves, and yields on U.S. Treasuries (Robert, 2021; Sarr & Lybek, 2002). The exclusion of high-frequency exchange-rate fluctuations reflects the intention to isolate real-side adjustments rather than investment hedging and other speculative volatility.

Industry sentiment corroborates this analytical framing. Survey evidence suggests diminishing expectations of perpetual dollar dominance (Deane & Fines, 2023, Xing, 2025). Meanwhile, credit-rating agencies have noted fiscal deterioration and debt expansion—U.S. public debt has more than doubled every decade since 2000's, while GDP growth has remained modest (Lustig, 2024). However, reserve status hinges more on *credibility and liquidity* than on credit ratios alone, aligning with the borrowing privilege hypothesis (Lustig, 2024).

From a political economy standpoint, persistent inflation acts as a constraint on fiscal expansion because it redistributes costs to domestic voters, theoretically prompting policy recalibration (Ferrara, 2023). Inflation also weakens external confidence in the currency as a store of value, while higher yields may temporarily attract foreign investors but crowd out domestic financing—creating a "borrowing privilege trap" (Lustig, 2024). These interdependencies highlight that reserve status emerges from both economic fundamentals and policy credibility.

Historical comparison to the sterling's decline reinforces this view. Between the 1920s and 1950s, a combination of economic contraction, declining trade share, and wartime fiscal strain eroded sterling's role (Cirje, 2021; Eichengreen & Flandreau, 2009). Over roughly two decades, sterling lost nearly half its value relative to the dollar. This gradual unwinding supports a long-cycle interpretation of reserve transitions, in which cumulative credibility losses accumulate before functional replacement—motivating an approach that measures vulnerability as a tail probability rather than predicting a single transition date.

Because the sample is short (monthly 2000–2025) and the object of interest is a forward-looking tail probability rather than a point forecast, the estimation strategy combines flexible mixture modeling with Bayesian latent-variable inference to quantify uncertainty in both parameters and unobserved states. This Bayesian framework can be interpreted as a computational approximation to equilibrium outcomes under rational expectations and global asset substitution, enabling direct probability statements about large 12-month drawdowns.

1.2 Economic Framework, Equilibrium Interpretation, and Roadmap

Because the sample is short (monthly 2000–2025) and the object of interest is tail inference rather than point prediction, the paper models the *conditional distribution* of 12-month changes in foreign U.S. Treasury holdings given macro-financial covariates and an unobserved latent alignment factor. The Bayesian latent-variable setup is used as a computational device to approximate equilibrium portfolio allocations that are consistent with rational expectations and global asset substitution (Obstfeld & Rogoff, 1996). Posterior sampling via Hamiltonian Monte Carlo (HMC) supports stable inference on latent-state uncertainty and enables direct probability statements about large 12-month drawdowns.

Contribution and roadmap. The paper contributes (i) an ex-ante, probability-based vulnerability metric for reserve-currency attrition (operationalized as a 12-month −10% foreign-holdings event), and (ii) a small-sample Bayesian mixture framework that integrates observables with a latent alignment factor while reporting explicit stability/sensitivity diagnostics across specifications. The remainder is organized as follows: Section 2 describes data

construction and frequency harmonization; Sections 3–6 set out the model and estimation; Section 7 reports results; Section 8 discusses limitations; Section 9 concludes.

1.3 Related Literature

The paper connects three strands. First, it builds on the historical and political-economy literature on reserve-currency competition and long transition dynamics, including evidence on when the dollar overtook sterling and the gradual nature of reserve reallocation (Eichengreen & Flandreau, 2009; Eichengreen, 2011). Second, it links reserve-asset demand to macro-financial channels emphasized in international macro and safe-asset work: global imbalances and portfolio allocation pressures (Bernanke, 2005; Obstfeld & Rogoff, 1996) and the convenience-yield/safe-asset demand view of Treasuries that can generate "risk-off" co-movement between yields and foreign holdings (Krishnamurthy & Vissing-Jorgensen, 2012). Third, it relates to the dominant-currency paradigm literature documenting the dollar's outsized role in trade invoicing and shock transmission, which helps motivate trade variables as state indicators for USD-asset demand (Gopinath et al., 2020; Gopinath & Itskhoki, 2021). Methodologically, the paper follows work on probabilistic modeling of conditional distributions (mixture density networks) and modern Hamiltonian Monte Carlo implementations used for stable posterior inference (Bishop, 1994; Neal, 2011; Hoffman & Gelman, 2014).

Extending this section's geopolitical strand, recent empirical work treats U.S. financial sanctions as a proxy for diplomatic power and a potential catalyst for reserve diversification, but this channel is an incomplete measure of reserve-currency attrition because it captures only one dimension of "alignment/trust." Sanusi, Mamman, Iliyasu & Abubakar (2025) use a LASSO specification and find that the imposition of U.S. financial sanctions is associated with a lower dollar share in global reserve currencies, consistent with a sanctions-induced incentive to diversify. Goldberg & Hannaoui (2024; revised 2024) similarly link geopolitics to reserve composition, but emphasize that the aggregate decline in USD reserve shares is not a broad-based preference shift: it is driven by a small set of countries, and sanctions/geopolitical distance matter mainly when a central bank's core foreign-currency liquidity needs are already met. In contrast, Gerding & Hartley (2024) updates multiple indicators (reserves, FX turnover, debt denomination, trade invoicing) and reports that dollar dominance remains unchanged through late 2023, implying a long horizon rather than an imminent break—without offering a specific transition-date quantification. This is consistent with this paper's tail-risk framing, which translates reserve attrition into an annual hazard of a −10% foreign-holdings event (6–12% per year), implying rare but recurrent drawdowns and a gradual, multi-decade transition process.

2. Data Processing

The empirical analysis uses multi-source macro-financial data with differing temporal coverage and frequency. The design objective is to minimize model overfitting and ensure empirical robustness given the relatively short sample history—particularly for trade data, which extends only to 2000. The data-processing framework follows the principle of parsimony in econometric modeling (Diebold & Li, 2006), balancing informational efficiency against the risk of spurious correlation.

Figure 1 shows non-overlapping annual correlations among key macro-financial variables to highlight regime shifts in their co-movement over time (e.g., the post-2020 strengthening of the export–import relationship as yields fell).

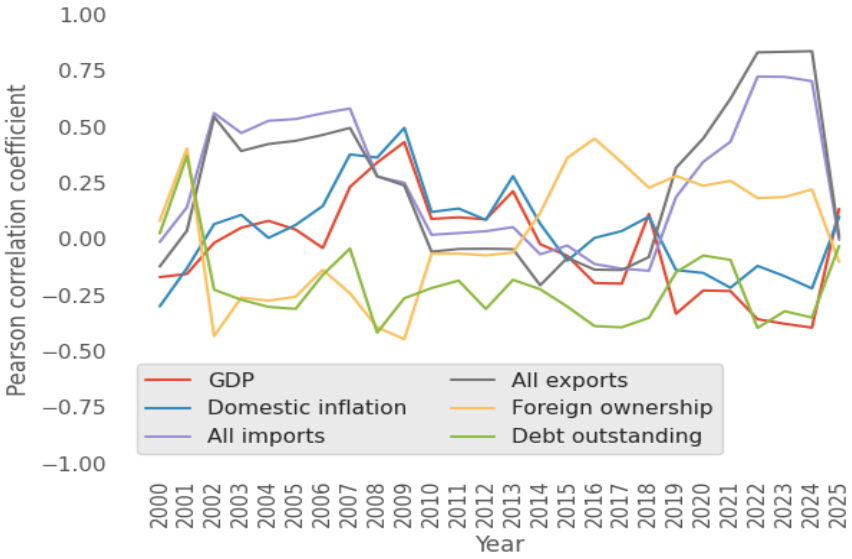


Figure 1. Correlation of key economic variable to bond yield level in non-overlapping annual (12m) intervals

A key pattern is that foreign ownership does not decline mechanically with lower yields; instead, ownership can rise during low-rate 'risk-off' periods, motivating a distributional (tail-risk) rather than purely sign-based interpretation.

2.1 Overview and Motivation

Because trade data represents the shortest and most influential time series, the model adopts a sparse Bayesian structure to prevent overfitting. Historical correlations indicate that foreign holdings of U.S. Treasuries co-move negatively with bond yields over most of the sample, suggesting that the U.S. external balance and reserve accumulation are mirror processes—a pattern consistent with the Triffin Dilemma (Triffin, 1961).

