Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Stablecoin Liquidity as a Crypto-Native Regime Signal: State-Dependent Density Forecasts for BTC, ETH, and SOL

Veliota Drakopoulou

drakopov@erau.edu

Higher Colleges of Technology

Research Article

Keywords: stablecoins, cryptocurrency returns, hidden Markov models, regime switching, density forecasting, Bitcoin, Ethereum, Solana

Posted Date: June 18th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9715935/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Stablecoin Liquidity as a Crypto-Native Regime Signal: State-Dependent Density Forecasts for BTC, ETH, and SOL

Abstract

This study examines whether stablecoin liquidity, network congestion, and macro-financial stress predicts major cryptoasset returns in a regime-dependent manner. Using daily data for BTC, ETH, and SOL from April 15, 2020, to March 17, 2026, the analysis estimates homogeneous and non-homogeneous hidden Markov models to capture shifts across latent low-, medium-, and high-volatility states. The paper's central contribution is to reframe stablecoin activity not as a constant linear return predictor, but as a crypto-native signal of changing market regimes. The final model selection supports BTC NHHM (3), ETH HHM (3), and SOL HHM (3) as the preferred specifications. These models rank first by average log score and continuous ranked probability score for their respective assets, indicating superior density-forecast performance, although mean squared forecast error gains are not uniform. The coefficient evidence is strongest for BTC and ETH. In medium-volatility regimes, stablecoin volume growth is positively associated with returns, whereas stablecoin transaction-count growth is negatively associated with returns, suggesting that different dimensions of stablecoin activity capture distinct liquidity and stress channels. BTC also exhibits stronger macro-stress sensitivity in the high-volatility regime. SOL shows clear regime separation and density-forecast gains, but its mean-equation coefficients are less precise. Overall, the findings show that stablecoin liquidity helps characterize latent crypto-market risk states and improves distributional forecasting, creating value for researchers studying crypto-market transmission and for investors managing dynamic volatility exposure.

Keywords: stablecoins; cryptocurrency returns; hidden Markov models; regime switching; density forecasting; Bitcoin; Ethereum; Solana

1. Introduction

Cryptoasset markets are difficult to model because their return dynamics are unstable across market conditions. Periods of calm trading can quickly give way to sharp volatility, liquidity stress, and cross-market repricing. This instability makes a single linear return model too restrictive. The central question of this paper is therefore not only whether stablecoin liquidity, network congestion, and macro-financial stress predict BTC, ETH, and SOL returns, but whether their predictive content changes across latent market regimes.

This paper argues that stablecoin activity is best understood as a crypto-native regime signal. Stablecoins are widely used as dollar-linked settlement and trading vehicles within crypto markets, making their supply, volume, transaction activity, and active-address growth natural indicators of internal liquidity conditions (Lyons & Viswanath-Natraj, 2023). Prior research shows that crypto markets can be segmented and liquidity constrained, with arbitrage frictions and stablecoin flows playing important roles in price formation (Griffin & Shams, 2020; Makarov & Schoar, 2020). Building on this literature, the present study asks whether stablecoin variables help identify the volatility states in which crypto return distributions shift.

The empirical framework estimates homogeneous and non-homogeneous hidden Markov models for BTC, ETH, and SOL. This approach allows dynamic returns to vary across unobserved low-, medium-, and high-volatility regimes, rather than imposing one average relationship over the full sample. The non-homogeneous specification also permits observed predictors to affect regime-transition probabilities, making it useful for testing whether liquidity and stress variables help explain movements between latent market states (Hamilton, 1989; Kim & Nelson, 1999; Rabiner, 1989). The final replication package supports regime-based preferred specifications for all three assets and shows that the main empirical evidence is strongest for density forecasting rather than uniform MSFE dominance.

The paper makes three contributions. First, it shifts the focus from constant return prediction to regime-dependent distributional forecasting. This is important because investors and researchers often care not only about expected returns, but also about whether the predictive distribution is widening, whether downside risk is increasing, and whether the market is moving toward a higher-volatility state. Second, it treats stablecoin variables as economically meaningful liquidity indicators rather than passive controls. Third, it distinguishes density-forecast gains from point-forecast gains by emphasizing proper scoring rules, including log score and CRPS, rather than relying only on MSFE (Gneiting & Raftery, 2007; Gneiting & Ranjan, 2011).

The core interpretation is that stablecoin liquidity helps reveal the market state in which return predictability becomes economically meaningful. Stablecoin volume growth can be interpreted as a proxy for deployable liquidity or capital availability, whereas stablecoin transaction-count growth may reflect stress, fragmentation, or repositioning. This distinction is difficult to observe in a single-regime model. It becomes visible when stablecoin variables are allowed to operate differently across latent volatility regimes. The remainder of the paper develops this argument by reviewing the related literature, describing the data and variable construction, presenting the hidden-state forecasting framework, and evaluating the empirical results.

2. Literature Review

The study connects three strands of literature: cryptoasset market structure, stablecoin liquidity, and regime-switching forecast evaluation. The first strand emphasizes that crypto markets are fragmented, volatile, and subject to liquidity frictions. Makarov and Schoar (2020) show that cryptocurrency markets exhibit large and recurring arbitrage opportunities across exchanges, highlighting the importance of capital constraints and market segmentation. Griffin and Shams (2020) examine whether Tether flows were associated with Bitcoin price movements during the 2017 boom, making stablecoin activity central to the debate over crypto price formation.

Together, these studies motivate the view that crypto-native liquidity variables may contain information not captured by standard financial-market controls.

The second strand examines stablecoins as liquidity infrastructure. Stablecoins are designed to reduce crypto trading frictions by maintaining a value close to a reference currency, typically the U.S. dollar. Lyons and Viswanath-Natraj (2023) show that stablecoins are used as vehicles for moving between national currencies and non-stable cryptocurrencies, with some stablecoins exhibiting trading volumes that are large relative to outstanding supply. This role makes stablecoins more than passive payment instruments. In crypto markets, they function as settlement assets, liquidity reservoirs, trading collateral, and bridges between fiat and digital-asset risk exposure. For this reason, stablecoin supply, volume, transaction counts, and active-address activity may proxy changes in internal crypto liquidity conditions.

The third strand concerns regime-switching models. Hamilton (1989) introduced a tractable econometric framework in which time-series dynamics change across unobserved Markov regimes. Hidden Markov models generalize this idea by allowing the observed process to depend on an unobserved state variable (Rabiner, 1989). In financial applications, regime-switching models are useful because volatility, expected returns, and cross-market sensitivities often differ between tranquil and stressed periods (Kim & Nelson, 1999). Non-homogeneous or time-varying transition-probability models extend this logic by allowing observed variables to influence the probability of moving between regimes (Filardo, 1994). This extension is especially relevant for the present study because stablecoin and macro variables may not only affect returns within a regime; they may also help identify when the market is transitioning from one volatility state to another.

The forecast-evaluation literature provides the final methodological anchor. If the model's purpose is to forecast the full conditional distribution, evaluation should not rely only on squared forecast errors. Proper scoring rules such as log score and CRPS evaluate predictive densities and reward calibrated probabilistic forecasts (Gneiting & Raftery, 2007). Pairwise forecast-loss comparisons, following Diebold and Mariano (1995), provide an additional way to assess whether competing models differ in predictive accuracy. This paper therefore treats log score and CRPS as primary criteria and MSFE as secondary, matching the evaluation method to the distributional objective of the regime model.

This literature leaves an important gap. Existing crypto studies show that crypto markets are segmented, liquidity sensitive, and connected to stablecoin activity, while regime-switching methods show how financial relationships can vary across latent states. However, less is known about whether stablecoin variables act as regime-state indicators that improve the predictive distribution of major cryptoasset returns. This paper addresses that gap by combining stablecoin liquidity measures, congestion variables, and macro-financial controls in a hidden-state forecasting framework for BTC, ETH, and SOL. The contribution is strong where empirical design shows that stablecoin information matters selectively across regimes rather than uniformly across the full sample.

3. Data and Variable Construction

This study uses daily data for BTC, ETH, and SOL to examine whether stablecoin liquidity, network congestion, and macro-financial stress predict cryptoasset returns in a regime-dependent manner. The final model-ready sample spans April 15, 2020, to March 17, 2026. As shown in Table 1, BTC and ETH each contain 2,163 observations, while SOL contains 1,695 observations because of its shorter usable history. The empirical design uses an 80/20 train-holdout split: BTC and ETH each have 1,730 training observations and 433 holdout observations, whereas SOL has 1,356 training observations and 339 holdout observations. Table 1 also reports on the frozen preferred model set: BTC NHHM (3), ETH HHM (3), and SOL HHM (3).

Table 1

Sample Construction, Forecast Split, and Frozen Preferred Models

Asset	Full sample start	Full sample end	Full N	Train N	Holdout N	Matched-sample N	Frozen preferred model
BTC	2020-04-15	2026-03-17	2,163	1,730	433	1,695	NHHM (3)
ETH	2020-04-15	2026-03-17	2,163	1,730	433	1,695	HHM (3)
SOL	2020-04-15	2026-03-17	1,695	1,356	339	1,695	HHM (3)

Note. Full-sample rows come from data/processed/processing_report.json. Train and holdout rows come from data/processed_train_80/processing_report.json. Matched-sample N comes from data/processed_matched_sol/processing_report.json. Preferred models come from outputs/final_model_freeze/frozen_preferred_models.csv.

