Trading Volume, Buying Pressure and Cryptocurrency Return Distributions: Evidence from Markov-Switching Models

Madueke Chinwe Monica [1], Ezenwobi Florence Ngozi [2], Nwokoye Mathew Okechukwu [3], Onwuka Irene Nkechi [4], Okaforocha Chika Maureen [5], & Mba Ifeoma Christy [6*]

[1-5] Department of Economics, Faculty of Social Sciences, Nnamdi Azikiwe University, Awka, Anambra State, Nigeria.

[6*(corresponding author)] Department of Economics, Faculty of the Social Sciences, University of Nigeria, Nsukka, Enugu State, Nigeria.

Abstract

Daily trading activity may offer insight into changing cryptocurrency risk, but not necessarily into future returns. This study examines that distinction for Bitcoin, Ethereum, Solana and BNB using daily Binance data from January 2021 to September 2026. Two-state Gaussian Markov-switching regressions relate returns to lagged returns, trading volume and taker-buy pressure, and trading variables affect the probabilities of regime switching. Forecasts are compared against historical-mean and Gaussian linear benchmarks over January 2025-September 2026. The estimates identify lower-volatility and higher-volatility states, with shorter expected durations in the higher-volatility regime for all four assets. For Ethereum and BNB, volatility transitions are driven by trading volume, while for Bitcoin and Solana, they are driven by taker-buy pressure. Bitcoin also shows a stronger positive relationship between lagged buy pressure and returns under higher-volatility conditions. These in-sample relationships do not carry over into consistent improvements in return magnitude or directional forecasts. Overall, regime models improve predictive density scores over the Gaussian linear benchmark, but buy pressure provides no statistically reliable incremental density-forecasting benefit across the four assets. The results support an asset-specific evaluation of trading indicators and favour assessing regime models by their ability to describe changing return distributions, rather than assuming that regime identification leads to accurate return forecasts.

Keywords: Cryptocurrency; Markov-switching; Trading volume; Taker-buy pressure; Volatility regimes

I. INTRODUCTION

Activity in cryptocurrency trading could offer a better insight into risk changes than the next day's performance. To assess its value, one must disentangle changes in projected return distributions from improvements in mean or directional forecasts. Evidence on momentum, investor attention and signed trade flows suggests that market behaviour contains information about cryptocurrency returns (Liu & Tsyvinski, 2021; Makarov & Schoar, 2020). Direct tests are needed to see whether this information remains useful in a changing market environment and to improve out-of-sample forecasts. Even if a model accurately describes turbulent periods, it may offer little insight into the next price move.

Trading volume and buying pressure are two different types of activity. Volume is the amount exchanged, and the taker-buy share is the fraction of orders initiated by buyers. Anastasopoulos et al. (2026) and Makarov and Schoar (2020) link order flow and signed volume to cryptocurrency returns, setting the context for exploring trading pressure relative to overall activity. These findings do not establish whether the Binance taker-buy imbalance examined in this study predicts next-day returns. International order flow aggregates data across multiple marketplaces, while exchange-level taker-buy volume reflects a smaller subset of trades.

The information these transactions carry may change between relatively calm and turbulent periods. A constant-coefficient regression imposes the same relationship throughout the sample, potentially concealing differences between market states. Markov-switching models allow selected return parameters to vary with an unobserved regime, while time-varying transition probabilities let observed predictors help explain regime changes (Filardo, 1994; Hamilton, 1989). Trading activity can consequently be examined both as a predictor of returns within a state and as a predictor of movement between states. These channels need not provide the same forecasting benefit.

Cryptocurrency applications support examining both channels. Researchers have incorporated volume and Google search activity into time-varying transition probability models of Bitcoin volatility, while hidden Markov models have identified return regimes and improved predictive densities for Bitcoin, Ether, and Ripple (Koki et al., 2022; Tan et al., 2021). These findings establish the relevance of regime dynamics but leave a specific question for the present analysis: does taker-based buying pressure add useful information once past returns, volume and regime changes are included? Addressing that question requires comparing otherwise similar models with and without buying pressure, rather than attributing every gain from regime switching to the trading signal.

We make this comparison for Bitcoin, Ether, Solana and BNB. The analysis addresses three questions: which return and volatility regimes characterize the four assets; whether trading volume and buying pressure predict transitions between them; whether buying pressure provides additional information about next-day returns, with effects that differ across regimes; and whether these

relationships improve out-of-sample forecasts relative to simpler benchmarks. The corresponding objectives are to characterize regime persistence and its predictors, estimate the additional information associated with buying pressure, and evaluate the forecasting gains from the competing specifications.