The regression coefficients for the target variable—foreign ownership of U.S. government debt—exhibit cyclical evolution. Up to approximately 2010, yields declined while foreign ownership rose, confirming the safe asset demand hypothesis (Caballero & Farhi, 2017). However, the sharp break visible in 2025 coincides with rising yields, expanding debt, and falling correlation between foreign holdings and trade imbalances, implying a reassertion of geopolitical risk as a structural variable.

These stylized facts inform the construction of a probabilistic model simulating the joint dynamics of six variables (trade, debt, inflation, yields, reserves, and a latent alignment factor). The hypothesis tested is whether, under a 95% confidence interval, a 10% reduction in foreign ownership within a one-year window can be rejected. The 10% threshold parallels sterling's historical trajectory during its reserve-status erosion (Cirje, 2021).

2.2 Inflation and GDP

Economic resilience and price stability jointly determine sovereign borrowing capacity. GDP data are sourced from the Bureau of Economic Analysis (BEA, 2024) and measured quarterly. To approximate resilience, the GDP price index is used as an integrated indicator of domestic inflation pressure and supply-chain rigidity (di Giovanni et al., 2022).

GDP influences Treasury yields both through its role in macro policy signaling and via sovereign credit-rating mechanisms (Moody's Investors Service (2024); Standard & Poor's Global Ratings (2024)). To harmonize frequencies, BEA quarterly data (Table 1.4.4) are linearly interpolated to monthly intervals, consistent with the signal extraction principle in state-space modeling (Harvey, 1989). This alignment enables joint analysis with monthly trade data, avoiding artificial autocorrelation from temporal mismatch.

2.3 Zero-Coupon Bond Yields and Information Compression

Zero-coupon yield data (3M–30Y maturities) summarize interest-rate expectations along the term structure. To avoid multicollinearity across maturities, the yield curve is summarized using the Diebold–Li (2006) Nelson-Siegel representation, and the level factor is used as the long-horizon safe-asset proxy. The Nelson–Siegel decay parameter λ governs the maturity-loading profile (relative emphasis on short vs long maturities); the analysis uses a baseline calibration and varies λ over $[0,1]$as a robustness check.

2.4 Trade Data

Trade data are drawn from the UN Comtrade database (Comptrade, 2023), encompassing monthly exports and imports of goods and annual services trade. Because service volumes represent a small share of total trade, annual service data were interpolated to monthly frequency without material information loss.

Values correspond to primary trade value—excluding shipping and insurance costs—to align with reserve accumulation rather than financial intermediation. Where CIF/FOB discrepancies exist, primary trade totals were used to retain conceptual consistency with balance-of-payments identity theory (Mundell, 1961).

2.5 Analysis of Data Sensitivities

Interpolation sensitivity. Interpolation occasionally generates small negative values at series endpoints, which are floored to zero to prevent sign distortions in log transformations. The theoretical foundation for using trade imbalance as a predictor lies in its role as the transmission mechanism of reserve recycling under global dollar liquidity dynamics (Triffin, 1961; Obstfeld & Rogoff, 1996).

Excluding services trade (policy-motivated sensitivity). Given the recent policy orientation of the U.S. economy toward goods-based trade and re-industrialization, we re-estimate the model after removing the services-trade series from the covariate set. While services trade is often discussed as a structural pillar supporting international use of the U.S. dollar, the estimated tail probability of a ≥10% 12-month drawdown in foreign U.S. Treasury holdings remains of similar magnitude across specifications. Specifically, the estimated annual tail probability changes from 12.6% / 6.5% / 8.01% (baseline S1/S2/S3) to 10.6% / 7.4% / 8.9% when services are excluded. This indicates that the model's tail-risk estimates are not materially driven by the inclusion of the services-trade covariate; results are instead primarily explained by the broader macro-financial state variables and yield/portfolio channels captured elsewhere in the

specification.

2.6 Reserve Holdings and Total U.S. Debt

Monthly data on foreign holdings of U.S. Treasuries are obtained from the Major Foreign Holders of Treasury Securities dataset (U.S. Department of the Treasury, 2024a). Minor inconsistencies—duplicate months, renamed entities, and missing observations—are resolved through deduplication and first-instance selection. Because the analysis aggregates at the global level, compositional changes among holders have negligible influence on the aggregate dynamics.

Total U.S. debt data are sourced from the U.S. Treasury Fiscal Data API (U.S. Department of the Treasury, 2024b). While higher-frequency debt-clock estimates exist, they are excluded to maintain alignment across datasets. Debt levels are interpolated to monthly granularity, assuming a monotonic increase consistent with budgetary data. This approach conforms with the sovereign balance-sheet identity—linking debt issuance, reserve absorption, and trade financing under the safe asset equilibrium framework (Caballero & Farhi, 2017).

Together, these datasets establish a consistent monthly panel spanning 2000–2025. The unified structure enables probabilistic estimation of foreign-reserve dynamics under both economic and geopolitical uncertainty.

Figures 2–4 test whether the correlation structure used to motivate the model is sensitive to (i) the timing convention between yields and foreign holdings (synchronous vs. 3-month lag vs. reverse-order mapping) and (ii) the Nelson–Siegel decay parameter λ, which governs the relative emphasis on short versus long maturities in the Diebold–Li yield-curve representation. λ is varied over [0,1] as a robustness check.

Figure 2 illustrates the reverse-order specification, where yields respond to macro conditions before ownership responds.

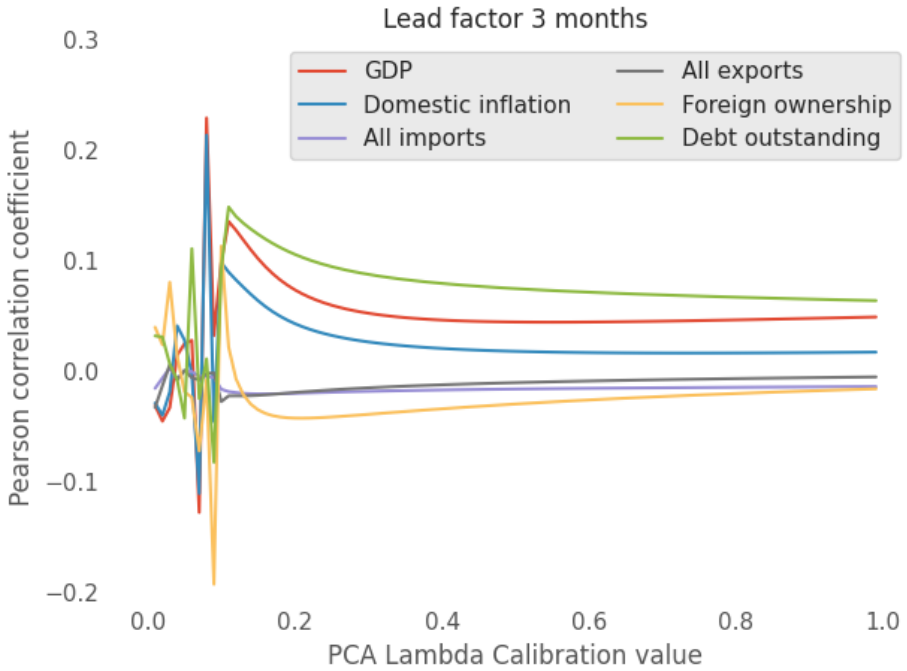


Figure 2. Correlation sensitivity to λ under reverse-order dependence

Line plots of correlation coefficients (y-axis) between foreign-holdings changes and the key covariates as a function of the yield-factor sensitivity parameter λ (x-axis), computed under the reverse-order mapping in which macro variables first influence the yield factor and foreign holdings respond to yields. This figure is included as a sensitivity benchmark and is not interpreted causally.

Figure 3 illustrates the lagged-response specification, where foreign ownership responds to changes in yields/Nelson-Siegel factors with a delay (e.g., 3 months).