The dependent variable is the daily percentage log return of each target asset, defined as

$$r_t = 100[\ln(P_t) - \ln(P_{t-1})],$$

where $P_t$denotes the daily reference price. The models also include the one-day lag of the target return to capture short-run return dependence. All predictors are lagged before estimation so that the forecasting exercise uses only information available prior to the return realization being predicted.

As summarized in Table 2, the explanatory variables are organized into three blocks. The first block captures stablecoin liquidity using aggregate USDT and USDC market-capitalization growth, spot-volume growth, transaction-count growth, and active-address growth. The second block captures network congestion and crypto activity using aggregate BTC and ETH fee growth and transaction-count growth. The third block captures macro-financial conditions using S&P 500 returns, VIX growth, and changes in the 10-year Treasury yield. Crypto and stablecoin variables are based on Coin Metrics data, while macro-financial variables are based on FRED series for the S&P 500, VIX, and 10-year Treasury yield (Coin Metrics, n.d.; Federal Reserve Bank of St. Louis, n.d.).

Table 2

Variable Definitions, Transformations, and Data Sources

Variable	Definition	Construction	Source
ret	Daily percentage log return of the target asset	100 * [ln(P_t) - ln(P_{t-1})]	Coin Metrics reference rate
lag_ret	One-day lag of the target asset return	ret_{t-1}	Derived
SCAPG	Growth in aggregate USDT and USDC market capitalization	Lagged 100 * Deltaln (CapMrktCurUSD_USDT + CapMrktCurUSD_USDC)	Coin Metrics
SVOLG	Growth in aggregate USDT and USDC spot volume	Lagged 100 * Deltaln (volume_reported_spot_usd_1d_USDT + volume_reported_spot_usd_1d_USDC)	Coin Metrics
STXG	Growth in aggregate USDT and USDC transaction counts	Lagged 100 * Deltaln (1 + TxCnt_USDT + TxCnt_USDC)	Coin Metrics
SADDRG	Growth in aggregate USDT and USDC active addresses	Lagged 100 * Deltaln (1 + AdrActCnt_USDT + AdrActCnt_USDC)	Coin Metrics
FEEG	Growth in aggregate BTC and ETH fee totals	Lagged 100 * Deltaln (1 + FeeTotNtv_BTC + FeeTotNtv_ETH)	Coin Metrics
BTXG	Growth in aggregate BTC and ETH transaction counts	Lagged 100 * Deltaln (1 + TxCnt_BTC + TxCnt_ETH)	Coin Metrics
SPXRET	S&P 500 daily percentage log return	Lagged 100 * [ln (SP500_t) - ln (SP500_{t-1})]	FRED / SP500
VIXG	Growth in the VIX level	Lagged 100 * [ln (VIX_t) - ln (VIX_t-1})]	FRED / VIXCLS
TY10	Daily change in the 10-year Treasury yield	Lagged (DGS10_t - DGS10_{t-1})	FRED / DGS10

Note. Variables are copied from data/processed/variable_dictionary.csv. Predictors are lagged and standardized before estimation in the model outputs.

Several construction choices are used to make the variables suitable for regime-based forecasting. Growth variables are expressed as percentage log changes. Count and fee variables are transformed using $\ln(1 + x)$before differencing, which accommodates zeros and reduces the influence of extreme observations. Predictors are

standardized before estimation so that coefficients can be interpreted as the return effect associated with a one-standard-deviation movement in the corresponding predictor.

The descriptive statistics in Table 3 show substantial heterogeneity across the three assets. BTC has a mean daily return of 0.110% and a standard deviation of 3.030. ETH has a mean daily return of 0.125% and a standard deviation of 4.089. SOL has a higher mean daily return of 0.355% and a much larger standard deviation of 6.728. SOL also has the most extreme downside observation, with a minimum daily return of -54.958%, compared with -18.422% for BTC and -28.391% for ETH. These differences justify estimating asset-specific models rather than imposing a common return process across BTC, ETH, and SOL.

Table 3

Descriptive Statistics for Model-Ready Returns and Predictors

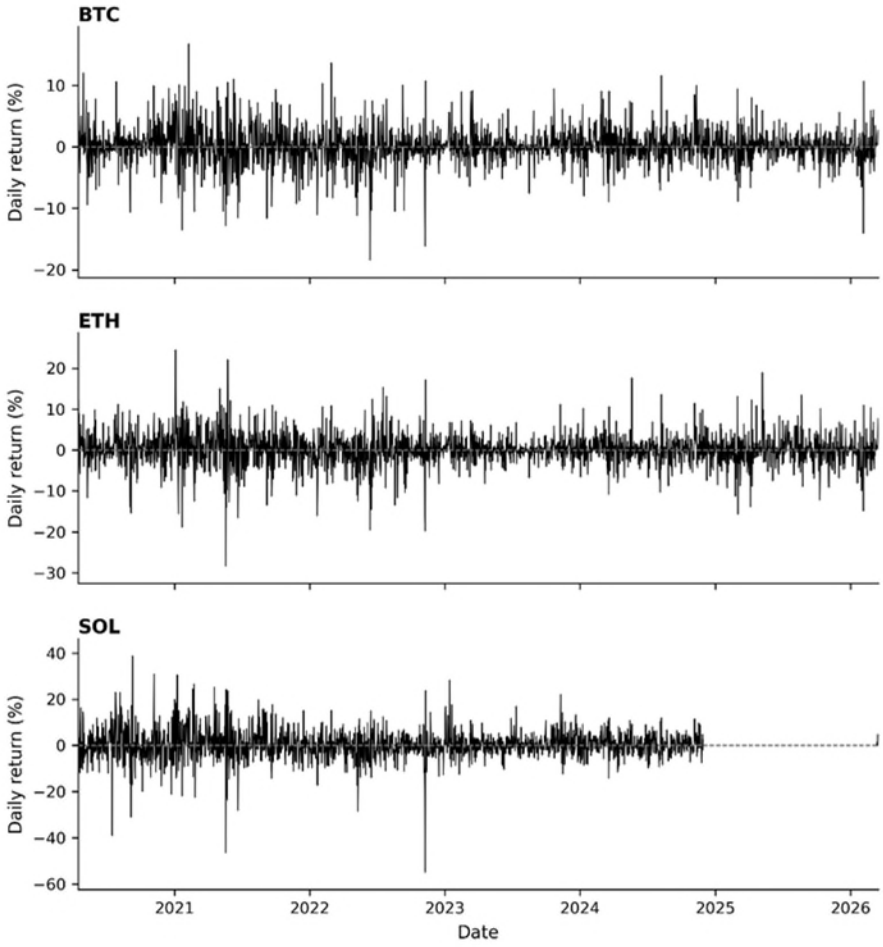
Asset	Variable	N	M	SD	Min	Median	Max
BTC	ret	2,163	0.110	3.030	-18.422	0.080	16.765
BTC	lag_ret	2,163	0.109	3.029	-18.422	0.080	16.765
BTC	SCAPG	2,163	0.163	0.541	-8.657	0.035	3.904
BTC	SVOLG	2,163	0.092	33.550	-107.062	-0.540	127.249
BTC	STXG	2,163	0.197	11.248	-70.917	-0.407	66.437
BTC	SADDRG	2,163	0.179	11.848	-95.511	-0.595	75.030
BTC	FEEG	2,163	-0.045	36.752	-246.962	-1.408	280.566
BTC	BTXG	2,163	0.042	6.854	-50.052	-0.149	61.162
BTC	SPXRET	2,163	0.041	0.904	-6.161	0.000	9.089
BTC	VIXG	2,163	-0.026	6.349	-44.245	0.000	55.411
BTC	TY10	2,163	0.002	0.049	-0.300	0.000	0.280
ETH	ret	2,163	0.125	4.089	-28.391	0.109	24.483
ETH	lag_ret	2,163	0.122	4.086	-28.391	0.109	24.483
ETH	SCAPG	2,163	0.163	0.541	-8.657	0.035	3.904
ETH	SVOLG	2,163	0.092	33.550	-107.062	-0.540	127.249
ETH	STXG	2,163	0.197	11.248	-70.917	-0.407	66.437
ETH	SADDRG	2,163	0.179	11.848	-95.511	-0.595	75.030
ETH	FEEG	2,163	-0.045	36.752	-246.962	-1.408	280.566
ETH	BTXG	2,163	0.042	6.854	-50.052	-0.149	61.162
ETH	SPXRET	2,163	0.041	0.904	-6.161	0.000	9.089
ETH	VIXG	2,163	-0.026	6.349	-44.245	0.000	55.411
ETH	TY10	2,163	0.002	0.049	-0.300	0.000	0.280
SOL	ret	1,695	0.355	6.728	-54.958	0.049	38.718
SOL	lag_ret	1,695	0.342	6.739	-54.958	0.037	38.718
SOL	SCAPG	1,695	0.183	0.590	-8.657	0.038	3.904
SOL	SVOLG	1,695	0.160	31.449	-101.062	-0.626	127.249
SOL	STXG	1,695	0.224	11.106	-70.917	-0.599	66.437
SOL	SADDRG	1,695	0.216	11.862	-95.511	-0.607	75.030
SOL	FEEG	1,695	0.072	32.177	-246.962	-1.204	280.566
SOL	BTXG	1,695	0.033	6.616	-50.052	-0.208	61.162
SOL	SPXRET	1,695	0.045	0.902	-6.075	0.000	5.395
SOL	VIXG	1,695	-0.066	6.006	-33.068	0.000	50.017
SOL	TY10	1,695	0.002	0.051	-0.300	0.000	0.280

Note. Statistics were computed directly from the model-ready files in data/processed. ret is the target daily percentage log return; predictors are in their pre-standardized percentage-change or change units.

