Persistence, duration and transition dynamics between stability and exchange-rate stress regimes in Bolivia's parallel dollar market (2024–2026): evidence from an autoregressive Markov-switching model

Walter Manuel Marañon Quiñones

La Paz, Bolivia · contacto@wmaranon.com

August 2026

Abstract

Since February 2023 the Bolivian economy has operated under a de facto multiple exchange rate system: an official rate pegged at BOB 6.96 per US dollar since 2011, and a parallel market that, facing foreign currency rationing, migrated towards trading the USDT stablecoin. This paper characterises the persistence, duration and transition dynamics between calm and stress phases in that market using a two-state Markov-switching autoregressive model, MS-AR(10), with regime-dependent variance, estimated by maximum likelihood on 462 daily returns of the USDT/BOB exchange rate between July 2024 and June 2026. The results identify two regimes that differ fundamentally in volatility rather than in trend: a speculative regime with 68 times higher variance and a mean duration of 4.77 trading days, and a calm regime with a mean duration of 11.25 trading days and a stationary probability of 70.2%. The resulting dynamics are cyclical rather than monotonically deteriorating: episodes of exchange rate pressure are intense but transitory and are followed by normalisation, a pattern consistent with a demand for digital dollars driven by temporary hedging against confidence shocks rather than by expectations of sustained depreciation. Out-of-sample evaluation on a 70/30 temporal split shows that the model substantially outperforms naive, constant-variance AR(10) and GARCH(1,1) specifications in predictive log-score (−1.37 versus −1.88 for the best competitor), and that its one-day Value at Risk is rejected neither by Kupiec coverage nor by Christoffersen independence tests. The one-day-ahead probability of stress discriminates well between stress and calm days in the test period (ROC-AUC of 0.86) and, with the threshold calibrated on the training set, flags 21 of the 23 stress days, albeit with a precision of 37.5%, which delimits the model's scope: probabilistic identification, early warning and risk quantification rather than exact prediction of the timing of regime breaks.

Keywords: parallel exchange market, Markov regime switching, stablecoins, Bolivia, Value at Risk, out-of-sample evaluation.

JEL Codes: F31, E58, C22, G17, O54.

1. Introduction

Bolivian exchange rate policy has remained essentially unchanged for more than a decade. In November 2011 the Central Bank of Bolivia (BCB) set the selling exchange rate at BOB 6.96 per US dollar, justifying the measure as a contribution to containing imported inflation in an economy where around 80% of imports consist of capital goods and productive inputs (Banco Central de Bolivia, 2011). That rate, held unchanged ever since, implicitly established a fixed exchange rate regime that operated during a prolonged terms-of-trade boom.

From February 2023 onwards, however, the depletion of hydrocarbon revenues, the sustained deterioration of the trade balance and the erosion of international reserves led to a widespread shortage of foreign exchange and to the rationing of foreign currency within the formal financial system. The consequence was the emergence of a de facto multiple exchange rate system, in which a parallel market with quotations substantially above the official rate came to channel a growing share of the demand for foreign exchange. A distinctive feature of this episode, which sets it apart from the parallel markets traditionally documented in the literature, is that trading migrated to the USDT stablecoin on peer-to-peer (P2P) platforms, giving rise to a digital market with quotations observable at daily frequency and with a degree of transparency unusual for an informal market (Amajaya, 2026).

This paper exploits that observability to address the following question: how can the persistence, duration and transition dynamics between phases of stability and exchange rate stress in Bolivia's parallel dollar market be characterised? The answer is not evident a priori. The classical literature on parallel markets in economies with managed regimes suggests a progressive accumulation of imbalances that eventually leads to the collapse of the official parity (Dornbusch, 1986) (Kiguel & O'Connell, 1995). An alternative reading, by contrast, would interpret the demand for foreign exchange as an episodic response to confidence shocks, with recurrent returns to normality. Both hypotheses imply radically different transition dynamics and, therefore, divergent policy recommendations.