The distinction between fitting and forecasting is central to these comparisons. A statistically significant coefficient does not, by itself, establish that a predictor improves subsequent forecasts. Likewise, lower point-forecast errors and better predictive densities measure different aspects of performance. We therefore evaluate return errors, the probability of a positive return and predictive log scores. Proper scoring rules provide a basis for evaluating the probabilities assigned to realized outcomes (Gneiting & Raftery, 2007). Together, these measures let us determine whether a model improves the expected return forecast, the description of uncertainty, or both.

The statistical tests examine null hypotheses that each trading predictor's transition coefficients are jointly zero, that buying-pressure coefficients in the return equations are jointly zero, and that those coefficients are equal across regimes. Forecast comparisons test equal predictive performance between specified models. We assess evidence at the 5% level and adjust the reported families of multiple comparisons. The estimated regimes describe variation under the chosen model; they do not establish the true number of market states. Relationships between trading activity and returns are interpreted as predictive associations.

The dataset contains 8,396 daily asset-day observations across four Binance spot trading pairs; BTC/USDT, ETH/USDT, SOL/USDT and BNB/USDT from 1 January 2021 to 30 September 2026. Returns are calculated from closing prices, while buying pressure is constructed from taker-buy and total quote volumes denominated in USDT. We estimate Markov-switching regressions with constant and time-varying transition probabilities and compare their forecasts with historical-mean and linear regression benchmarks. All predictors enter with a one-day lag. For the final out-of-sample evaluation, model parameters are estimated using observations through 31 December 2024, with 1 January 2025 to 30 September 2026 reserved as the test period.

The results show clearer gains in distribution forecasts than in mean or directional forecasts. Markov models improve predictive log scores relative to the fitted Gaussian linear benchmark, but improvements in exact-return and directional accuracy are inconsistent. In the full sample, buying pressure is regime-dependent in its association with Bitcoin returns but does not deliver a reliable incremental forecasting gain across the four assets. The study thus identifies where trading information is useful and where its apparent explanatory value fails to carry into forecasts. Its contribution rests on separating the information supplied by buying pressure from the gains attributable to modelling changing market regimes.

II. Literature review

Trading activity can contain information about cryptocurrency returns, but its usefulness depends on what the measure captures. Makarov and Schoar (2020) connect signed trading volume to Bitcoin price movements, while Anastasopoulos et al. (2026) report predictive information in international cryptocurrency order flows. These findings support examining transaction direction alongside total trading volume. However, international flows aggregate information across markets, whereas a taker-buy imbalance from one exchange describes a narrower set of transactions. We must therefore test whether this exchange-level measure improves subsequent forecasts after accounting for information already contained in returns and volume.

The persistence of market conditions affects how such relationships should be interpreted. Hamilton's (1989) framework allows the return-generating process to change across unobserved regimes. Within Bitcoin research, Mba et al. (2018) examine persistence through mean sojourn and recurrence times in buy, hoard and sell states. Eze et al. (2026) approach the problem through a Gaussian hidden Markov model, identifying bullish, stable and bearish regimes with different persistence patterns. The classifications and estimation approaches differ, but both studies make duration central to understanding market behaviour. For risk assessment, identifying a turbulent condition becomes more informative when its likely persistence is also considered.

Identifying these conditions and demonstrating forecasting gains nevertheless require different evidence. Mba et al. (2021) characterize Bitcoin's bullish and bearish phases through turning points and durations, describing market transitions. Koki et al. (2022) extend the forecasting comparison by evaluating predictive densities for Bitcoin, Ether and Ripple, reporting stronger one-step-ahead performance from a four-state non-homogeneous hidden Markov specification. Their results illustrate how regime analysis can move from describing historical conditions to evaluating predictions. This distinction matters for the present study: an estimated regime, or persistent market phase, does not show that buying pressure adds information to forecasts. Its contribution must be isolated by comparing it with models that exclude it.

Observable predictors can also help explain movement between regimes. Filardo (1994) develops a framework in which transition probabilities depend on available information, and Tan et al. (2021) incorporate trading volume and Google search activity into time-varying transition probability Markov-switching GARCH models of Bitcoin volatility. This approach lets us examine trading activity through both transition probabilities and the return equation. For our analysis, it raises a testable possibility: volume or buying pressure may help anticipate regime changes even when its relationship with the conditional mean is weak. Whether either channel improves forecasts is an empirical question that requires separate comparisons of constant-transition and predictor-dependent specifications.

Forecast evaluation determines what those comparisons reveal. Proper scoring rules assess the probabilities assigned to realized outcomes, complementing measures of point-forecast error (Gneiting & Raftery, 2007). Evaluating predictive log scores, return errors, and positive-return probabilities together can establish whether a model improves its description of uncertainty, its expected-return forecast, or both. The present study applies

these comparisons to Bitcoin, Ether, Solana and BNB. It examines whether buying pressure contributes information after past returns, volume and regime dynamics are considered, and whether that contribution survives outside the estimation sample. This connects earlier evidence on market-state persistence to a practical forecasting question: which trading signals remain useful when predicting observations that were unavailable during model estimation?