The predictor variables in Table 3 also exhibit substantial dispersion, especially stablecoin volume growth, stablecoin transaction-count growth, active-address growth, and fee growth. This variation supports the use of a regime-dependent model because stablecoin and congestion variables are unlikely to have a single constant effect across tranquil, transitional, and stressed market environments.

Figure 1

Daily Returns and Volatility Clustering for BTC, ETH, and SOL.



Note. The figure plots daily percentage log returns for BTC, ETH, and SOL from the model-ready data. Returns are expressed in percentage points. The panels use asset-specific vertical scales to make volatility clustering visible within each asset.

As shown in Figure 1, the return series should display periods of volatility clustering, with relatively calm trading intervals interrupted by sharp high-volatility episodes. This visual pattern motivates the use of hidden-state models because a single-regime specification would average across very different volatility environments.

4. Econometric Framework and Forecast Evaluation

This section defines the econometric framework and forecast-evaluation design. The sample construction, variable definitions, and descriptive statistics were presented in Section 3 using Tables 1-3; those tables are not repeated here. The modelling strategy treats BTC, ETH, and SOL returns as regime dependent rather than governed by a single stable return process. Hidden Markov and regime-switching models are appropriate for this setting because they allow an unobserved state variable to govern shifts in conditional means, variances, and transition dynamics (Hamilton, 1989; Kim & Nelson, 1999; Rabiner, 1989). The empirical framework links lagged stablecoin, congestion, and macro-financial predictors to state-dependent cryptoasset returns. Let $r_{i,t}$denote the daily percentage log return for asset $i \in$ {BTC,ETH,SOL}. The observation equation is

$$r_{i,t} = \mu_{i,s_t} + \mathrm{x}'_{t-1}\beta_{i,s_t} + \varepsilon_{i,t},$$
$$\varepsilon_{i,t} \mid s_t = k \sim N(0, \sigma^2_{i,k}),$$

where s_t is the latent regime, x_{t-1} is the lagged predictor vector, β_{i,s_t} is the regime-specific coefficient vector, and $\sigma^2_{i,k}$ is the state-specific innovation variance. This specification allows the same predictor to have different effects in low-, medium-, and high-volatility regimes.

Two hidden-state structures are estimated. In the homogeneous hidden Markov model, transition probabilities are fixed:

$$P(s_t = j \mid s_{t-1} = k) = p_{kj}.$$

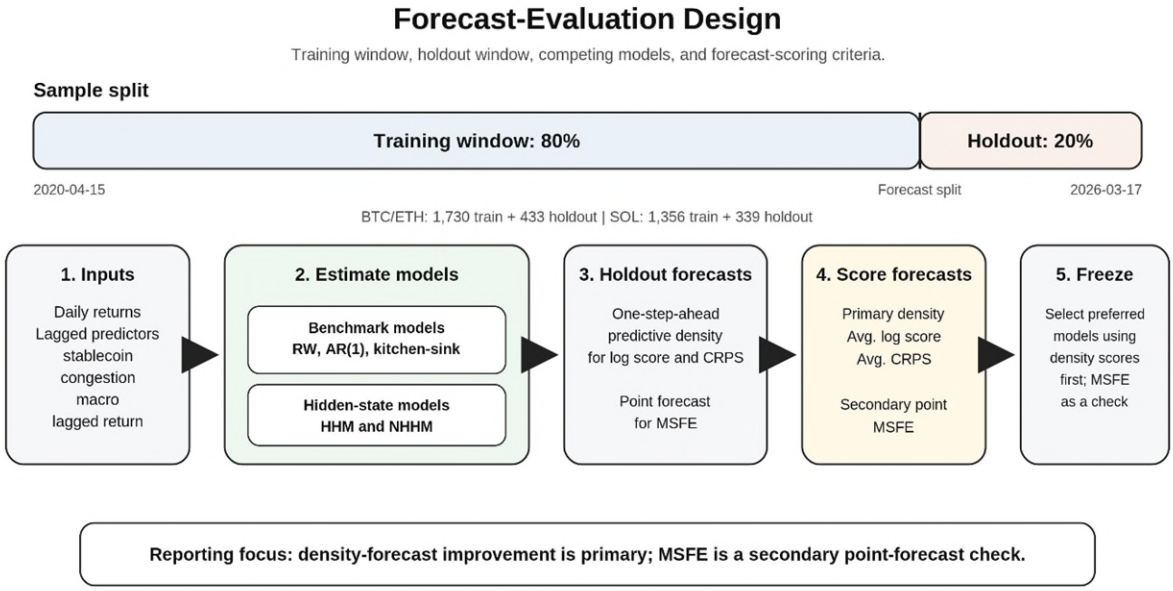
In the non-homogeneous hidden Markov model, transition probabilities vary with lagged predictors:

$$P(s_t = j \mid s_{t-1} = k, z_{t-1}) = \frac{\exp(\alpha_{kj} + z'_{t-1}\gamma_{kj})}{\sum_{\ell=1}^{K} \exp(\alpha_{k\ell} + z'_{t-1}\gamma_{k\ell})}.$$

The NHHM is more flexible because predictors can influence not only returns within a regime but also the probability of moving between regimes. This distinction is central to the paper's contribution. The HHM asks whether predictor effects differ across states; the NHHM additionally asks whether observable market conditions help explain the evolution of the latent state itself. Time-varying transition-probability models are useful when regime shifts are expected to depend on economic or financial conditions (Filardo, 1994; Kim & Nelson, 1999).

The models are estimated in a Bayesian framework. Posterior summaries are reported using posterior means, standard deviations, and credible intervals. Regimes are ordered by posterior variance and interpreted as low-, medium-, and high-volatility states. This ordering avoids arbitrary state labels and allows the volatility structure to be compared across assets.

Figure 2
Forecast-Evaluation Design.



Note. The figure shows the 80% training window, 20% holdout window, competing benchmark models, hidden-state models, and forecast-scoring criteria. RW = random walk; HHM = homogeneous hidden Markov model; NHHM = non-homogeneous hidden Markov model; CRPS = continuous ranked probability score; MSFE = mean squared forecast error.

Figure 2 summarizes the forecast-evaluation design. Forecast performance is evaluated out of sample using the holdout periods defined in Table 1. The hidden-state models are compared with random walk, AR (1), kitchen-sink regression, and alternative hidden-state specifications. The primary forecast metrics are average log score and average continuous ranked probability score (CRPS), while MSFE is treated as a secondary point-forecast metric. Log score and CRPS are appropriate because the regime models generate full predictive distributions rather than only conditional-mean forecasts (Gneiting & Raftery, 2007; Gneiting & Ranjan, 2011).

The forecast-evaluation design therefore matches the paper's substantive objective. The purpose is not only to determine whether stablecoin, congestion, and macro variables improve next day mean forecasts. The more important question is whether these variables help identify shifts in the conditional return distribution. Pairwise forecast comparisons use Diebold-Mariano style loss differentials to assess whether preferred models reduce forecast loss relative to benchmarks (Diebold & Mariano, 1995).

5. Empirical Results

This section reports the empirical findings from the regime-switching forecast framework. The sample, variables, and descriptive statistics were presented in Section 3 using Tables 1-3. Section 5 therefore begins with Table 4, which reports the preferred frozen model set and forecast metrics. The main result is that BTC NHHM (3), ETH HHM (3), and SOL HHM (3) provide the strongest density forecasts for their respective assets. These models rank first by average log score and CRPS, while MSFE rankings are not uniformly best. The evidence should therefore be interpreted as support for state-dependent density predictability, not universal point-forecast superiority.

5.1 Preferred Models and Forecast Performance

As shown in Table 4, BTC is modeled as NHHM (3), while ETH and SOL are modeled as HHM (3). BTC NHHM (3) records an average log score of -2.216, average CRPS of 1.250, and MSFE of 5.637. ETH HHM (3) records an average log score of -2.739, average CRPS of 2.119, and MSFE of 15.998. SOL HHM (3) records an average log score of -2.931, average CRPS of 2.444, and MSFE of 18.459. These statistics establish the preferred model set used for the rest of the empirical analysis.

Table 4

Frozen Preferred Model Set and Final Forecast Metrics

Asset	Model	Avg. Log score	Avg. CRPS	MSFE	Rationale
BTC	NHHM (3)	-2.216	1.250	5.637	Frozen by holdout density performance among BTC hidden-state candidates (log score and CRPS primary, MSFE secondary).
ETH	HHM (3)	-2.739	2.119	15.998	Only one hidden-state candidate was contested for this asset in the final freeze pass; retained as preferred.
SOL	HHM (3)	-2.931	2.444	18.459	Only one hidden-state candidate was contested for this asset in the final freeze pass; retained as preferred.