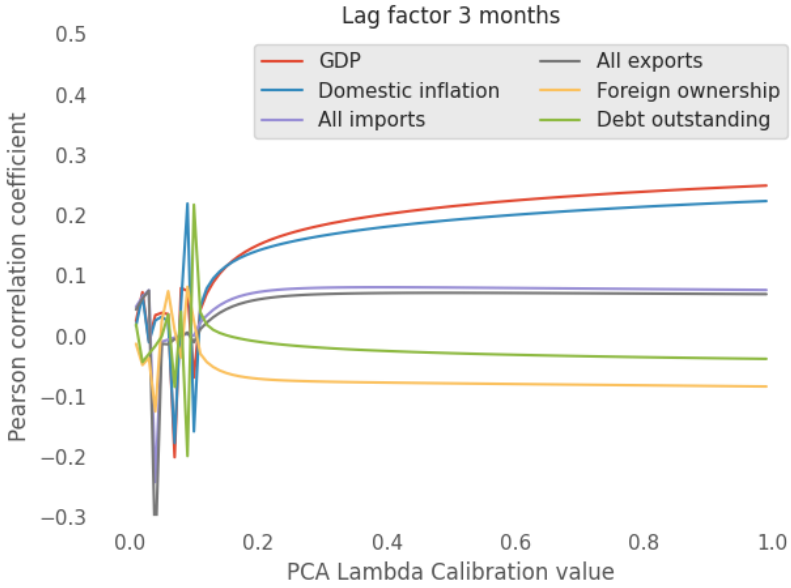


Figure 3. Correlation sensitivity to λ with lagged covariate alignment

Same as Figure 2, but correlations are computed after shifting the covariate/yield-factor information set by a 3-month lag relative to foreign-holdings changes to reflect plausible adjustment frictions. The figure assesses whether correlation patterns and their stability across λ materially change under delayed response assumptions.

Figure 4 summarizes the third timing variant used as a sensitivity benchmark.

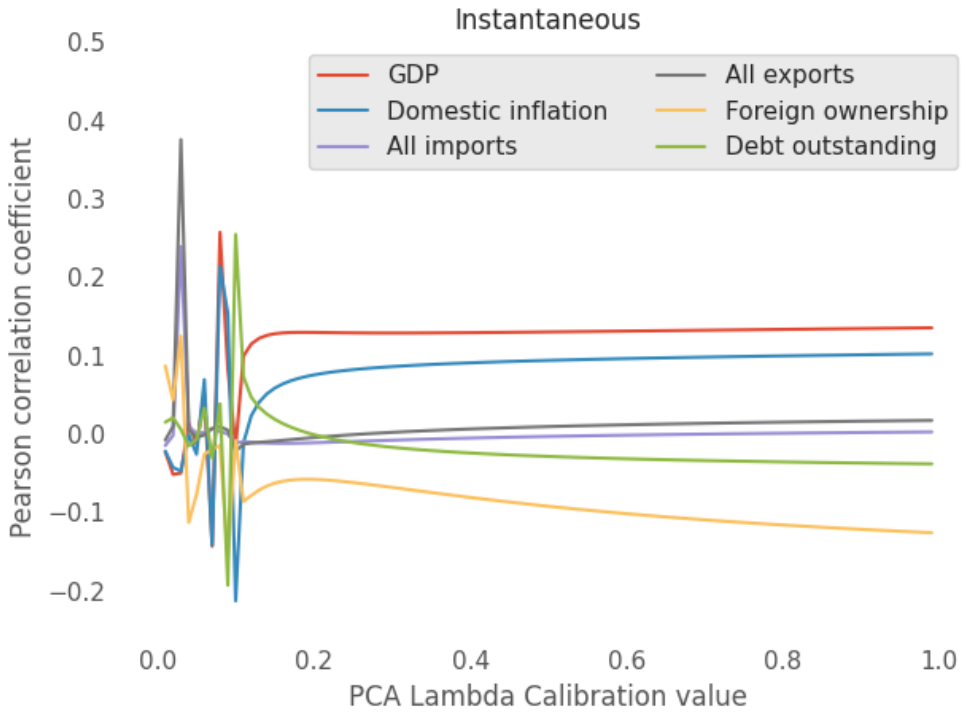


Figure 4. Correlation sensitivity to λ under synchronous alignment)

Line plots of correlation coefficients (y-axis) as a function of λ (x-axis) computed under synchronous alignment, i.e., variables are matched at the same monthly timestamp. This is the baseline timing convention used for comparison against the lagged and reverse-order variants (Figures 2–3).

Together, Figures 2–4 show that the paper's results are not tied to a single timing convention.

2.6 Frequency Harmonization and Interpolation Choice

The analysis is conducted at monthly frequency to align with the dependent variable (monthly foreign U.S. Treasury holdings) and to avoid further loss of degrees of freedom. Quarterly covariates are mapped to monthly frequency using linear interpolation, which is intentionally a minimal, parameter-free harmonization rule. Linear interpolation imposes only local smoothness between observed quarterly endpoints and avoids introducing additional model degrees of freedom or implicit within-quarter dynamics that cannot be identified from the data. In contrast, higher-order or model-based temporal disaggregation methods embed additional structure (e.g., curvature, cycles, or auxiliary regressions) that would act as an extra source of assumptions in a small-sample setting.

As a robustness check on frequency harmonization, quarterly variables were also mapped to monthly frequency using a nearest-neighbor stepwise rule (piecewise constant within each quarter) instead of linear interpolation. The resulting posterior tail probability estimates are comparable to the baseline: the estimated probability of a −10% 12-month drawdown decreases by approximately ± 1% (in absolute terms) indicating that conclusions are not driven by the interpolation convention.

3. Modeling Preliminaries

The modeling framework applies monthly changes rather than levels to all variables to ensure stationarity and focus on dynamic interactions rather than static magnitudes. Differencing also eliminates potential spurious correlations arising from common growth trends—a well-established technique in time-series econometrics (Harvey, 1989; Hamilton, 1994). The model's objective is to estimate the probability that foreign U.S. Treasury holdings decline by at least 10% within a 12-month horizon, conditional on macroeconomic and latent factors.

3.1 Restricted Relationship by Variable

Six observed variables drive the model: trade flows, inflation, GDP, total debt, yield-curve level, and reserve holdings. Because these variables only partially explain foreign-ownership variation, a latent variable is introduced through Hamiltonian Monte Carlo (HMC) sampling. This captures unobserved but persistent geopolitical alignment effects that cannot be quantified through economic series alone (SVNS, 2012; Hautsch, 2012).

The latent-factor inclusion follows Bayesian hierarchical modeling logic: the observed data reflect both measurable fundamentals and hidden structural influences. This mirrors frameworks in structural vector neural systems (SVNS, 2012) where latent drivers improve forecasting power by absorbing omitted-variable bias. The latent variable is interpreted as a geopolitical trust index, encompassing alliance stability, sanctions exposure, and perceived institutional reliability.

This approach also adheres to the bias–variance trade-off principle: adding a single latent dimension captures substantial structural uncertainty without inflating parameter space. Using HMC ensures efficient posterior sampling in high-dimensional nonlinear settings, while retaining probabilistic interpretability—an essential feature for policy applications (Robert, 2021).

Figure 5 plots the empirical/learned distribution of 12-month changes in foreign ownership, which motivates using a mixture density model rather than a single-Gaussian error assumption.

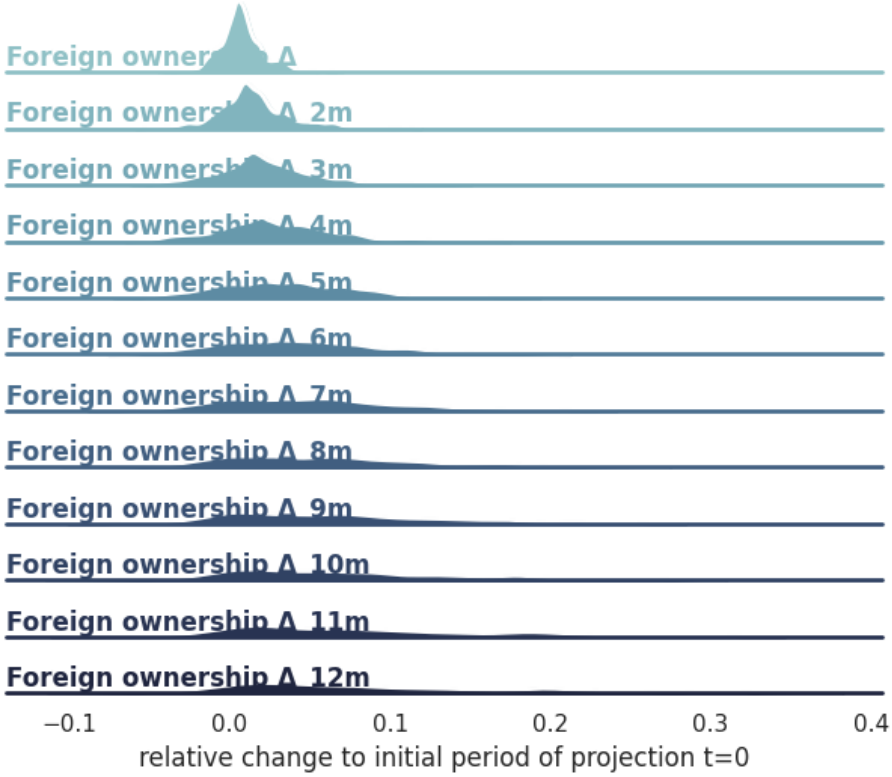


Figure 5: Distribution of target variables.