III. Methodology

3.1 Data and sample preparation

We analyse daily observations for Bitcoin, Ether, Solana and BNB traded against USDT on Binance. The sample runs from 1 January 2021 to 30 September 2026 and contains 2,099 observations per asset, giving 8,396 asset-days. The source workbook records closing prices, total quote volume and taker-buy quote volume, together with the underlying archive filenames. These fields correspond to Binance's documented spot kline structure (Binance, n.d.). Models are estimated separately for each pair, allowing coefficients and transition dynamics to differ across assets.

Records were ordered by UTC date and checked for duplicate and missing dates, missing required numeric values, invalid price relationships and inconsistent trading volumes. Taker-buy volume was required to lie between zero and total volume. No internal violations were found in these checks, and no missing dates required filling. The checks establish consistency within the supplied dataset; its observations were not independently reconciled with exchange archives or official checksum files. Constructing returns removes the first price observation, and lagging returns removes one further observation, leaving 2,097 regression observations per asset from 3 January 2021.

3.2 Returns and trading measures

Daily returns are measured as percentage log changes in closing prices:

$$r_t = 100 \ln\left(\frac{C_t}{C_{t-1}}\right) \quad (1)$$

Here, C_t is the closing price on day t. Trading activity is represented by the logarithm of total quote volume, Q_t, measured in USDT:

$$V_t = \ln(Q_t) \quad (2)$$

Buying pressure is calculated from taker-buy quote volume, B_t, using the same units as total quote volume:

$$BP_t = 2\left(\frac{B_t}{Q_t}\right) - 1 \quad (3)$$

BP_t ranges from −1 to 1. Positive values indicate that buyer-initiated trading exceeds seller-initiated trading in quote-volume terms. The measure indicates which side initiates transactions; it does not directly measure net capital inflows or investor sentiment. Log volume and buying pressure are standardized using the means and population standard deviations of their lagged values within the relevant estimation window. The notation $z_{v,t-1}$ and $z_{\beta,t-1}$ denotes these standardized predictors. The estimation-window scaling is retained during forecast evaluation, and all predictors enter with a one-day lag.

3.3 Markov-switching specification

We use a two-state Markov-switching regression following Hamilton's (1989) regime framework. For clarity, the asset index is omitted:

$$r_t = \alpha_{S_t} + \varphi r_{t-1} + \beta_{S_t} z_{V,t-1} + \gamma_{S_t} z_{B,t-1} + \varepsilon_t \quad (4)$$

$$\varepsilon_t \mid S_t = j, I_{t-1} \sim \mathrm{N}\left(0, \sigma_j^2\right),\ j \in \{1,2\} \quad (5)$$

S_t denotes the unobserved state and I_{t-1} the information available before the return on day t is observed. The intercept, coefficients on included trading predictors, and innovation variance vary across states. The lagged-return coefficient, φ, is common to both states. We identify states by ordering their innovation standard deviations so that $\sigma_1 < \sigma_2$. State 1 is therefore the lower-volatility state and state 2 the higher-volatility state; we do not assign either a bullish or bearish label in advance. Two states are fixed for this analysis rather than selected through a formal test of the number of regimes.

The constant-transition models estimate $p_{ij} = \Pr(S_t = j \mid S_{t-1} = i)$. Following Filardo (1994), the variable-transition models allow lagged trading predictors to enter the probability of moving into state 2:

$$p_{i2,t} = \frac{\exp(\eta_{i,t})}{1 + \exp(\eta_{i,t})},\ i \in \{1,2\} \quad (6)$$

$$\eta_{i,t} = a_i + b_i z_{V,t-1} + c_i z_{B,t-1} \quad (7)$$

The remaining transition probability is $p_{i1,t} = 1 - p_{i2,t}$. The volume-only specification omits buying pressure from both the return and transition equations. This produces five Markov specifications:

Model	Predictors in the return equation	Transition probabilities
MS AR	Lagged return	Constant
MS V	Lagged return and volume	Constant
MS V BP	Lagged return, volume and buying pressure	Constant
TVTP V	Lagged return and volume	Depend on lagged volume
TVTP V BP	Lagged return, volume and buying pressure	Depend on lagged volume and buying pressure

For constant-transition models, expected state duration is calculated as:

$$D_j = \frac{1}{1 - p_{jj}} \quad (8)$$

This expression is not applied to the variable-transition models, whose persistence changes with the predictors. Comparisons

between the two variable-transition specifications assess the joint effect of buying pressure on returns and transitions, rather than identifying either channel separately.