Note. Avg. log score and avg. CRPS are the primary density-forecast criteria; MSFE is secondary. Higher Log score is better; lower CRPS and MSFE are better. Source: outputs/final_model_freeze/frozen_preferred_models.csv.

The cross-asset model choices are economically meaningful. BTC receives the more flexible NHHM specification, indicating that its regime dynamics benefit from allowing predictors to affect transition probabilities. ETH and SOL retain homogeneous transition matrices, suggesting that their forecast gains are captured primarily through state-dependent means and variances rather than covariate-driven state transitions.

Table 5

Final Holdout Forecast Contest Across Preferred Models and Benchmarks

Asset	Model	Avg. Log score	Avg. CRPS	MSFE	N	Log score rank	CRPS rank	MSFE rank	Density rank sum
BTC	NHHM (3)	-2.216	1.250	5.637	433	1	1	4	2
BTC	NHHM (3)-SmallCore	-2.226	1.252	5.612	433	2	2	1	4
BTC	HHM (3)	-2.239	1.257	5.672	433	3	3	6	6
BTC	Kitchen Sink	-2.352	1.316	5.618	433	4	4	2	8
BTC	RW	-2.353	1.318	5.637	433	5	5	3	10
BTC	AR (1)	-2.354	1.321	5.662	433	6	6	5	12
ETH	HHM (3)	-2.739	2.119	15.998	433	1	1	4	2
ETH	RW	-2.795	2.124	15.630	433	2	2	1	4
ETH	AR (1)	-2.797	2.130	15.689	433	3	3	2	6
ETH	Kitchen Sink	-2.798	2.133	15.725	433	4	4	3	8
SOL	HHM (3)	-2.931	2.444	18.459	339	1	1	2	2
SOL	Kitchen Sink	-3.069	2.619	18.472	339	2	3	3	5
SOL	AR (1)	-3.070	2.618	18.386	339	3	2	1	5
SOL	RW	-3.072	2.624	18.502	339	4	4	4	8

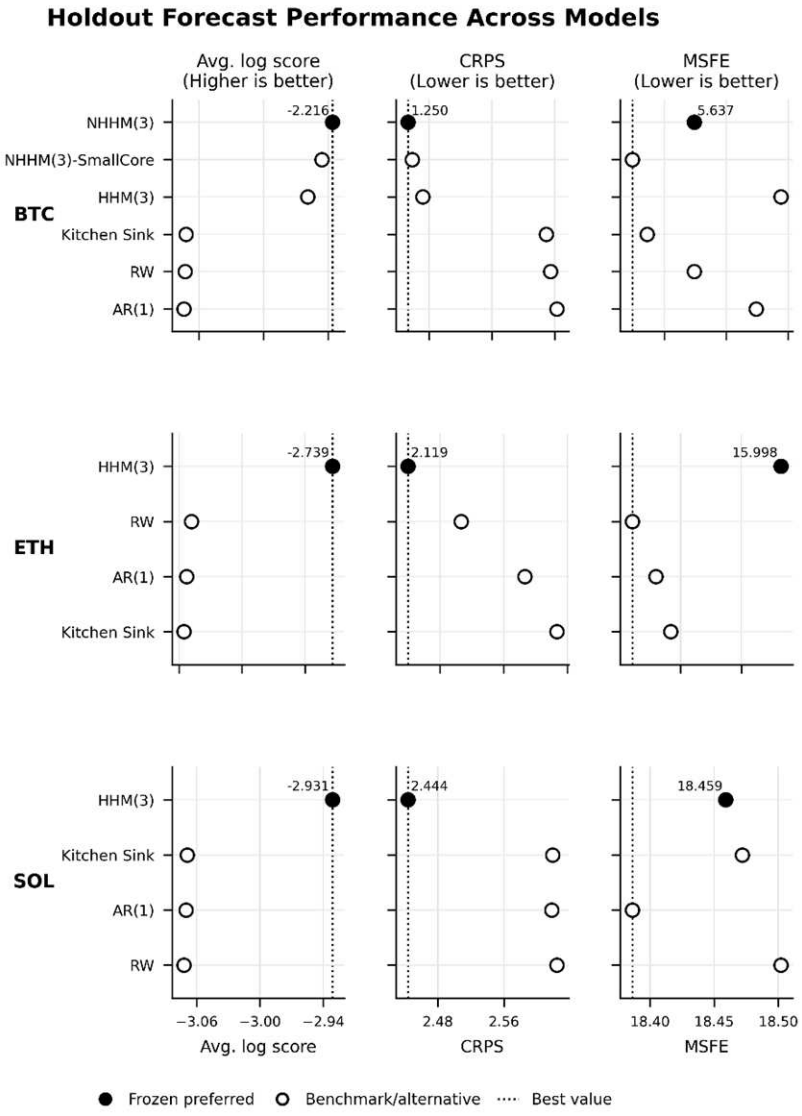
Note. This is the controlling forecast table for manuscript reporting because it includes the BTC NHHM (3)-SmallCore candidate. The frozen preferred models rank first by log score and CRPS for all three assets. Source: outputs/final_model_freeze/table_6_final_forecast_contest.csv.

Table 5 provides the controlling holdout forecast contest. For BTC, NHHM (3) ranks first by average log score and CRPS, even though the NHHM (3)-SmallCore candidate has the lowest MSFE. For ETH, HHM (3) ranks first by log score and CRPS, although the random walk has the lowest MSFE. For SOL, HHM (3) ranks first by log score and CRPS, although AR (1) has the lowest MSFE. Thus, Table 5 shows that the preferred models improve predictive distributions rather than uniformly improving conditional-mean forecasts.

Figure 3 visualizes the same forecast contrast shown in Table 5. The expected pattern is that the preferred hidden-state models dominate by log score and CRPS, while MSFE winners sometimes differ. This distinction supports the central claim that the regime models are most valuable for forecasting the full return distribution.

Figure 3

Holdout Forecast Performance Across Models.



Note. Values are from Table 5. Filled circles indicate the frozen preferred model for each asset. Open circles indicate benchmark or alternative models. The dotted vertical line marks the best value within each asset-metric panel. Higher average log score is better; lower CRPS and MSFE are better

The pairwise comparisons in Table 6 confirm the density-forecast interpretation. Negative differences in negative log score generally favor the preferred hidden-state models over benchmarks. The BTC small-core comparison is especially informative: NHHM (3)-SmallCore slightly improves squared-error loss but worsens negative log-score loss relative to the frozen BTC NHHM (3). Because the model-selection rule prioritizes log score and CRPS, the original BTC NHHM (3) remains the preferred specification.

Table 6

Pairwise Forecast-Loss Comparisons

Asset	Comparison	DM, squared error	Mean diff, squared error	DM, neg. log score	Mean diff, neg. log score	N
BTC	NHHM (3) vs. RW	0.001	0.000	-5.473	-0.138	
BTC	NHHM (3) vs. AR (1)	-0.331	-0.025	-5.576	-0.139	
BTC	NHHM (3) vs. Kitchen Sink	0.318	0.019	-5.368	-0.136	
ETH	HHM (3) vs. RW	1.980	0.368	-2.169	-0.056	
ETH	HHM (3) vs. AR (1)	2.341	0.309	-2.289	-0.058	
ETH	HHM (3) vs. Kitchen Sink	2.737	0.273	-2.337	-0.059	
SOL	HHM (3) vs. RW	-0.176	-0.043	-8.751	-0.141	
SOL	HHM (3) vs. AR (1)	0.281	0.073	-8.520	-0.140	
SOL	HHM (3) vs. Kitchen Sink	-0.069	-0.013	-8.444	-0.138	
BTC	NHHM (3)-SmallCore vs. NHHM (3)	-0.533	-0.025	1.055	0.010	433
BTC	NHHM (3)-SmallCore vs. HHM (3)	-0.643	-0.060	-1.502	-0.013	433
BTC	NHHM (3)-SmallCore vs. RW	-0.386	-0.025	-5.359	-0.128	433
BTC	NHHM (3)-SmallCore vs. AR (1)	-0.924	-0.049	-5.427	-0.129	433
BTC	NHHM (3)-SmallCore vs. Kitchen Sink	-0.209	-0.006	-5.233	-0.126	433

Note. DM = Diebold-Mariano style statistics. Negative differences in negative log score indicate that the first model has better density loss than the benchmark. The BTC small-core comparisons show that the smaller-core candidate slightly lowers squared-error loss but worsens negative log-score loss relative to the frozen BTC NHHM (3). Sources: outputs/forecast_eval/diebold_mariano_results.csv and outputs/final_model_freeze/btc_smallcore_pairwise_dm.csv.

5.2 Regime Structure

Table 7 reports regime-specific variance and volatility estimates. The states are ordered by posterior variance and labeled low, medium, and high volatility. BTC's posterior mean daily standard deviations are 1.149, 2.595, and 4.864 percentage points. ETH's corresponding values are 1.571, 3.897, and 7.614. SOL has the highest conditional volatility, with state-specific standard deviations of 4.378, 5.821, and 12.999. These estimates show that the hidden states capture economically meaningful shifts in conditional risk rather than minor statistical fluctuations.