To discriminate between them, a two-state Markov-switching autoregressive model with regime-dependent variance is estimated on 462 daily returns of the USDT/BOB exchange rate between July 2024 and June 2026. The class of Markov-switching models is particularly appropriate in this context because market states are not directly observable and because, under state intervention, the parameters governing exchange rate behaviour vary with the prevailing regime in an unannounced manner (Hamilton, 1990) (Lee & Chen, 2006).

The paper makes three contributions. First, it provides the first quantitative characterisation of the regime dynamics of the Bolivian parallel market in its digitalisation phase, documenting that the alternation between states reflects differences in volatility rather than in trend, and that the system exhibits cyclical dynamics rather than a monotonic deterioration. Second, it extends the analysis beyond the in-sample inference customary in this literature by incorporating an out-of-sample predictive validation protocol with a temporal split, comparison against benchmark models and formal backtesting of the risk measures. Third, it documents precisely the scope of the approach: probabilistic identification of the state and risk quantification are sound, and the one-day-ahead probability of stress provides a useful early warning,

although with a share of false alarms that prevents pinpointing the exact timing of a regime break, a finding that steers the agenda towards specifications with heavy tails and time-varying transition probabilities.

The remainder of the paper is organised as follows. Section 2 reviews the relevant literature. Section 3 describes the data. Section 4 presents the model specification and the evaluation protocol. Section 5 reports the results. Section 6 concludes.

2. Literature review

The literature on parallel exchange markets in developing economies established early on that the gap with respect to the official rate reflects both imbalances in fundamentals and expectations about the sustainability of the parity, and that its dynamics respond to the quantitative restrictions imposed on access to foreign exchange (Dornbusch, 1986) (Agénor, 1991) (Kiguel & O'Connell, 1995). A recurring result of this literature is that the relationship between the parallel market and its determinants is not stable: it changes with the prevailing policy regime, which motivates non-linear specifications.

The application of Markov-switching models to foreign exchange markets is motivated precisely by this parameter instability. Under managed floating schemes, intervention by the monetary authority introduces unobservable changes in the relationship between the exchange rate and its fundamentals, which can be modelled through state-dependent parameters (Lee & Chen, 2006). The evidence for emerging economies confirms this diagnosis: regime-switching models applied to Central and Eastern European markets reject the linearity hypothesis and show that the mechanism of adjustment towards long-run equilibrium depends on the prevailing exchange rate regime (Syllignakis & Kouretas, 2011); in East Asian markets, specifications with shifts in mean and variance better represent the dynamics following the 1997 crisis and improve out-of-sample predictive ability (Caporale & Spagnolo, 2004); and for the real exchange rate, two-regime models outperform the random walk and offer an alternative explanation for the persistence of unit roots (Bergman & Hansson, 2005). A cross-cutting finding is the high persistence of the states, with high probabilities of remaining in the same state that imply prolonged periods of stability or dislocation (Frömmel, MacDonald, & Menkhoff, 2005). The evidence also cautions that greater flexibility does not guarantee systematic predictive superiority over the random walk (Stillwagon & Sullivan, 2019).

A second, more recent strand applies this class of models to crypto-assets. The evidence consistently identifies low- and high-volatility states with high persistence and co-movement across assets (Figà, Focardi, & Patacca, 2019), and shows that latent-state models outperform alternatives such as VAR models and neural networks in statistical and predictive performance (Giudici & Hashish, 2020). In the estimation of risk measures, Markov-switching GARCH specifications perform better than single-regime models, particularly when heavy tails and asymmetry are incorporated (Caporale & Zekokh, 2019), while time-varying transition probabilities significantly improve the forecasting of realised volatility (Ma, Liang, Ma, & Wahab, 2020). The literature specifically on stablecoins is still limited, but it shows that their hedging properties depend on the prevailing market regime and weaken during episodes of extreme stress (Wang, Ma, & Wu, 2020).