3.4 Estimation and statistical inference

Parameters are estimated by maximizing the Gaussian mixture likelihood evaluated through the Hamilton forward filter. Each estimation starts with equal probabilities of occupying the two states. We use bounded L-BFGS-B optimization, retaining the solution with the highest likelihood across four starting values for initial and calibration fits and five for full-sample fits. We cap iterations at 700. Innovation standard deviations are parameterized as $\sigma_1 = \exp(u)$ and $\sigma_2 = \exp(u) + \exp(v)$, preserving their positivity and ordering.

The common autoregressive coefficient is bounded between −0.98 and 0.98. State intercepts and slopes on standardized predictors are bounded between −30 and 30; transition logit coefficients between −8 and 8; u between −3 and 3.8; and v between −3 and 4.4. Convergence, starting-value sensitivity and proximity to parameter bounds are checked for every fit. AIC and BIC describe in-sample fit after allowing for model complexity. We use full-sample estimates to describe regimes and examine coefficient restrictions, keeping them separate from the estimates used for held-out forecasts.

Coefficient inference uses a numerical observed Hessian and a heteroskedasticity- and autocorrelation-consistent score sandwich with Bartlett weights and seven lags, following the covariance-estimation principle of Newey and West (1987). Joint Wald tests assess whether each predictor's transition coefficients are zero, whether buying-pressure coefficients in the return equations are jointly zero, and whether those return coefficients are equal across states. The tests are approximate and conditional on the fitted two-state specification. We use a 5% significance level and report Benjamini–Hochberg adjustments within the stated test families (Benjamini & Hochberg, 1995). The two return-equation hypothesis families each contain eight tests across four assets and two buying-pressure models; the buying-pressure transition family contains four tests.

3.5 Forecast design

The chronological design preserves the information available at each forecast origin:

Stage	Period	Observations per asset
Initial estimation	3 January 2021–31 December 2023	1,093
Validation	1 January–31 December 2024	366
Final calibration	3 January 2021–31 December 2024	1,459
Final evaluation	1 January 2025–30 September 2026	638

The validation period provides an initial assessment of forecast behaviour. For final evaluation, model parameters and predictor scaling are fixed at their calibration estimates. State probabilities are updated sequentially as each test-period return becomes available. We do not re-estimate parameters during the evaluation period. All specifications are reported. The volume-only variable-transition model was added after inspecting the initial results as an exploratory comparison to assess the additional contribution of buying pressure; it was not a prespecified comparison.

One-step-ahead state probabilities are calculated from the previous filtered probabilities:

$$\pi_{j,t|t-1} = \sum_{i=1}^{2} p_{ij,t}\pi_{i,t-1|t-1} \quad (9)$$

The return forecast is the weighted average of the state-conditional means:

$$\hat{r}_{t|t-1} = \sum_{j=1}^{2} \pi_{j,t|t-1}\hat{\mu}_{j,t} \quad (10)$$

The corresponding predictive density is:

$$\hat{f}(r_t \mid I_{t-1}) = \sum_{j=1}^{2} \pi_{j,t|t-1}\phi\left(r_t;\ \hat{\mu}_{j,t}, \hat{\sigma}_j^{\,2}\right) \quad (11)$$

Here, $\hat{\mu}_{j,t}$ is the conditional mean implied by Equation (4), and ϕ is the Gaussian density. The predictive mixture also supplies the probability that the next return is positive. Forecasts use filtered information available at the time; smoothed probabilities incorporating later returns are excluded.

3.6 Forecast evaluation and verification

We compare the Markov forecasts with a historical-mean benchmark and a constant-coefficient linear regression containing lagged returns, volume and buying pressure. Both benchmarks are estimated within the same calibration window. Their predictive densities use a single Gaussian distribution, with variance estimated from the calibration returns or residuals, as appropriate. Point forecasts are evaluated using RMSE and MAE, distribution forecasts using mean predictive log scores, and positive-return probabilities using Brier scores. Directional accuracy classifies a return as positive when its predicted probability is at least 0.5. This is distinct from classifying the sign of the conditional-mean forecast. The scoring approach follows Gneiting and Raftery (2007).

Forecast comparisons use paired daily log-score or squared-error differences with seven-lag HAC standard errors. Positive log-score gains favour the first model; positive squared-error improvements indicate lower error for that model. Reported confidence intervals and two-sided p-values are approximate and condition on the estimation-window fits. Benjamini–Hochberg adjustments cover the 16 test-period log-score comparisons concerning buying pressure and variable transitions across the four assets. Benchmark comparisons and other metrics are reported separately. These procedures do not replace a bootstrap accounting for estimation uncertainty.

The likelihood filter was checked against an independently written implementation, and all 20 asset-by-model Markov forecast paths were recomputed independently. We also verified that predictor scaling used calibration data only and that changing future observations left earlier filtered predictions unchanged. The analysis retains Gaussian conditional errors and a fixed two-state specification.