Table 7

Regime-Specific Variance and Volatility Estimates for Preferred Models

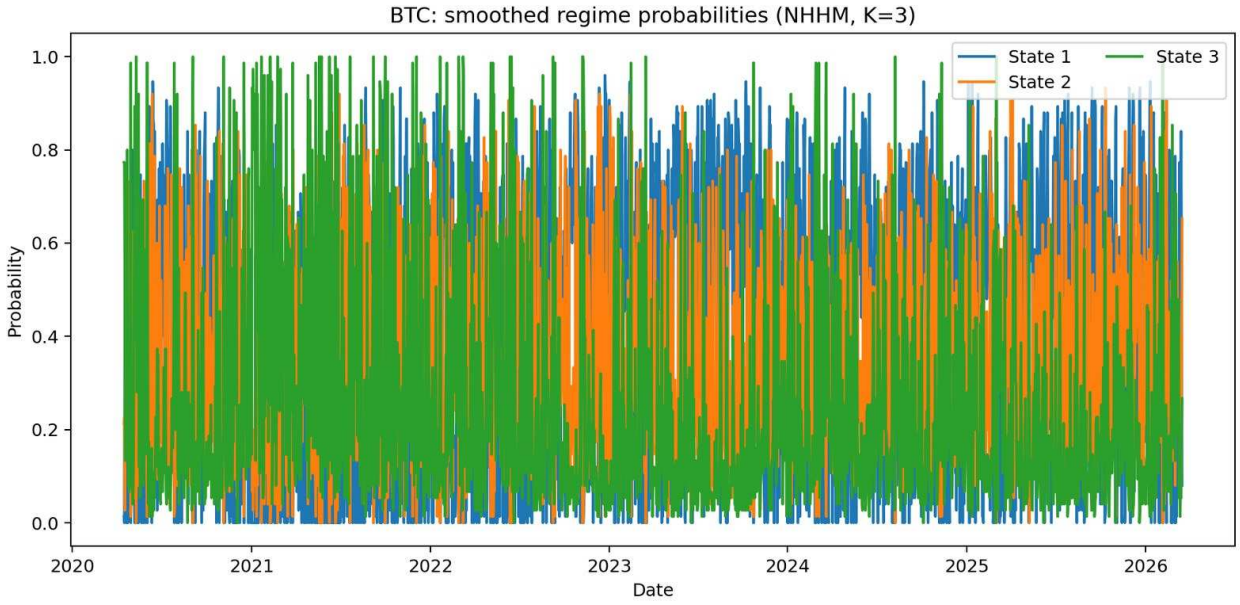
Asset	Model	State	Regime label	Variance M	Variance SD	Variance 90% CrI	sigma M
BTC	NHHM (3)	1	Low volatility	1.321	0.149	[1.086, 1.567]	1.149
BTC	NHHM (3)	2	Medium volatility	6.732	0.691	[5.596, 7.781]	2.595
BTC	NHHM (3)	3	High volatility	23.657	1.722	[20.888, 26.413]	4.864
ETH	HHM (3)	1	Low volatility	2.467	0.399	[1.685, 3.083]	1.571
ETH	HHM (3)	2	Medium volatility	15.184	1.824	[12.374, 18.537]	3.897
ETH	HHM (3)	3	High volatility	57.973	13.937	[40.915, 74.137]	7.614
SOL	HHM (3)	1	Low volatility	19.166	1.406	[16.887, 21.773]	4.378
SOL	HHM (3)	2	Medium volatility	33.881	11.574	[18.395, 53.960]	5.821
SOL	HHM (3)	3	High volatility	168.977	23.307	[134.206, 211.529]	12.999

Note. States are ordered by posterior variance and labeled low, medium, and high volatility. Sigma M is the posterior mean of the daily return standard deviation in percentage-point return units. Sources: preferred model sigma_summary.csv files.

The regime separation shown in Table 7 helps explain why the preferred models improve density scores even when they do not always minimize MSFE. A model can produce a similar conditional mean but a much better predictive distribution if it correctly adjusts uncertainty across volatility states.

Figure 4

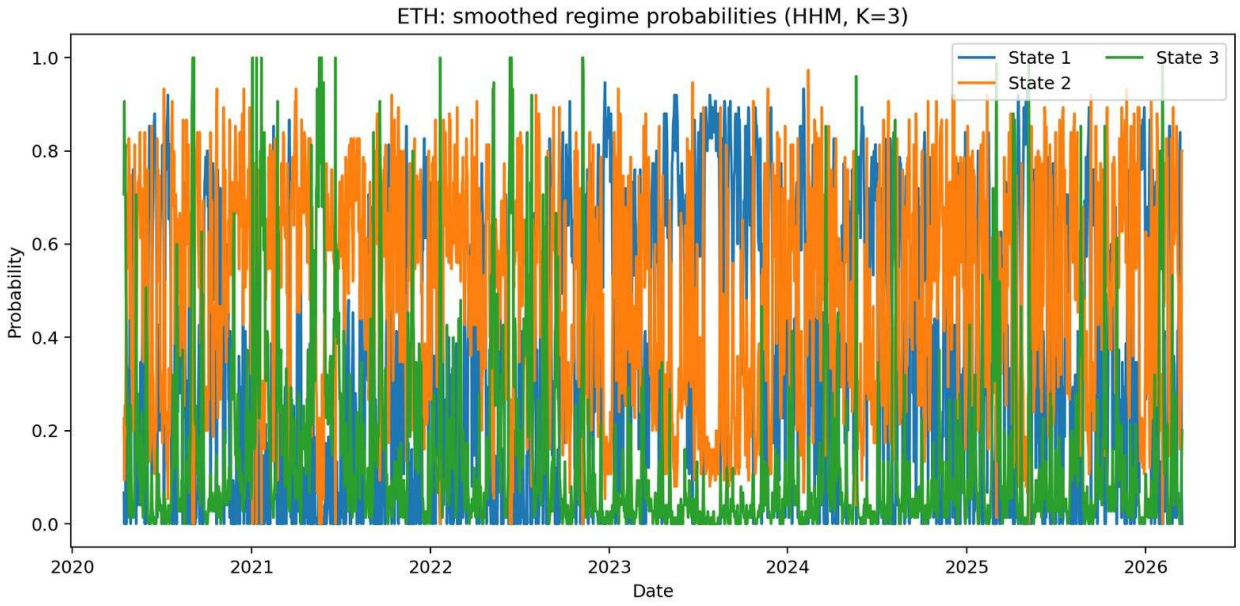
BTC Posterior Regime Probabilities from NHHM (3).



Note: This figure plots the smoothed posterior probabilities of BTC low-, medium-, and high-volatility states.

Figure 5

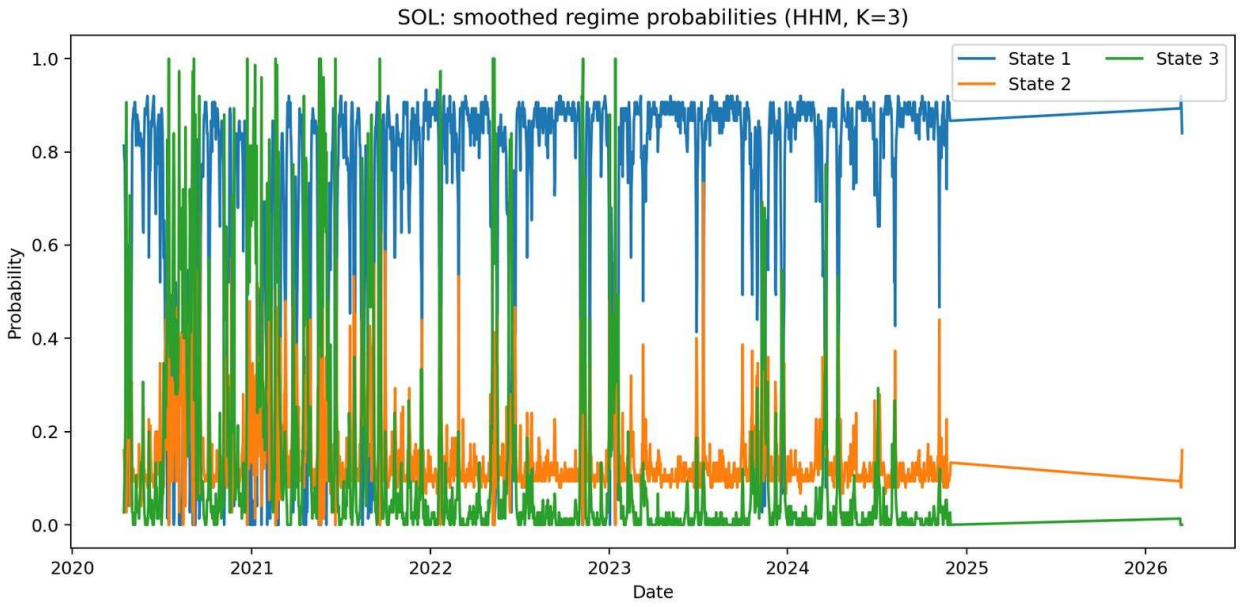
ETH Posterior Regime Probabilities from HHM (3).