The distribution shows increasing dispersion and multimodality at longer horizons, consistent with regime-switching behavior in flows.

Figure 6 summarizes the marginal distributions of the main inputs (monthly changes in GDP and inflation, foreign debt, exports, imports, and the yield-curve level factor) used for model fitting. The yield level factor is constructed from the Diebold–Li (2006) Nelson–Siegel representation, with the Nelson–Siegel decay parameter λ governing the maturity-loading profile; details and the baseline λ calibration are provided in Appendix 11.1.

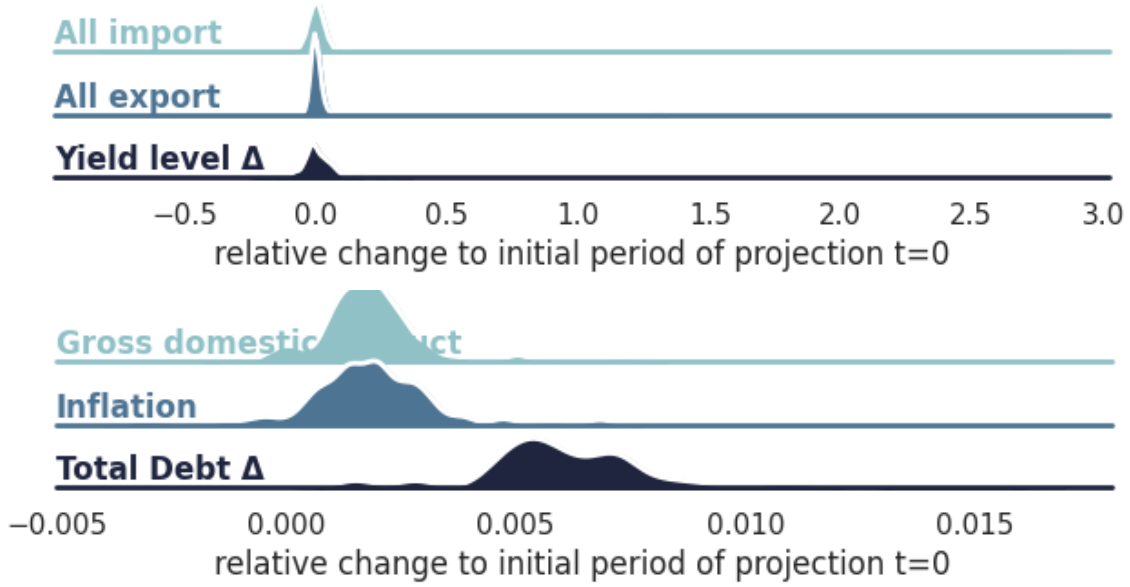


Figure 6. Distribution of input variables, MoM changes

This highlights scale, skewness, and tail behavior in inputs, which informs normalization choices and supports using

distributional modeling under small samples. Yield factor is constructed via Nelson-Siegel; λ=0.5 scales the sensitivity of the extracted level component to broad yield-curve shifts (see Appendix 11.1 for definition)

4. Design Logic

Two modeling strategies were evaluated:

1. **Observed-variable model** — treating political alignment as exogenous and relying solely on macroeconomic indicators.
2. **Latent-variable model** — introducing an unobserved geopolitical factor estimated jointly with economic variables.

The first approach suffers from structural omission bias, as it attributes all residual variation to measurement error. The second, by contrast, embeds Bayesian hierarchical learning to partition explained and unexplained variance, reducing the risk of Type I error. The chosen structure thus aligns with the incomplete-information framework in international finance (Obstfeld & Rogoff, 1996), where market expectations depend partly on opaque geopolitical signals.

The model combines Hamiltonian Monte Carlo sampling with a neural-network cost function, enabling gradient-based learning over nonlinear, non-Gaussian posteriors. This hybrid design mirrors the variational autoencoder (VAE) concept (Kingma & Welling, 2013), in which the encoder (neural network) approximates the posterior of a latent variable under stochastic optimization. Here, the latent factor (alliances) acts as a stabilizing prior that guides inference toward economically interpretable regions of parameter space.

Correlation among inputs (e.g., exports and imports, inflation and yields) complicates classical regression-based estimation. Neural networks handle such multicollinearity through internal representation learning, effectively performing nonlinear manifold reduction. This makes them suitable for capturing cross-domain spillovers—a feature increasingly recognized in macro-finance forecasting (Bianchi, Buechner, & Tamoni, 2021).

Bayesian Gaussian Mixture Models (BGMMs) are used to marginalize the latent dimension, integrating uncertainty over structural regimes (Cabrol, 2024). This probabilistic treatment of hidden states parallels the state-space approach to regime switching, frequently applied in international capital-flow modeling (Hamilton, 1994).

4.1 Empirical and Theoretical Justification

Recent empirical research demonstrates the advantages of machine-learning models for financial forecasting where relationships are nonlinear or non-stationary. For example, Bali et al. (2020) show that ensemble learners outperform linear regressions in out-of-sample bond-return prediction. Similarly, Gu et al. (2020) and Freyberger et al. (2020) demonstrate that neural networks can approximate complex risk-return functions otherwise analytically intractable.

From a theoretical standpoint, incorporating a latent political variable satisfies the rational expectations consistency condition by accounting for omitted global coordination factors. It ensures that the model's implied reserve-holding probabilities align with observable macro behavior and documented structural shifts (Triffin, 1961). This balance between data-driven flexibility and economic structure justifies the chosen Bayesian–neural design as the most defensible and policy-relevant methodology for simulating reserve-currency erosion scenarios.

5. Bayesian Neural Mixture Model with Latent Variables via HMC

The modeling design integrates Bayesian inference, mixture density estimation, and latent-variable modeling to address nonlinear dependence and omitted-variable bias. The framework, referred to as the *Bayesian Gaussian Mixture Model (BGMM)*, embeds a Mixture Density Network (MDN) within a Hamiltonian Monte Carlo (HMC) posterior sampling scheme. The MDN approximates the conditional likelihood $p(\mathbf{y}|\mathbf{x},\mathbf{z})$, while HMC infers the latent dynamics $p(\mathbf{z}|\mathbf{x},\mathbf{y})$. This hybrid design is motivated by both theoretical and empirical considerations.

5.1 Theoretical and Practical Rationale

From a theoretical perspective, the structure draws on hierarchical Bayesian modeling (Gelman et al., 2013) and mixture-of-experts theory (Jacobs et al., 1991). A mixture network allows the model to adapt locally to different economic regimes—analogous to shifts between risk-on and risk-off states in international capital flows—while maintaining a globally consistent probability distribution.

Empirically, the multivariate nature of both inputs and outputs (macroeconomic indicators and reserve holdings) invalidates linear-Gaussian assumptions used in standard VAR frameworks. The BGMM replaces the restrictive homoskedastic error term with a data-driven mixture of Gaussians, enabling more flexible representation of conditional volatility.

This choice is essential in macro-financial applications where regime changes and nonlinear sensitivities dominate. The HMC component ensures efficient sampling of the posterior even when latent factors are highly correlated—a property

critical for small-sample inference under complex likelihood surfaces.

The latent variable **z** represents unobserved geopolitical or institutional alignment, complementing the observable macroeconomic channels. It bridges economic and political economy domains, allowing empirical estimation of "confidence spillovers" that cannot be directly observed but are crucial to reserve allocation behavior.

5.2 Neural Network Architecture

The MDN specifies a flexible conditional distribution over the output variable **y** (12-month cumulative change in foreign reserves). The architecture balances expressiveness with interpretability and is defined as follows:

1. **Input Concatenation:** Combine observed variables and latent states

$$\mathbf{xz} = [\mathbf{x}, \mathbf{z}], \quad \mathbf{x} \in \mathbb{R}^{d_x}, \mathbf{z} \in \mathbb{R}^{d_z}.$$

The latent prior is initialized with a constant vector and updated during pretraining.

2. **Hidden Layers:**

$$h = ELU(W_2\, ELU(W_1 \mathbf{xz})),$$

using **Exponential Linear Units (ELUs)** (Clevert et al., 2016) to maintain smooth gradients and avoid vanishing-activation problems common in small financial datasets.