IV. Results and discussion

4.1 Descriptive statistics

Table 1 summarizes daily returns, trading volume and taker-buy pressure for the four cryptocurrencies. The estimation sample contains 2,097 observations per asset after constructing returns and lagged predictors.

Table 1. Descriptive statistics

Statistic	Bitcoin	Ethereum	Solana	BNB
Observations	2,097	2,097	2,097	2,097
Mean return	0.046	0.059	0.200	0.143
Return standard deviation	2.982	4.036	5.688	4.084
Minimum return	−16.700	−32.486	−54.901	−41.671
Maximum return	17.845	23.375	30.967	53.240
Return skewness	−0.119	−0.179	−0.362	0.833
Return excess kurtosis	4.121	5.587	9.198	25.121
Mean quote volume, million USDT	2,502.422	1,326.814	405.505	324.872
Quote-volume standard deviation, million USDT	1,956.523	989.725	445.989	468.245
Mean buy pressure	−0.0176	−0.0074	−0.0057	0.0028
Buy-pressure standard deviation	0.0434	0.0415	0.0411	0.0472

Source: Authors' calculations from the supplied Binance dataset.

Average daily returns are positive but small relative to their standard deviations. Solana has the highest mean return and the greatest return variability, while Bitcoin has the lowest standard deviation. BNB exhibits the largest excess kurtosis, indicating particularly pronounced departures from the normal distribution in the sample.

The return distributions therefore differ in both dispersion and tail behaviour. These characteristics motivate examining changes in conditional volatility, although descriptive statistics alone cannot establish the appropriate number of regimes. The substantial excess kurtosis also makes heavier-tailed innovations an important extension to the Gaussian specifications examined here.

Bitcoin records the highest average quote volume, followed by Ethereum. These figures describe trading activity on the supplied exchange and should not be treated as complete measures of market liquidity. Mean buy pressure is close to zero across assets, but this average can conceal daily variation and associations that differ across volatility states.

4.2 Volatility regimes and persistence

Table 2 reports regime characteristics from the constant-transition specification containing lagged returns, volume, and buy pressure. A common specification permits direct comparison across assets; it is not assumed to be the preferred model for every cryptocurrency.

Table 2. Estimated volatility regimes and persistence

Cryptocurrency	Lower-state SD	Higher-state SD	Lower-state persistence	Higher-state persistence	Lower-state duration, days	Higher-state duration, days
Bitcoin	1.594	4.439	0.870	0.770	7.69	4.35
Ethereum	2.258	6.080	0.910	0.832	11.14	5.94
Solana	3.903	10.522	0.970	0.852	33.09	6.73
BNB	2.083	7.643	0.957	0.851	23.40	6.70

Source: Authors' calculations from the supplied Binance dataset.

The estimated shock standard deviations differ substantially between regimes. Bitcoin's estimate rises from 1.594 percentage points in the lower-volatility state to 4.439 in the higher-volatility state. Solana has the largest conditional standard deviation in both states, while BNB exhibits the widest proportional difference between them.

Lower-volatility conditions persist longer than higher-volatility conditions for every asset. Solana's expected lower-state duration is 33.09 days, compared with 6.73 days in the higher-volatility state. BNB shows a similar asymmetry. Bitcoin has shorter expected durations in both states, implying more frequent switching under this specification.

These results separate shock magnitude from the duration of volatility conditions. Solana can experience relatively persistent lower-volatility periods while still exhibiting the largest shocks when conditions change. The regime labels concern conditional volatility and do not identify bullish or bearish markets. They also remain conditional on the imposed two-state structure.

4.3 Trading activity and conditional mean returns

Table 3 presents the return-equation estimates and the principal tests of buy pressure. Volume and buy pressure are standardized, so their coefficients measure associations with a one-standard-deviation increase in the previous day's predictor.

Table 3. Return-equation estimates and buy-pressure tests

Variable or test	Bitcoin	Ethereum	Solana	BNB
Panel A: coefficient estimates				
Lagged return	−0.097 (0.028)	−0.062 (0.031)	−0.040 (0.026)	−0.093 (0.027)
Intercept: lower-volatility state	0.058 (0.068)	0.110 (0.092)	−0.027 (0.114)	0.105 (0.065)
Intercept: higher-volatility state	0.034 (0.202)	0.006 (0.287)	1.286 (0.591)	0.061 (0.450)
Volume: lower-volatility state	−0.013 (0.049)	−0.039 (0.081)	0.034 (0.108)	−0.056 (0.074)
Volume: higher-volatility state	−0.071 (0.184)	−0.105 (0.257)	−1.249 (0.531)	0.252 (0.329)
Buy pressure: lower-volatility state	0.075 (0.055)	0.014 (0.077)	0.126 (0.108)	0.036 (0.049)
Buy pressure: higher-volatility state	0.674 (0.217)	0.296 (0.275)	0.788 (0.643)	0.820 (0.459)
Panel B: adjusted joint-test results				
Both buy-pressure coefficients equal zero: q	0.0263	0.5278	0.3082	0.3082
Buy-pressure coefficients equal across states: q	0.0188	0.3416	0.3416	0.1991

Source: Authors' calculations from the supplied Binance dataset.