Note: This figure plots the smoothed posterior probabilities of ETH low-, medium-, and high-volatility states.

Figure 6

SOL Posterior Regime Probabilities From HHM (3).



Note: This figure plot the smoothed posterior probabilities of SOL low-, medium-, and high-volatility states.

Figures 4-6 provide the dynamic interpretation of the volatility regimes reported in Table 7. Figure 4 should show when BTC is likely to occupy low-, medium-, or high-volatility states under the preferred NHHM (3). Figure 6 should provide the corresponding regime-probability path for ETH under HHM (3), and Figure 7 should do the same for SOL. For BTC and ETH, the medium-volatility state is especially important because it is where the most meaningful stablecoin-related coefficient evidence appears. For SOL, Figure 7 is more important for documenting volatility-state separation than for supporting coefficient-level claims.

5.3 State-Dependent Coefficient Evidence

Table 8 reports only mean-equation coefficients whose 90% credible intervals exclude zero. Because predictors are standardized, the coefficients can be interpreted as the return effect of a one-standard-deviation movement in the predictor within a given regime. The coefficient evidence is strongest for BTC and ETH, while SOL should be interpreted mainly through regime separation and density-forecast performance because its mean-equation intervals are imprecise.

Table 8

Mean-Equation Coefficients With 90% Credible Intervals Excluding Zero

Asset	Model	State	Regime label	Parameter	Direction	M	SD	90% CrI
BTC	NHHM (3)	1	Low volatility	lag_ret	Negative	-0.227	0.110	[-0.379, -0.033]
BTC	NHHM (3)	2	Medium volatility	SVOLG	Positive	0.555	0.231	[0.231, 0.968]
BTC	NHHM (3)	2	Medium volatility	STXG	Negative	-0.540	0.301	[-0.965, -0.056]
BTC	NHHM (3)	2	Medium volatility	SADDRG	Positive	0.436	0.241	[0.100, 0.831]
BTC	NHHM (3)	2	Medium volatility	FEEG	Negative	-0.523	0.238	[-0.898, -0.150]

Asset	Model	State	Regime label	Parameter	Direction	M	SD	90% CrI
BTC	NHHM (3)	2	Medium volatility	VIXG	Positive	0.466	0.287	[0.069, 0.996]
BTC	NHHM (3)	3	High volatility	SPXRET	Negative	-1.231	0.414	[-1.854, -0.564]
BTC	NHHM (3)	3	High volatility	VIXG	Negative	-1.155	0.441	[-1.813, -0.489]
ETH	HHM (3)	1	Low volatility	lag_ret	Negative	-0.621	0.188	[-0.936, -0.349]
ETH	HHM (3)	2	Medium volatility	SVOLG	Positive	0.714	0.322	[0.245, 1.364]
ETH	HHM (3)	2	Medium volatility	STXG	Negative	-0.817	0.469	[-1.645, -0.091]

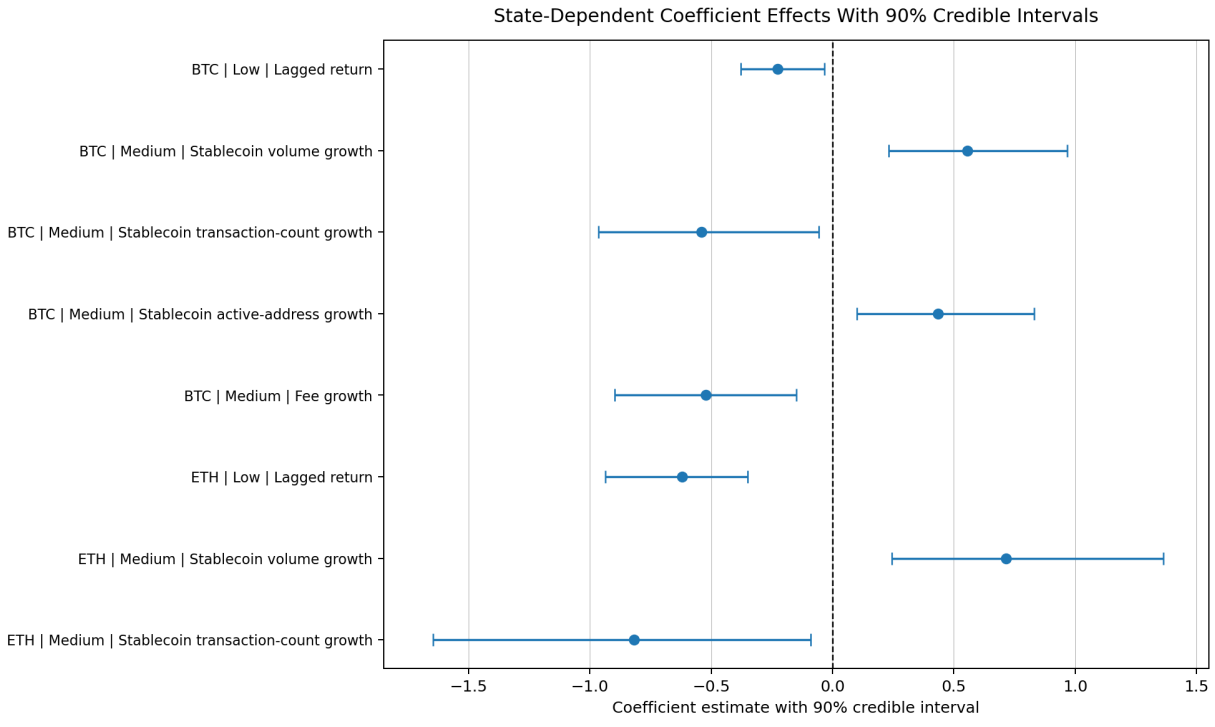
Note. The table excludes intercepts and lists only mean-equation coefficients whose 90% credible interval does not include zero. Coefficients are return effects per one-standard-deviation predictor movement because predictors are standardized.

For BTC, Table 8 shows short-run reversal in the low-volatility regime: the lagged-return coefficient is -0.227, with a 90% credible interval of [-0.379, -0.033]. In the medium-volatility regime, stablecoin and congestion variables become economically meaningful. Stablecoin volume growth is positive, 0.555 $[0.231, 0.968]$, while stablecoin transaction-count growth is negative, -0.540 $[-0.965, -0.056]$. Stablecoin active-address growth is also positive, 0.436 $[0.100, 0.831]$, whereas fee growth is negative, -0.523 $[-0.898, -0.150]$. BTC's high-volatility regime is more macro-sensitive, with negative SPXRET and VIXG coefficients.

For ETH, Table 8 shows a narrower but related pattern. In the low-volatility regime, lagged return is negative, -0.621 $[-0.936, -0.349]$. In the medium-volatility regime, stablecoin volume growth is positive, 0.714 $[0.245, 1.364]$, while stablecoin transaction-count growth is negative, -0.817 $[-1.645, -0.091]$. ETH therefore reinforces the interpretation that stablecoin liquidity is most informative in transitional volatility conditions.

Figure 7

State-Dependent Coefficient Effects With 90% Credible Intervals.



Note. This figure plot only the coefficients reported in Table 8, grouped by asset, regime, and predictor.

Figure 7 should visually summarize the coefficient evidence from Table 8. It should show that meaningful coefficient effects are concentrated in BTC and ETH, especially in the medium-volatility regime. This figure is important because it prevents overinterpretation of weaker posterior estimates and highlights the paper's central mechanism: stablecoin liquidity is most informative when markets are neither fully calm nor fully stressed.

5.4 Robustness and Diagnostics

Table 16 evaluates whether the preferred state counts are robust. For BTC and ETH, four-state models improve in-sample likelihood, but the three-state specifications are retained because they provide a more parsimonious and interpretable low-, medium-, and high-volatility structure. For SOL, the three-state HHM performs better than both the two-state and four-state alternatives, providing direct support for HHM (3).

Table 16

State-Count Robustness for Preferred Model Families

Asset	Family	K	Source	Mean log likelihood	Runtime (s)
BTC	NHHM	2	full_sample	-5247.190	195.496
BTC	NHHM	3	full_sample	-5226.359	233.634
BTC	NHHM	4	robustness_run	-5223.992	122.708
ETH	HHM	2	full_sample	-5933.445	110.178
ETH	HHM	3	full_sample	-5899.157	111.587
ETH	HHM	4	robustness_run	-5883.551	77.622
SOL	HHM	2	full_sample	-5442.919	88.120
SOL	HHM	3	full_sample	-5429.781	90.093
SOL	HHM	4	robustness_run	-5432.724	64.192

Note. For BTC and ETH, K=4 improves in-sample log likelihood, but the final freeze retains K=3 for parsimony and out-of-sample density performance. For SOL, HHM (3) dominates HHM (2) and HHM (4) in this robustness table. Source: outputs/statecount_robustness/statecount_preferred_comparison.csv.

Table 17 addresses the shorter SOL sample by evaluating the preferred model structure on a matched sample. The matched-sample evidence supports the same broad model architecture, indicating that the preferred specifications are not an artifact of unequal sample lengths.