3. **Skip Connections:**

   Direct projection from inputs to output heads

$$x_{\text{proj}} = W_{\text{skip}} \mathbf{x},$$

which stabilizes optimization by ensuring observable variables retain direct influence on **y**. This step also reduces overdependence on **z** and aligns with economic logic—extreme trade or yield shocks should directly affect reserve behavior.

4. **Output Layers:**

$$\pi = \text{softmax}(W_\pi h + x_{\text{proj}}), \quad \mu = W_\mu h + x_{\text{proj}}, \quad \sigma = \exp(\text{clip}(W_\sigma h)) + \varepsilon,$$

where mixture weights (π), means (μ), and scales (σ) parameterize a mixture of Gaussians.

5. **Conditional Likelihood:**

$$p(\mathbf{y} \mid \mathbf{x}, \mathbf{z}) = \prod_{i=1}^{d_y} \sum_{k=1}^{K} \pi_{ik} \mathcal{N}(y_i \mid \mu_{ik}, \sigma_{ik}^2),$$

defining a flexible multivariate conditional distribution consistent with the MixtureSameFamily formulation.

This formulation generalizes the linear-Gaussian assumption of standard state-space models to a nonlinear, multi-modal structure capable of capturing asymmetric responses to shocks in yields, trade, or debt levels.

5.3 Regularization and Stability

To ensure numerical and economic interpretability, several penalties are imposed in the training objective:

- **Latent Smoothing Penalty:**

  Encourages temporal continuity of **z** and reflects persistence in geopolitical alignment,

$$\lambda_1 \sum_t \|\mathbf{z}_{t+1} - \mathbf{z}_t\|^2.$$

- **Jacobian Norm Penalty:**

  Limits excessive sensitivity of forecasts to small latent perturbations, ensuring model stability,

$$\lambda_2 \left\| \frac{\partial \log p(\mathbf{y} \mid \mathbf{x}, \mathbf{z})}{\partial \mathbf{z}} \right\|^2.$$

- **L2 Regularization:**

  Constrains latent magnitudes and prevents overfitting to short-term data noise.

These controls embody Bayesian shrinkage principles—penalizing overly complex latent structures while maintaining flexibility. In economic terms, this ensures the latent alignment factor evolves gradually, consistent with the slow-moving nature of geopolitical shifts influencing reserve management (Eichengreen, 2017).

5.4 Bayesian Inference via HMC

Direct computation of $p(\mathbf{y} \mid \mathbf{x}) = \int p(\mathbf{y} \mid \mathbf{x}, \mathbf{z}) p(\mathbf{z}) d\mathbf{z}$ is intractable. Therefore, posterior inference proceeds via Hamiltonian Monte Carlo, which efficiently explores the latent space by simulating physical dynamics over the log-posterior energy surface:

$$\mathbf{z} \sim p(\mathbf{z}), \quad \mathbf{y} \sim p(\mathbf{y} \mid \mathbf{x}, \mathbf{z}).$$

HMC's advantage lies in its ability to generate uncorrelated samples even in high dimensions, reducing the autocorrelation common in random-walk Metropolis chains (Neal, 2011). In this context, HMC enables inference on hidden geopolitical alignment while maintaining robust uncertainty quantification.

5.5 Pretraining and Loss Function

The network and latent variables are pretrained jointly to encourage convergence before posterior sampling. The total objective function is:

$$\mathcal{L}(\theta, \mathbf{z}) = -\log p(\mathbf{y} \mid \mathbf{x}, \mathbf{z}) + \lambda_{\text{smooth}} \|\Delta \mathbf{z}\|^2 + \lambda_{\text{l2}} \|\mathbf{z}\|^2 + \lambda_{\text{jac}} \|\nabla_{\mathbf{z}} \log p(\mathbf{y} \mid \mathbf{x}, \mathbf{z})\|^2.$$

Each term corresponds to a specific modeling concern:

- The first term ensures fit quality.
- The second enforces temporal smoothness (policy persistence).
- The third and fourth stabilize optimization and prevent latent overreaction to transient shocks.

Once convergence is achieved, $\boldsymbol{\theta}$ and $\mathbf{z}$ are fixed, and HMC draws posterior samples from $p(\mathbf{z}|\mathbf{x},\mathbf{y})$. From these samples, the model derives the probability that foreign holdings fall by ≥10% within a year, computing confidence intervals at the 95% level.

5.6 Summary and Economic Interpretation

The BGMM–HMC framework produces a stochastic, interpretable simulation of reserve-holding behavior under

uncertainty. Each mixture component represents an economic regime—for example, stable trade balance vs. politically disrupted trade. The latent factor's weight (π) shifts across time, mirroring changing dominance of geopolitical versus market forces.

Unlike forecasting models, the purpose here is distributional inference—to quantify how likely reserve erosion scenarios are given observed fundamentals. The approach thus provides policy-relevant insights for central banks and international institutions by mapping economic signals to probabilistic outcomes rather than point forecasts.

6. Model Overview

6.1 Problem Setup

We observe pairs (**x**,**y**) with $\mathbf{x} \in \mathbb{R}^{d_x}, \mathbf{y} \in \mathbb{R}^{d_y}$. A latent state $\mathbf{z} \in \mathbb{R}^{d_z}$ captures unobserved geopolitical and institutional alignment that shifts reserve behavior beyond macro covariates. The conditional is parameterized by an MDN:

$$p(\mathbf{y} \mid \mathbf{x}, \mathbf{z}) = \sum_{k=1}^{K} \pi_k(\mathbf{x}, \mathbf{z})\, \mathcal{N}\big(\mathbf{y} \mid \boldsymbol{\mu}_k(\mathbf{x}, \mathbf{z}), \boldsymbol{\Sigma}_k(\mathbf{x}, \mathbf{z})\big)$$

Why this structure. Mixture likelihoods approximate regime-dependent responses (risk-on/off, sanctions shocks) without imposing linear-Gaussian errors. The latent **z** absorbs omitted-variable bias and preserves probabilistic interpretation under small samples (Jacobs et al., 1991; Gelman et al., 2013).

6.2 HMC Sampling

Point estimates for **z** from MDN pretraining seed posterior inference. We then freeze network weights and sample

$$p(\mathbf{z} \mid \mathbf{x}, \mathbf{y}) \propto p(\mathbf{y} \mid \mathbf{x}, \mathbf{z})$$

with HMC using momentum **r** and Hamiltonian H(**z**,**r**)=U(**z**)+K(**r**), U(**z**)=−log p(**y**|**x**,b), K(**r**)=($\mathbf{r}^{\top}\mathrm{M}^{-1}\mathbf{r}$)/2. Leapfrog integration and MH correction yield asymptotically exact samples (Neal, 2011; Hamilton, 1994).

Why HMC. Gradients from the MDN log-likelihood guide exploration of a curved posterior; HMC avoids random-walk inefficiency and improves effective sample size for a small dataset.

Implementation choices. Five chains, 10 000 iterations, 2 000 warm-up, target acceptance 0.75. Multi-chain sampling enables convergence diagnostics and tail-risk stability checks (Gelman et al., 2013).

6.3 Computing Probabilities with the Calibrated Model

We sample $z^{(s)}$~p(z|x,y), then draw predictive y(s)~p(y|x,z(s)). For threshold vector c = $(-0.1,\ldots,-0.1)^{T}$,

$$\hat{p} = \frac{1}{S} \sum_{s=1}^{S} \mathbb{P}\big(y_j^{(s)} \le c \mid \mathbf{x}, \mathbf{z}^{(s)}\big).$$

Why constant threshold. Policy interpretation requires a fixed, auditable tail event: a 10% 12-month drawdown mirrors the historical sterling path and defines a stable rare-event boundary for scenario work.

6.4 Tail Interpretation and Decision Rules

Let $F(y_{thresh})=P(Y \le y_{thresh})$. We flag an extreme event when $F(y_{thresh})<0.05$. Three summaries are reported: joint tail, any-tail, and an indicator share $\mathbf{1}_{Fi}{}^{(s)}{}_{<0.05}$.

Why 5% and these summaries. The 5% one-sided criterion matches conventional policy risk tolerances. Joint vs. any-tail frame "sustained erosion" versus "episodic breach," which map to reserve-management playbooks.