Lagged-return coefficients are negative across the four assets. The estimates are individually significant at the conventional 5% level for Bitcoin, Ethereum and BNB, suggesting modest daily reversal after controlling for trading activity and regimes. Solana's coefficient is not statistically distinguishable from zero.

Volume has weak associations with subsequent conditional mean returns for Bitcoin, Ethereum and BNB. Solana differs in its higher-volatility state, where the volume coefficient is −1.249 (p=0.019). Nevertheless, the joint test that both Solana volume coefficients equal zero gives (p=0.063). The evidence is therefore qualified, rather than broadly establishing a volume effect.

Bitcoin provides the clearest buy-pressure result. Its higher-state coefficient is 0.674, compared with 0.075 in the lower-volatility state. Both the joint-zero and cross-state equality hypotheses are rejected after adjustment for multiple testing. Stronger buyer-initiated trading is thus more clearly associated with higher subsequent conditional mean returns during Bitcoin's higher-volatility conditions.

The corresponding adjusted tests do not reject either hypothesis for Ethereum, Solana or BNB. Consequently, positive coefficient estimates across all four assets do not establish a common predictive relationship. Nor does Bitcoin's in-sample association demonstrate that buy pressure improves forecasts or generates trading profits.

4.4 Trading activity and volatility transitions

Table 4 examines trading activity in the time-varying transition model. Both equations describe the probability of the next state being the higher-volatility regime.

Table 4. Transition-equation estimates and joint tests

Predictor or test	Bitcoin	Ethereum	Solana	BNB
Panel A: transition coefficients				
Volume: lower → higher	0.351 (0.210)	0.881 (0.206)	−0.083 (0.281)	1.232 (0.228)
Volume: higher → higher	0.147 (0.380)	1.033 (0.265)	−0.484 (0.300)	0.883 (0.772)
Buy pressure: lower → higher	0.418 (0.158)	0.112 (0.155)	−0.107 (0.249)	−0.041 (0.142)
Buy pressure: higher → higher	−1.307 (0.560)	−0.618 (0.549)	−1.323 (0.575)	−0.770 (0.518)
Panel B: joint tests across transition equations				
Volume coefficients jointly zero: p	0.1026	<0.0001	0.0767	<0.0001
Buy-pressure coefficients jointly zero: q	0.0345	0.5265	0.0345	0.4072

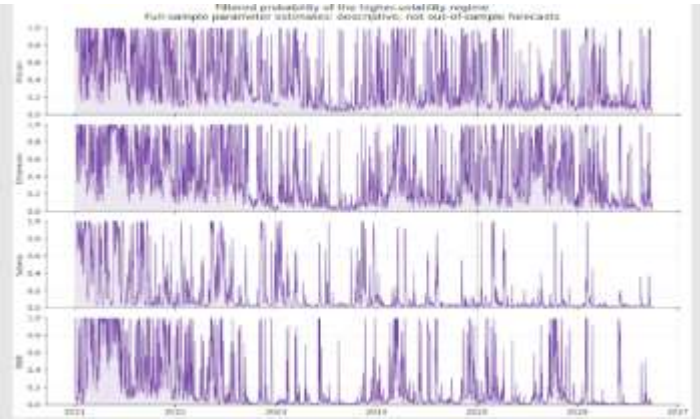
Notes: Panel A reports logistic coefficients with HAC standard errors using seven lags in parentheses. Predictors are standardized and lagged by one day. Positive coefficients increase the log odds of entering or remaining in the higher-volatility state. Transition intercepts are omitted. Panel B reports approximate Wald tests; volume values are unadjusted, while buy-pressure values are adjusted across the four assets.

Source: Authors' calculations from the supplied Binance dataset.

Volume contributes to Ethereum's and BNB's transition equations. For Ethereum, higher volume is associated with both entering and staying in the higher-volatility state. For BNB, the individual evidence is concentrated in entry from the lower-volatility state. These findings contrast with their weak volume associations in the return equation: trading activity appears more informative about volatility conditions than about the direction of subsequent returns.

Buy pressure contributes jointly to the transition equations for Bitcoin and Solana after adjustment. For Bitcoin, its positive lower-to-higher coefficient and negative higher-to-higher coefficient indicate that the relationship depends on the starting state. For Solana, the individual evidence is concentrated in a negative association with higher-state persistence. These patterns do not identify the economic mechanism responsible for the transitions.