Table 17

Matched-Sample Preferred-Model Robustness

Asset	Model	Full sample mean log likelihood	Matched sample mean log likelihood	Full runtime (s)	Matched runtime (s)	Full saved draws	Matched saved draws
BTC	NHHM (3)	-5226.359	-4176.438	233.634	99.659	75	75
ETH	HHM (3)	-5899.157	-4626.796	111.587	69.035	75	75
SOL	HHM (3)	-5429.781	-5429.781	90.093	90.093	75	75

Note. Matched-sample log likelihoods should be compared cautiously because the sample length changes. The check retains the same preferred model structure for BTC, ETH, and SOL. Source: outputs/matched_sample_robustness/matched_sample_preferred_comparison.csv.

Table 18 isolates the BTC predictor-block mechanism. The stablecoin-only specification has a stronger mean log likelihood than the macro-only and full specifications, indicating that stablecoin variables carry substantial information for BTC's regime-dependent return process. This supports the paper's broader contribution: stablecoin liquidity should be treated as a crypto-native regime signal, not merely as a control variable.

Table 18

BTC Mechanism Robustness Across Predictor Blocks

Specification	Mean log likelihood	Runtime (s)	Nonzero mean coefficients	Nonzero transition coefficients
Stablecoin only	-5222.188	109.679	6	28
Full	-5226.359	233.634	9	49
Macro only	-5229.099	110.820	3	13

Note. The stablecoin-only block has the best mean log likelihood among the three BTC mechanism specifications, supporting the claim that stablecoin variables carry predictive content. Source: outputs/mechanism_robustness/btc_mechanism_comparison.csv.

Table 19 qualifies the strength of the parameter interpretation. ETH has the strongest convergence profile, SOL is more mixed, and BTC's NHHM diagnostics require caution, especially for transition coefficients. These diagnostics do not overturn the forecast results, but they affect how the coefficient evidence should be discussed. The strongest claims should emphasize density performance, regime separation, and the clearest BTC and ETH mean-equation effects.

Table 19

Multi-Chain Convergence Diagnostics for Final and Deeper Publication Runs

Source run	Label	Asset	Model	Chains	Saved / chain	Median R-hat	Max R-hat	Median bulk ESS	Median tail ESS	Max MCSE mean	Runtime total (s)
Regularized publication run	BTC__NHHM__K3_REGSM	BTC	NHHM (3)	4	400	1.830	3.190	6.000	19.000	0.632	1420.194
Regularized publication run	ETH__HHM__K3	ETH	HHM (3)	4	400	1.020	1.370	219.000	451.000	0.944	1119.565
Regularized publication run	SOL__HHM__K3	SOL	HHM (3)	4	400	1.090	1.800	40.000	118.000	11.287	945.514
BTC small-core rerun	BTC__NHHM__K3_SMALLCORE_SCAPG_STXG_PR025	BTC	NHHM (3)	4	600	1.050	1.910	83.000	250.000	0.710	2425.520
ETH/SOL deeper run	ETH__HHM__K3_DEEPER	ETH	HHM (3)	4	600	1.020	1.210	205.000	569.000	1.254	14528.132
ETH/SOL deeper run	SOL__HHM__K3_DEEPER	SOL	HHM (3)	4	600	1.140	1.410	22.000	91.000	9.505	12850.500
SOL-only deeper run	SOL__HHM__K3_DEEPER_V3	SOL	HHM (3)	4	800	1.180	1.850	20.000	118.000	9.139	7894.818

Note. R-hat, ESS, and MCSE are from the retained publication diagnostic outputs. BTC regularized diagnostics remain mixed; the BTC small-core run improves median R-hat but has elevated max R-hat. ETH deeper diagnostics are comparatively strong, whereas SOL diagnostics remain mixed.

Table 21 formalizes the final reporting discipline. The paper should report BTC NHHM (3), ETH HHM (3), and SOL HHM (3); emphasize log-score and CRPS gains; avoid claiming uniform MSFE dominance; exclude removed Student-t robustness outputs; treat SOL coefficients cautiously; and present BTC transition coefficients as supplementary rather than central. This keeps the empirical claims aligned with the strongest evidence in the replication package.

Table 21

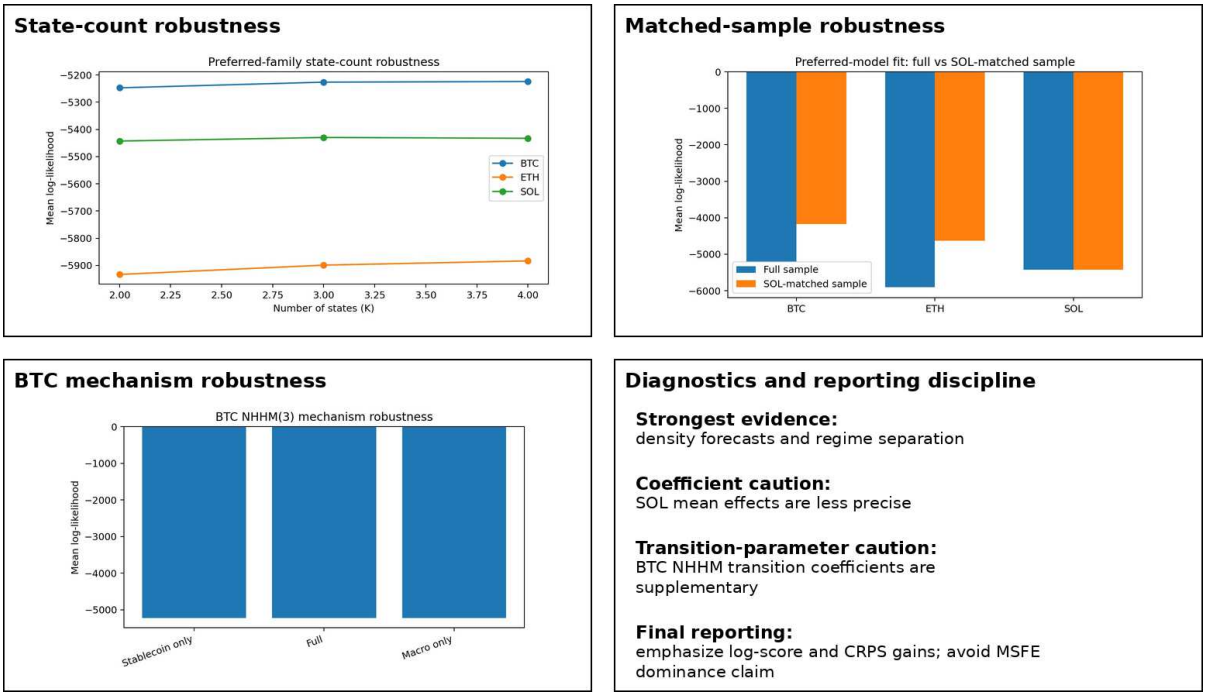
Final Reporting Decisions and Exclusions

Item	Decision	Reason
Controlling forecast table	Use outputs/final_model_freeze/table_6_final_forecast_contest.csv	It incorporates the BTC small-core candidate and defines the frozen preferred set.
Preferred models	Report BTC NHHM (3), ETH HHM (3), and SOL HHM (3)	These are the models frozen in outputs/final_model_freeze/frozen_preferred_models.csv.
Forecast interpretation	Emphasize density-forecast gains	Preferred models rank first by log score and CRPS; MSFE gains are not uniform.
Student-t robustness	Do not report	The top-level package status says those outputs were removed because activation was not confirmed.
SOL coefficients	Report with caution	Base mean-equation credible intervals do not exclude zero for non-constant coefficients.
BTC transition coefficients	Append as supplementary, interpret cautiously	Full-transition NHHM diagnostics are mixed; small-core diagnostics improve median R-hat, but max R-hat remains elevated.

Note. This table translates the replication-package status files into manuscript reporting decisions. The removed Student-t robustness outputs should not be cited.

Figure 8

Robustness and Diagnostic Summary.



Note. This figure summarizes the state-count, matched-sample, mechanism, and convergence evidence from Tables 16-19 and Table 21.

Figure 8 should present robustness and diagnostic evidence as a hierarchy. It should distinguish between the strongest evidence, which is the density-forecast and regime-separation result, and the more cautious evidence, which concerns SOL coefficient interpretation and BTC transition-parameter interpretation.

5.5 Summary of Results

The empirical results support three conclusions. First, BTC, ETH, and SOL returns exhibit economically meaningful latent volatility regimes, as shown in Table 7 and Figures 5-7. Second, the preferred hidden-state models improve out-of-sample predictive densities, as shown in Tables 4-6 and summarized visually in Figure 4. Third, stablecoin liquidity has state-dependent predictive content, especially for BTC and ETH in medium-volatility regimes, as shown in Table 8 and Figure 8. The strongest interpretation is therefore not that stablecoin variables generate uniform return predictability. Rather, the evidence shows that stablecoin liquidity, network congestion, and macro-financial stress matter most when their effects are allowed to vary across latent crypto-market regimes.