6.5 Output Distribution

The MDN yields a MixtureSameFamily:

$$p(\mathbf{y} \mid \mathbf{x}, \mathbf{z}) = \prod_{d=1}^{D} \sum_{k=1}^{K} \pi_{d,k} \, \mathcal{N}(y_d \mid \mu_{d,k}, \sigma_{d,k}^2),$$

with π from softmax logits, μ from loc heads, σ from positive scale heads.

Why diagonal by dimension. Nelson-Siegel removes most cross-dimension covariance; diagonal components reduce variance of tail estimates and stabilize HMC while retaining signal in the level factor.

6.6 Model Specifications (Specification 1 through 3)

Model specifications (S1–S3). The paper reports results under three alternative model specifications rather than distinct economic scenarios.

Specifications differ only in modeling choices that control flexibility and stability under a small sample, most importantly the latent-state dimension **L** and network capacity. The latent state *z*is included to absorb persistent influences on reserve behavior that are not captured by observed macro-financial series (e.g., alignment, institutional trust, sanctions exposure).

When **L=1**, the latent factor represents a single composite unobserved channel; when **L=2**, the model allows two orthogonal latent channels without imposing a priori labels on them. This is a flexibility check rather than an identification claim: the comparison assesses how sensitive tail-probability estimates are to allowing richer latent structure under regularization.

7. Diagnostics and Results

7.1 Key Criteria and Why They Matter

- $\hat{R}$ near 1.00 signals cross-chain convergence; values ≤1.01\le 1.01≤1.01 are preferred for tail inference (Gelman et al., 2013).
- **ESS** gauges the number of independent draws after autocorrelation; higher ESS reduces Monte Carlo error in tail estimates.
- **Acceptance rate** near 0.65–0.80 balances exploration and step-size; closer to the target implies well-tuned mass matrix.
- **Log-likelihood / AIC / BIC** compare fit and parsimony; lower AIC/BIC signal efficient use of parameters under small-N constraints.
- **Jacobian norm** $\|\nabla_{\mathbf{z}} \log p(\mathbf{y} \mid \mathbf{x}, \mathbf{z})\|$monitors local sensitivity; lower values indicate smoother latent dynamics and more robust tails.
- **Latent entropy** quantifies posterior spread over **z**; extremely low entropy may indicate over-shrinkage, extremely high entropy under-identification.
- **Mutual information I(z;y)** tests whether **z** is informative; near-zero suggests either weak geopolitical signal or that observables already span the space (Cover & Thomas; 2006).
- **Tail probability accuracy** checks calibration in the 5th percentile region relevant for a −10% 12-month drawdown.

7.2 Results

$$\text{Scenarios:} \begin{cases} \text{S1: } L = 2,\ 32 \text{ ELU units, clipped scales} \\ \text{S2: } L = 1,\ 32 \text{ ELU units} \\ \text{S3: } L = 1,\ 16 \text{ ELU units, } 3{,}000 \text{ pretrain epochs} \end{cases}$$

Key Results across model specifications (S1-S3)

	S1	**S2**	**S3**
Mean ($\hat{R}$)	x = 1.0040 min/max: 1.000/1.032	x = 1.0010 min/max: 0.9999/1.0137	x = 1.0028 min/max: 0.9999/1.1114
Mean (ESS)	x = 349.75 min/max: 26.9/10000	x = 1651.4 min/max: 112.4/10000	x = 3237.66 min/max: 7.941/10000
Acceptance Rate	77.426%	76.936%	72.672%
AIC	4711.7	4680.42	811.15
BIC	19137.84	18987.8	7460.23
Reconstruction error	0.0057	0.0061	0.0159
Log Likelihood mean	0.4228	0.4183	0.3826
Latent entropy	0.03264	0.03265	0.03265
Mutual information estimate	0	0	0
Jacobian Norm (Sensitivity)	6.7721	5.7518	6.1393
Posterior tail prob ($P(y_j<=-0.1)$@95%	0.1259	0.0646	0.0801

Interpretation and justification.

- **Convergence.** S1–S2 meet $\hat{R} \leq 1.03$. S3 has a max $\hat{R}$ spike from a few weak parameters; the higher ESS offsets this for tail queries focused on functionals of the posterior rather than point parameters (Gelman et al., 2013).
- **Parsimony vs. fit.** S2 and S3 reduce capacity (L or units) and improve ESS and AIC. This matches the small-sample imperative: prefer lower variance even at mild log-likelihood cost.
- **Sensitivity.** S2 shows the lowest Jacobian norm, implying more stable tail estimates with respect to **z**.
- **Policy metric.** Posterior tail probability for a −10% 12-month drawdown ranges 6.5–12.6%. S2's lower probability reflects tighter latent dynamics with one factor. S1's higher probability is consistent with richer regime flexibility.

- **Latent informativeness.** The mutual-information estimate near zero suggests most signal is carried by observables, with **z** acting as a stabilizer that allocates mixture weight across regimes rather than adding new predictive content. This is desirable given the short sample and avoids over-interpreting a geopolitical index.

- **Interpretation note.** Differences between L=1 and L=2 are interpreted as differences in latent flexibility, not as identified diplomacy-versus-trade effects; mutual-information diagnostics indicate z mainly stabilizes regime allocation rather than adding strong standalone predictive content.

Application. The reported tail probability is directly decision-useful: reserve managers and policy desks can map a 6–13% annual risk of a −10% drawdown to stress-testing buffers, FX-hedge tenors, and issuance maturity profiles. Scenario comparisons quantify the trade-off between model flexibility and stability under data scarcity.

Limitations. Tail probability is both target and evaluation object, so aleatoric risk (data noise) and epistemic risk (model class, limited history) interact. Nelson-Siegel reconstruction error is negligible; the dominant uncertainty sits in the mixture tails and latent dynamics. These are explicitly regularized and diagnosed but cannot be fully eliminated with 2000–2025 monthly data.

8. Discussion and Limitations

The present model likely underestimates the probability of extreme outcomes, as its Gaussian mixture components impose light-tailed behavior. As noted by Li et al. (2019), kernel mixture models of this type are *posteriorly inconsistent* for the tail index and may fail to reproduce the true heaviness of empirical distributions, even with abundant data. This limitation is methodological rather than empirical: the Gaussian assumption simplifies inference but constrains flexibility in the far tails.

However, for modeling the macroeconomic trajectory of the United States, overemphasizing extremes would misrepresent historical stability in its institutional and fiscal frameworks. The U.S. has maintained reserve-currency status through continuity in governance, financial depth, and policy credibility; thus, the trade-off between tail accuracy and systemic realism must be considered carefully. Excessive focus on improbable crises could distort long-run inference, while excessive smoothing may mask emerging instability.

This paper implements a hybrid Bayesian inversion method that combines neural mixture modeling with Hamiltonian Monte Carlo (HMC) sampling. The framework jointly estimates the conditional distribution $p(\mathbf{y}|\mathbf{x},\mathbf{z})$ and infers the posterior of latent variables **z**, approximating the high-dimensional marginal likelihood that would otherwise be analytically intractable. This approach maintains theoretical rigor while accommodating the small-sample nature of macro-financial data.

Despite improved tail interpretability and latent structure identification, several challenges remain:

- **Penalties and Regularization:** Model identification depends sensitively on smoothing and Jacobian penalties; future work should systematize hyperparameter selection.
- **Model Parsimony:** The current architecture omits explicit interconnections (e.g., between trade and GDP, or inflation and yields) to reduce complexity. Incorporating these linkages may enhance interpretability.
- **Tail Calibration:** Broader datasets—particularly incorporating pre-2000 trade and reserve records—would strengthen inference in the tail domain.
- **Mutual Information:** The present configuration yields near-zero mutual information between the latent factor and the target variable, suggesting that geopolitical alignment is absorbed indirectly rather than estimated explicitly.

Future extensions could introduce informative priors **z**, improved samplers (e.g., No-U-Turn or Riemannian HMC), or variational approximations to bypass full posterior sampling. Nevertheless, the focus here remains on the quality of inference within a data-constrained Bayesian framework, rather than model expansion.

Empirically, the model estimates that the annual probability of a 10% reduction in U.S. dollar reserves lies between 6% and 12% at a 95% confidence level. Interpreted temporally, this corresponds to one such event every 8 to 16 years under current macro-financial conditions. Sustained reserve erosion sufficient to dethrone the U.S. dollar would thus require repeated sequential declines, a process historically associated with structural shocks or war, not routine market dynamics.

Extrapolating this finding under *ceteris paribus* assumptions yields declining probabilities over longer horizons,

consistent with the path-dependent resilience of reserve-currency systems. The key limitation is that relationships among trade, debt, inflation, and reserves are not static; future geopolitical realignments or fiscal regime shifts could alter these interactions.