Figure 1. Filtered probabilities of the higher-volatility regime.



Notes: Probabilities cover 3 January 2021–30 September 2026 and use the model allowing volume and buy pressure to influence transitions. Values approaching one indicate a greater estimated probability of the higher-volatility state. Parameters are estimated from the full sample, so these paths are descriptive fitted probabilities. The dashed line marks 1 January 2025, the beginning of the separate forecast evaluation period.

Source: Authors' calculations from the supplied Binance dataset.

Figure 1 shows frequent elevated higher-state probabilities during 2021 across all four markets. Later observations display differences in their timing and concentration. Solana and BNB generally show extended stretches of low probabilities interrupted by sharp increases, while Bitcoin and Ethereum fluctuate more often. Some increases occur around similar dates, but separate asset models cannot establish volatility transmission or contagion.

4.5 Model comparison and forecasting performance

Model comparison does not favour the same specification for every asset. Full-sample BIC selects the constant-transition model containing only lagged returns for Bitcoin and Solana. For Ethereum and BNB, it selects the specification allowing volume to enter both return and transition equations. Adding buy pressure to that specification increases BIC for Ethereum, Solana and BNB. For Bitcoin, the reduction is insufficient to justify moving away from the simpler autoregressive model.

These findings qualify the coefficient results: statistically significant individual associations do not necessarily justify a larger model after penalizing complexity. Forecast evaluation provides a separate assessment using 638 observations from 1 January 2025 to 30 September 2026. Parameters and standardization values are estimated using observations through December 2024 and remain fixed during evaluation, while regime probabilities update sequentially.

Table 5. Out-of-sample forecasting performance

Model	Bitcoin	Ethereum	Solana	BNB
Panel A: return RMSE				
Historical mean	2.2808	3.6330	4.0622	2.5708
Gaussian linear model	2.2805	3.6353	4.0356	2.6052
MS–AR	2.2782	3.6362	4.0611	2.5961
MS–V	2.2784	3.6422	4.0371	2.5936
MS–V–BP	2.2803	3.6388	4.0456	2.5916
TVTP–V	2.2783	3.6406	4.0394	2.5942
TVTP–V–BP	2.2956	3.6417	4.0477	2.5848
Panel B: mean predictive log score				
Historical mean	−2.3423	−2.7281	−2.9643	−2.5997
Gaussian linear model	−2.3409	−2.7282	−2.9609	−2.6010
MS–AR	−2.1954	−2.6231	−2.8102	−2.3048
MS–V	−2.1955	−2.6263	−2.8094	−2.3043
MS–V–BP	−2.1971	−2.6259	−2.8097	−2.3050
TVTP–V	−2.1845	−2.6134	−2.8069	−2.2878
TVTP–V–BP	−2.1889	−2.6245	−2.8042	−2.2864

Notes: Lower RMSE and higher log scores indicate better performance. RMSE is measured in return percentage points. MS–AR includes lagged returns with constant transition probabilities; MS–V adds volume to the return equation; MS–V–BP also adds buy pressure. TVTP–V includes volume in the return and transition equations, while TVTP–V–BP includes both trading variables in both sets of equations. The Gaussian linear benchmark includes lagged returns, volume, and buy pressure. Each entry uses 638 observations.

Source: Authors' calculations from the supplied Binance dataset.

Conditional mean forecasts show limited improvement. Bitcoin's lowest RMSE comes from MS–AR, although its advantage over the historical mean is small. The historical mean performs best for Ethereum and BNB, while the Gaussian linear model obtains Solana's lowest RMSE. The full time-varying transition model has a higher RMSE than the historical mean for all four assets.

Predictive density results favour regime modelling. Every Markov-switching specification achieves a higher mean log score than the Gaussian linear benchmark for each asset. The volume-based transition model yields the highest scores for Bitcoin and Ethereum; the specification that adds buy pressure yields the highest scores for Solana and BNB. This pattern indicates that modelling changes in return distributions contributes more consistently than forecasting their conditional means.

Incremental comparisons help identify the source of those gains. Allowing volume to influence transitions improves log scores relative to the constant-transition volume model for Bitcoin q=0.0006, Solana q=0.0429 and BNB q=0.0081. After adjustment, Ethereum's improvement is not statistically distinguishable from zero (q=0.4280).

Adding buy pressure to the constant-transition volume model produces no statistically reliable positive density gain for any

asset. Adding it to the volume-based transition model also produces no reliable positive gain. Ethereum instead records a decline of 0.01107 log-score units per day, q=0.0429. The higher raw scores of the larger model for Solana and BNB therefore do not establish an incremental benefit from buy pressure.