6. Discussion

The main contribution of the paper is not a general claim that stablecoin, network, and macro-financial variables predict cryptocurrency returns in all settings. The stronger result is that stablecoin liquidity helps identify latent market conditions, and that this information is most useful when returns are modeled as regime dependent rather than as outcomes of a single linear process. The preferred specifications: BTC NHHM (3), ETH HHM (3), and SOL HHM (3) rank first on the paper's primary density criteria, average log score and CRPS, even though the

MSFE evidence is mixed. This matters because, in markets characterized by volatility clustering, abrupt repricing, and asymmetric tail risk, the quality of the predictive distribution is often more informative than the conditional mean alone (Geweke & Amisano, 2010; Gneiting & Raftery, 2007; Gneiting & Ranjan, 2011).

This interpretation also sharpens the economic role of stablecoins. The results do not support a single average effect of stablecoin activity across all assets and all states. Instead, they suggest that stablecoin variables are informative because they reveal different dimensions of crypto-market liquidity. In the preferred BTC and ETH models, stablecoin volume growth is positively associated with returns in the medium-volatility regime, while stablecoin transaction-count growth is negatively associated with returns in that same regime. A plausible reading is that volume growth captures deployable trading liquidity or capital availability, whereas transaction growth reflects stress-related repositioning, settlement fragmentation, or congestion pressure. The medium-volatility regime is therefore economically important because it appears to be a transition state in which liquidity conditions become informative before extreme stress fully dominates the return process, a pattern consistent with broader theories of state-dependent liquidity and funding pressure (Brunnermeier & Pedersen, 2009).

The cross-asset evidence reinforces that conclusion by showing that the same mechanism does not operate identically in every market. Bitcoin provides the clearest case of transition-sensitive regime dynamics, which is consistent with its role as the most liquid and institutionally integrated cryptoasset. Ether supports the same mechanism more through state-dependent mean effects than through exogenous transition dynamics. Solana contributes primarily through regime separation and density-forecast improvement, not through sharply estimated coefficients. The paper should therefore avoid an overly uniform interpretation. The most defensible conclusion is that stablecoin-linked variables provide the clearest predictor evidence for BTC and ETH, while latent volatility regimes improve distributional forecasting across BTC, ETH, and SOL. In that sense, the findings extend the hidden-state cryptocurrency literature (Hamilton, 1989; Kim & Nelson, 1999; Koki et al., 2022) and complement the stablecoin literature that increasingly treats stablecoins as part of crypto-market infrastructure rather than as peripheral trading instruments (Barucci et al., 2022; Lyons & Viswanath-Natraj, 2023; Gubareva et al., 2023; Paeng et al., 2024).

A future extension would test stablecoin liquidity directly as an early-warning signal for volatility-regime transitions. The current evidence shows that stablecoin variables are informative in state-dependent return models, especially for BTC and ETH in medium-volatility regimes, but it does not directly test whether stablecoin variables anticipate transitions from low-volatility states into medium- or high-volatility states better than congestion or macro-financial variables. Future work could define a transition event as $Transition_{i,t} = 1$ when $s_{i,t-1} = 1$ and $s_{i,t} \in \{2,3\}$, where State 1 is the low-volatility regime and States 2 and 3 are the medium- and high-volatility regimes. Stablecoin-only, congestion-only, macro-only, and full predictor-block models could then be compared using transition AUC, average precision, Brier score, and log loss. This extension would directly test whether stablecoins function not only as state-dependent predictors, but as forward-looking signals of latent crypto-market risk.

The scope of the current findings should remain disciplined. The preferred models do not uniformly dominate on MSFE, Solana's coefficient evidence is weaker than its regime evidence, and Bitcoin's richer transition structure should be interpreted cautiously because non-homogeneous models are more parameter intensive. These qualifications do not weaken the contribution; they clarify it. The strongest claim supported by the paper is that stablecoin-linked liquidity conditions help reveal latent crypto-market risk states and that incorporating those signals improves the density forecasting of cryptocurrency returns.

7. Conclusion

This paper examined whether stablecoin liquidity, network congestion, and macro-financial stress help explain the daily returns of Bitcoin, Ether, and Solana when market conditions change discretely over time. The results support a clear conclusion: cryptocurrency returns are better described by a three-state regime structure than by a simpler two-state alternative, and the most useful forecasting gains come through improved predictive densities rather than uniform gains in point-forecast accuracy. In the final freeze, the preferred models are BTC NHHM (3), ETH HHM (3), and SOL HHM (3), implying that exogenous transition dynamics are economically relevant but not universal.

The substantive contribution of the paper is therefore to show that stablecoin liquidity acts as a crypto-native regime signal. Stablecoin variables do not have one fixed effect across the full sample. Instead, they become informative in states, especially transition-like medium-volatility regimes—where liquidity conditions, market positioning, and stress are changing. This moves the contribution beyond generic return prediction. Relative to the hidden-state literature, the paper links regime dynamics to a sharper market-structure mechanism; relative to the stablecoin literature, it shows that stablecoins matter not only as trading instruments or safe assets, but also as part of the information structure through which crypto-market risk is revealed (Hamilton, 1989; Kim & Nelson, 1999; Barucci et al., 2022; Koki et al., 2022).

The paper also defines its own limits. The preferred models are strongest for distributional forecasting, risk-state identification, and warning content, not for universal MSFE dominance or uniformly strong trading gains. Solana's evidence is most persuasive at the regime level, while Bitcoin provides the clearest case for transition-sensitive dynamics. Taken together, the findings support one central conclusion: stablecoin-linked liquidity conditions contain information about latent crypto-market states, and regime-based models that incorporate that information improve the characterization of cryptocurrency return risk.

Funding Declaration
This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

References

Barucci, E., Giuffra Moncayo, G., & Marazzina, D. (2022). Cryptocurrencies and stablecoins: A high-frequency analysis. *Digital Finance, 4*, 217–239. https://doi.org/10.1007/s42521-022-00055-9

Brunnermeier, M. K., & Pedersen, L. H. (2009). Market liquidity and funding liquidity. *The Review of Financial Studies, 22*(6), 2201–2238. https://doi.org/10.1093/rfs/hhn098

Coin Metrics. (n.d.). *Coin Metrics API v4* [Data set]. Retrieved May 14, 2026, from https://docs.coinmetrics.io/api/v4/

Diebold, F. X., & Mariano, R. S. (1995). Comparing predictive accuracy. *Journal of Business & Economic Statistics, 13*(3), 253–263. https://doi.org/10.1080/07350015.1995.10524599

Federal Reserve Bank of St. Louis. (n.d.). *FRED: Federal Reserve Economic Data* [Data set]. Retrieved May 14, 2026, from https://fred.stlouisfed.org/

Filardo, A. J. (1994). Business-cycle phases and their transitional dynamics. *Journal of Business & Economic Statistics, 12*(3), 299–308. https://doi.org/10.1080/07350015.1994.10524545

Geweke, J., & Amisano, G. (2010). Comparing and evaluating Bayesian predictive distributions of asset returns. *International Journal of Forecasting, 26*(2), 216–230. https://doi.org/10.1016/j.ijforecast.2009.10.007

Gneiting, T., & Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association, 102*(477), 359–378. https://doi.org/10.1198/016214506000001437

Gneiting, T., & Ranjan, R. (2011). Comparing density forecasts using threshold- and quantile-weighted scoring rules. *Journal of Business & Economic Statistics, 29*(3), 411–422. https://doi.org/10.1198/jbes.2010.08110

Griffin, J. M., & Shams, A. (2020). Is Bitcoin really untethered? *The Journal of Finance, 75*(4), 1913–1964. https://doi.org/10.1111/jofi.12903

Gubareva, M., Bossman, A., & Teplova, T. (2023). Stablecoins as the cornerstone in the linkage between the digital and conventional financial markets. *The North American Journal of Economics and Finance, 68*, Article 101979. https://doi.org/10.1016/j.najef.2023.101979

Hamilton, J. D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica, 57*(2), 357–384. https://doi.org/10.2307/1912559

Kim, C.-J., & Nelson, C. R. (1999). *State-space models with regime switching: Classical and Gibbs-sampling approaches with applications*. MIT Press. https://doi.org/10.7551/mitpress/6444.001.0001

Koki, C., Leonardos, S., & Piliouras, G. (2022). Exploring the predictability of cryptocurrencies via Bayesian hidden Markov models. *Research in International Business and Finance, 59*, Article 101554. https://doi.org/10.1016/j.ribaf.2021.101554

Lyons, R. K., & Viswanath-Natraj, G. (2023). What keeps stablecoins stable? *Journal of International Money and Finance, 131*, Article 102777. https://doi.org/10.1016/j.jimonfin.2022.102777

Makarov, I., & Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics, 135*(2), 293–319. https://doi.org/10.1016/j.jfineco.2019.07.001

Paeng, S., Senteney, D., & Yang, T. (2024). Spillover effects, lead and lag relationships, and stable coins time series. *The Quarterly Review of Economics and Finance, 95*, 45–60. https://doi.org/10.1016/j.qref.2024.03.003

Rabiner, L. R. (1989). A tutorial on hidden Markov models and selected applications in speech recognition. *Proceedings of the IEEE, 77*(2), 257–286. https://doi.org/10.1109/5.18626

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- StablecoinRegimeFinalSubmissionBundle20202026.zip