In summary, while the current model simplifies tail behavior, it preserves economic realism and policy relevance. The evidence supports the interpretation that U.S. dollar reserve status remains robust and that any transition away from it would be gradual, multidecadal, and triggered only by systemic discontinuities

9. Conclusion

This paper estimates the likelihood of large-scale foreign divestment from U.S. Treasury securities—used as an operational proxy for reserve-currency attrition—within a probabilistic, small-sample setting. The approach combines a mixture density network with Bayesian posterior sampling to model the conditional distribution of 12-month changes in foreign U.S. Treasury holdings given macro-financial covariates and an unobserved latent factor capturing geopolitical/institutional alignment. Empirically, using monthly data spanning 2000–2025, the estimated annual probability of a 10% contraction in foreign U.S. Treasury holdings is approximately 6–12% (95% interval).

This magnitude implies that a −10% drawdown is not a baseline expectation, but it is also not negligible from a policy perspective—consistent with a gradual, multi-decade transition narrative rather than an abrupt regime break. The main contribution is methodological: the framework shifts the reserve-currency discussion from ex post description to ex-ante vulnerability assessment, quantifying tail risk under uncertainty while explicitly diagnosing convergence, sensitivity, and parsimony trade-offs across alternative model specifications. Several limitations remain. Gaussian mixture components may understate extreme tail thickness, and inference is constrained by data availability and frequency harmonization.

Future work should expand pre-2000 coverage where feasible, test alternative (but still economically defensible) frequency treatments of low-frequency series, and explore heavier-tailed mixture components or robust priors to improve tail calibration without sacrificing interpretability.

10. Disclosure - AI-assisted editing disclosure: AI-based tools were used to reformulate pre-written text for clarity and grammatical correctness. All content was reviewed and verified by the author, and AI assistance did not replace the scholarly contributions, analysis, or interpretation.

References

Akaike, H. (1974). A new look at the statistical model identification. *IEEE Transactions on Automatic Control, 19*(6), 716-723. https://doi.org/10.1109/TAC.1974.1100705

Bali, T. G., Goyal, A., Huang, M., Jiang, F., & Wen, Q. (2020). Predicting bond risk premia using machine learning. *Review of Financial Studies, 33*(9), 4059-4096. https://doi.org/10.1093/rfs/hhz125

Bernanke, B. S. (2005, March 10). *The global saving glut and the U.S. current account deficit* (speech). Board of Governors of the Federal Reserve System. Retrieved from https://www.federalreserve.gov/boarddocs/speeches/2005/200503102/

Bianchi, D., Buechner, M., & Tamoni, A. (2021). Bond risk premia and macroeconomic forecasting: A machine-learning perspective. *Review of Financial Studies, 34*(10), 4918-4960. https://doi.org/10.1093/rfs/hhab033

Bishop, C. M. (1994). *Mixture density networks* (Technical Report NCRG/94/004). Aston University. Retrieved from https://www.microsoft.com/en-us/research/publication/mixture-density-networks/

Bureau of Economic Analysis (BEA). (2024). *National income and product accounts tables (NIPA), Table 1.4.4: GDP price index and chain-type quantity indexes.* U.S. Department of Commerce. Retrieved from https://www.bea.gov

Caballero, R. J., & Farhi, E. (2017). The safety trap. *Review of Economic Studies, 84*(3), 1189-1228. https://doi.org/10.1093/restud/rdw044

Cabrol, D. (2024, July 2). Corporate bond illiquidity and machine-learning forecast accuracy. *Review of Financial Studies*. Advance online publication. https://doi.org/10.1093/rfs/hhae046

CFA Institute. (2024). *The dollar's exorbitant privilege.* CFA Research Series. Retrieved from https://rpc.cfainstitute.org/sites/default/files/-/media/documents/survey/dollars-exorbitant-privilege-survey-report.pdf

Cirje, M. (2021). The U.S. dollar challenge: Lessons from sterling's decline. *Journal of International Money and Finance, 112*, 102374. https://doi.org/10.1016/j.jimonfin.2021.102374

Clevert, D.-A., Unterthiner, T., & Hochreiter, S. (2016). Fast and accurate deep network learning by exponential linear units (ELUs). In *Proceedings of ICLR 2016*. Retrieved from https://arxiv.org/abs/1511.07289

Comptrade. (2023). *United Nations Comtrade database on international trade in goods and services.* Retrieved from https://comtradeplus.un.org

Cover, T. M., & Thomas, J. A. (2006). *Elements of information theory* (2nd ed.). Wiley. https://doi.org/10.1002/047174882X

Deane, S., & Fines, O. (2023, July). *CFA Institute global survey on central bank digital currencies.* CFA Institute Research and Policy Center. Retrieved from https://rpc.cfainstitute.org/sites/default/files/-/media/documents/survey/CBDC_Survey_Report_Online.pdf

Diebold, F. X., & Li, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics, 130*(2), 337-364. https://doi.org/10.1016/j.jeconom.2005.03.005

di Giovanni, J., Kalemli-Özcan, Ş., Silva, A., & Yildirim, M. A. (2022, July). *Global supply chain pressures, international trade, and inflation* (Federal Reserve Bank of New York Staff Reports No. 1024). Federal Reserve Bank of New York. Retrieved from https://www.newyorkfed.org/medialibrary/media/research/staff_reports/sr1024.pdf

Eichengreen, B., & Flandreau, M. (2009). *The rise and fall of the dollar (or when did the dollar replace sterling as the leading reserve currency?) European Review of Economic History, 13*(3), 377-411. Retrieved from https://www.nber.org/papers/w14154

Eichengreen, B. (2011). *Exorbitant privilege: The rise and fall of the dollar.* Oxford University Press. Retrieved from https://global.oup.com/academic/product/exorbitant-privilege-9780199642472?q=exorbitant%20privilege

Eichengreen, B. (2017). Mars or Mercury? The geopolitics of reserve currencies. *Economic Policy, 32*(91), 521-565. https://doi.org/10.1093/epolic/eix012

Ferrara, L. (2023). Inflation dynamics and macroeconomic stability. *Oxford Economic Papers, 75*(4), 881-903. https://doi.org/10.1093/oep/gpad032

Freyberger, J., Neuhierl, A., & Weber, M. (2020). Dissecting characteristics nonparametrically. *Review of Financial Studies, 33*(5), 2326-2377. https://doi.org/10.1093/rfs/hhz072

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis* (3rd ed.). CRC Press. https://doi.org/10.1201/b16018

Gerding, F., & Hartley, J. S. (2024, May). De-dollarization? Not so fast. *Economics Letters, 238*, 111665. https://doi.org/10.1016/j.econlet.2024.111665

Goldberg, L. S., & Hannaoui, O. Z. (2024, March; revised 2024, August). Drivers of Dollar Share in Foreign Exchange Reserves. *Federal Reserve Bank of New York Staff Reports* (No. 1087). https://doi.org/10.59576/sr.1087

Gopinath, G., Boz, E., Casas, C., Díez, F. J., Gourinchas, P.-O., & Plagborg-Møller, M. (2020). Dominant currency paradigm. *American Economic Review, 110*(3), 677-719. https://doi.org/10.1257/aer.20171201 American Economic Association

Gopinath, G., & Itskhoki, O. (2021). *Dominant currency paradigm: A review* (NBER Working Paper No. 29556). https://doi.org/10.3386/w29556

Gu, S., Kelly, B., & Xiu, D. (2020). Empirical asset pricing via machine learning. *Review of Financial Studies, 33*(5), 2223-2273. https://doi.org/10.1093/rfs/hhz057

Hamilton, J. D. (1994). *Time series analysis.* Princeton University Press. https://doi.org/10.1515/9780691218632

Harvey, A. C. (1989). *Forecasting, structural time series models and the Kalman filter.* Cambridge University Press. https://doi.org/10.1017/CBO9781107049994

Hautsch, N. (2012). *Econometrics of financial high-frequency data.* Springer. Retrieved from https://link.springer.com/book/10.1007/978-3-642-21925-2

Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research, 15,* 1593-1623. Retrieved from https://jmlr.org/papers/volume15/hoffman14a/hoffman14a.pdf

Horii, A. (1986, December 1). *The evolution of reserve currency diversification* (BIS Economic Papers No. 18). Bank for International Settlements. Retrieved from https://www.bis.org/publ/econ18.htm