Directional accuracy provides little further support for the full specification: its rates are 50.00% for Bitcoin, 50.78% for Ethereum, 46.87% for Solana and 52.51% for BNB. These rates alone do not demonstrate superiority over simple directional benchmarks.

The forecast comparisons are exploratory because the volume-only transition specification was added after inspecting initial results. The reported q-values adjust 16 density comparisons using approximate HAC tests, which do not fully incorporate parameter-estimation uncertainty. Within these limits, the principal finding is consistent: Markov-switching models improve predictive distributions more clearly than conditional mean or directional forecasts, while taker-buy pressure offers limited incremental forecasting value.

V. Discussion and recommendations

The findings suggest that trading activity contains more consistent information about cryptocurrency volatility conditions than about subsequent mean returns. All four assets differ between lower- and higher-volatility regimes, with higher-volatility episodes generally lasting less. Solana combines the largest estimated return shocks with the longest expected duration in the lower-volatility state. This combination shows why average volatility alone does not fully capture market risk: relatively settled periods can be interrupted by substantial disturbances.

The contribution of trading activity differs across assets. Volume helps explain volatility transitions for Ethereum and BNB, despite its weak association with their subsequent conditional mean returns. Taker-buy pressure contributes to the transition equations for Bitcoin and Solana, while Bitcoin also exhibits a stronger positive return association during higher-volatility conditions. These differences suggest that a common trading indicator need not carry the same information across cryptocurrency markets. They also caution against interpreting stronger buyer-initiated trading as a uniform signal of rising prices.

The forecasting results place these in-sample findings in perspective. Regime models generally improve predictive density scores relative to the Gaussian linear benchmark, but improvements in return magnitude and directional forecasts are inconsistent. Adding taker-buy pressure does not produce a statistically reliable positive forecasting gain across the four assets; for Ethereum, it worsens the predictive density score relative to the volume-based transition model. The distinction matters: a model can describe changing return distributions more successfully without forecasting the next day's return more accurately.

Risk managers should therefore assess these models primarily as tools for monitoring changing return distributions and volatility conditions. Volume-based transition models deserve consideration for density forecasting, while the inclusion of buy pressure should depend on demonstrated improvements using observations excluded from estimation. Model choice should remain asset-specific, retaining simpler specifications where additional variables add little value.

Investors should avoid treating the estimated buy-pressure associations as sufficient trading rules. The analysis does not establish profitability after transaction costs. Further research should examine heavier-tailed innovations, alternative regime counts and repeated model estimation through time. Evaluation across exchanges would also help determine whether the findings extend beyond the supplied Binance data. These checks would clarify how stable the estimated relationships remain as market conditions change.

References

Anastasopoulos, A., Gradojevic, N., Liu, F., Maynard, A., & Tsiakas, I. (2026). Order flow and cryptocurrency returns. *Journal of Financial Markets*, *79*, Article 101047. https://doi.org/10.1016/j.finmar.2026.101047

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Methodological)*, *57*(1), 289–300. https://doi.org/10.1111/j.2517-6161.1995.tb02031.x

Binance. (n.d.). *Binance public data* [Data documentation]. GitHub. Retrieved October 28, 2026, from https://github.com/binance/binance-public-data

Eze, C., Mba, I. C., Ameh, C., Mba, E. I., & Ugwu, P. N. (2026). Regime-aware Bitcoin price forecasting using a Gaussian hidden Markov model. *Journal of Xi'an Shiyou University, Natural Science Edition*, *22*(3), 187–195. https://www.xisdxjxsu.asia/V22I03-12.pdf

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

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

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

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

Liu, Y., & Tsyvinski, A. (2021). Risks and returns of cryptocurrency. *The Review of Financial Studies*, *34*(6), 2689–2727. https://doi.org/10.1093/rfs/hhaa113

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

Mba, E. I., Mba, I. C., Arazu, O. W., Ugah, T. E., & Urama, C. E. (2021). The bull or bear characterization of Bitcoin market conditions. *International Journal of Supply Chain Management*, *10*(3), 67–73. https://doi.org/10.59160/ijscm.v10i3.5859

Mba, I. C., Mba, E. I., Ogbuabor, J. E., & Arazu, W. O. (2018). Mean sojourn and mean return time of the buy-hoard-sell strategy of Bitcoin exchange prices. *International Journal of Economics and Financial Issues*, *8*(5), 276–282. https://econjournals.com.tr/index.php/ijefi/article/view/6777

Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, *55*(3), 703–708. https://doi.org/10.2307/1913610

Tan, C.-Y., Koh, Y.-B., Ng, K.-H., & Ng, K.-H. (2021). Dynamic volatility modelling of Bitcoin using time-

varying transition probability Markov-switching GARCH model. *The North American Journal of Economics and Finance*, *56*, Article 101377. https://doi.org/10.1016/j.najef.2021.101377