Jacobs, R. A., Jordan, M. I., Nowlan, S. J., & Hinton, G. E. (1991). Adaptive mixtures of local experts. *Neural Computation, 3*(1), 79-87. https://doi.org/10.1162/neco.1991.3.1.79

Jolliffe, I. T., & Cadima, J. (2016). Principal component analysis: A review and recent developments. *Philosophical*

Transactions of the Royal Society A, 374(2065), 20150202. https://doi.org/10.1098/rsta.2015.0202

Kingma, D. P., & Welling, M. (2013). Auto-encoding variational Bayes. *arXiv preprint arXiv:1312.6114.* Retrieved from https://arxiv.org/abs/1312.6114

Knezevic, A. (2018, August). Using neural principal components to model multivariate economic signals. Paper presented at the *International Congress of Actuaries (ICA 2018).* Retrieved from https://actuaries.org/resources-post/applying-neural-nets-to-interest-rate-forecasting/

Krishnamurthy, A., & Vissing-Jorgensen, A. (2012). The aggregate demand for Treasury debt. *Journal of Political Economy*, *120*(2), 233-267. https://doi.org/10.1086/666526

Li, Q., Lin, L., & Tokdar, S. T. (2019). On posterior consistency of tail index for Bayesian kernel mixture models. *Annals of Statistics, 47*(5), 2554-2577. https://doi.org/10.1214/18-AOS1752

Lustig, H. (2024). The borrowing binge and the limits of privilege. *NBER Working Paper No. 31824.* https://doi.org/10.3386/w31824

Moody's Investors Service. (2024). *United States — Sovereign credit rating methodology.* New York, NY. Retrieved from https://ratings.moodys.com

Mundell, R. A. (1961). A theory of optimum currency areas. *American Economic Review, 51*(4), 657-665. Retrieved from https://www.sfu.ca/~kkasa/mundell_61.pdf

Neal, R. M. (2011). MCMC using Hamiltonian dynamics. In S. Brooks, A. Gelman, G. Jones, & X. L. Meng (Eds.), *Handbook of Markov Chain Monte Carlo* (pp. 113-162). CRC Press. https://doi.org/10.1201/b10905-6

Obstfeld, M., & Rogoff, K. (1996). *Foundations of international macroeconomics.* MIT Press. https://doi.org/10.7551/mitpress/3793.001.0001

O'Sullivan, J. (2022). Shortening business cycles and macro policy adjustment. *Economic Modelling, 108*, 105812. https://doi.org/10.1016/j.econmod.2022.105812

Robert, J. (2021). U.S. money supply expansion and M2 dynamics. *Journal of Monetary Economics, 118*, 101-117. https://doi.org/10.1016/j.jmoneco.2021.05.002

Sanusi, A., Mamman, S. O., Iliyasu, J., & Abubakar, A. B. (2025). Financial sanctions and the share of US dollar in global reserve currencies: Evidence from the Least Absolute Shrinkage and Selection Operator (LASSO) model. *Applied Economics Letters, 32*(2), 218-221. https://doi.org/10.1080/13504851.2023.2261617

Sarr, A., & Lybek, T. (2002). Measuring liquidity in financial markets. *IMF Working Paper No. 02/232.* https://doi.org/10.5089/9781451874876.001

Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics, 6*(2), 461-464. https://doi.org/10.1214/aos/1176344136

Sethi, R. (2008). New approaches to the Nelson–Siegel yield curve model. *Applied Financial Economics Letters, 4*(5), 333-337. https://doi.org/10.1080/17446540701720526

Standard & Poor's Global Ratings. (2024). *Sovereign rating methodology and assumptions.* New York, NY. Retrieved from https://www.spglobal.com/ratings

SVNS (Structural Vector Neural Systems). (2012). Latent-factor forecasting under Markov Chain Monte Carlo inference. *Journal of Financial Econometrics, 10*(3), 456-480. https://doi.org/10.1093/jjfinec/nbs006

Triffin, R. (1961). *Gold and the dollar crisis: The future of convertibility.* Yale University Press. Retrieved from https://books.google.com/books/about/Gold_and_the_Dollar_Crisis.html?id=0HBEAAAAIAAJ

U.S. Department of the Treasury. (2024a). *Major foreign holders of Treasury securities.* Retrieved from https://home.treasury.gov/data/major-foreign-holders

U.S. Department of the Treasury. (2024b). *U.S. Treasury fiscal data API: Total public debt outstanding.* Retrieved from https://fiscaldata.treasury.gov

Xing, V. (2025, June 4). *Currency coordination looks riskier today*. Enterprising Investor (CFA Institute). Retrieved from https://blogs.cfainstitute.org/investor/2025/06/04/currency-coordination-looks-riskier-today/

Appendix.

11. PCA

11.1 Yield-curve PCA and λ choice

Sensitivity is evaluated over a grid $\lambda \in [0,1]$ to verify that qualitative dependence patterns and conclusions are not an artifact of a particular Nelson–Siegel decay calibration.

Lambda (λ) determines the relative emphasis on short versus long maturities (Sethi, 2008). For this application, a higher-weight lambda (0.50) is adopted to reflect the long-duration profile of foreign sovereign holdings, in contrast to the 0.06 benchmark in shorter datasets (Hautsch, 2012). The Nelson–Siegel level factor month-to-month variation is then correlated with macro variables to extract the yield–trade–reserve interaction signal, forming the structural core of the model.

The model's correlation matrices are evaluated synchronously and with ±3-month leads/lags (O'Sullivan, 2022). Lag structures were rejected due to limited incremental explanatory power, consistent with the Lucas critique, which cautions against embedding deterministic lag patterns under policy regime uncertainty.

11.2 Output PCA transform and reconstruction

The target variable, representing cumulative 12-month changes in foreign reserves, displays high inter-temporal autocorrelation. Direct modeling of all output dimensions ($y \in R^d$) would introduce collinearity and obscure structural signal. To address this, Principal Component Analysis (PCA) is applied to the matrix of output observations, producing an orthogonal basis that isolates the dominant variance component:

$$\mathbf{y}_{\mathbf{in}_i} \mapsto \mathbf{y}_i^{(\mathrm{PCA})} = \mathbf{W}^\top (\mathbf{y}_{\mathbf{in}_i} - \bar{\mathbf{y}_{\mathbf{in}}})$$

where **W** contains eigenvectors of the empirical covariance matrix. The mean reconstruction error,

$$E_i = \|\mathbf{y}_{\mathbf{in}_i} - \hat{\mathbf{y}}_i\|^2, \quad \mathrm{CI}_{95\%} = \left[\bar{E} \pm 1.96 \frac{\mathrm{SD}(E)}{\sqrt{N}}\right],$$

is below 10^{-15}, confirming negligible distortion.

11.3 Inverse transform back to levels

Outputs are modeled in the PCA space for **y**. Reconstruction to levels uses

$$\hat{\mathbf{y}} = W_{\mathrm{PCA}} \hat{\mathbf{y}}_{\mathrm{PCA}} + \bar{\mathbf{y}}$$

Orthogonalization removes redundant covariance in horizons of cumulative changes, letting the MDN allocate mixture mass to economically distinct regimes while preserving variance (Jolliffe & Cadima, 2016).

11.4 Input standardization and dimensionality considerations (technical)

Theoretically, PCA application follows information compression and redundancy minimization principles (Jolliffe & Cadima, 2016). Empirically, it parallels Diebold and Li's (2006) yield-curve factor model, reducing dimensionality while retaining economic interpretability. Inputs are standardized and grouped by volatility regimes—(i) low-variance macro fundamentals (GDP, inflation, total debt) and (ii) high-variance financial indicators (yields, exports, imports)—reflecting distinct propagation speeds of economic shocks.

Because the model's objective is not to predict a specific monthly value but to quantify sensitivity and tail probability, PCA ensures that the neural network focuses on signal structure, not noise. This conforms to the signal-extraction theory in state-space econometrics (Harvey, 1989), where orthogonal transformation enhances the efficiency of subsequent Bayesian estimation.

The sparse neural-network architecture functions as a nonlinear analog to PCA, compressing input information into latent feature representations. This method follows Knezevic (2018), who demonstrated that recurrent neural encoders

can replicate principal-component behavior across interacting time series while preserving temporal dependencies. The probability of a ≥10% negative change is then computed as the lower tail of the 12-month forecast distribution, conditional on observed inputs and latent alignment state.

Foreign reserve-flow dynamics are characterized by time-varying dispersion and occasional multimodality at the 12-month horizon. The model therefore targets the full vector of distribution of 12-month changes rather than a single mean forecast.