The present paper lies at the intersection of both strands: it analyses a parallel exchange market in an economy with a fixed official exchange rate, where trading is conducted through a stablecoin. This configuration has not previously been documented using regime-switching tools. In addition, and unlike much of the applied literature cited, the analysis incorporates an explicit protocol of out-of-sample predictive validation and backtesting of the risk measures, in response to the need to assess the generalisation ability of models beyond their in-sample fit.

3. Data

The dataset comprises 463 daily observations of the parallel market selling rate (USDT/BOB) and of the official exchange rate, recorded on business days between 22 July 2024 and 1 June 2026, from which 462 returns are obtained. The parallel market series comes from public tracking of quotations for peer-to-peer transactions; the official exchange rate remains fixed at BOB 6.96 throughout the period. The analysis is carried out on the daily percentage log returns of the parallel series.

The series in levels records a mean of BOB 11.63, approximately 1.67 times the official rate; a median of BOB 11.12, below the mean, which indicates positive skewness; a minimum of BOB 9.04; and a maximum of BOB 18.23, reached in May 2025. Returns exhibit a practically zero mean of −0.0054% and a daily volatility of 1.83%, with extreme values of −11.16% and 12.26% that reveal heavy tails and episodes of abrupt variation (Table 1).

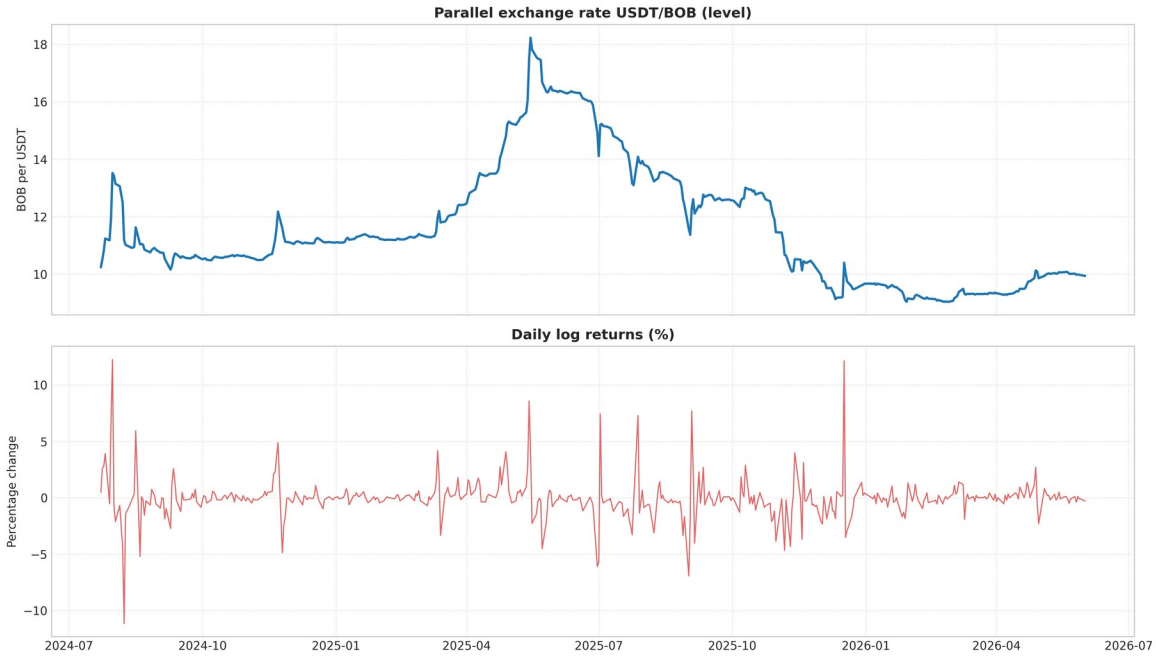
Table 1. Descriptive statistics

	Level (BOB/USDT)	Returns (%)
Observations	462	462
Mean	11.6295	−0.0054
Standard deviation	2.0747	1.8301
Minimum	9.0400	−11.1602
25th percentile	10.0700	−0.4306
Median	11.1200	0.0000
75th percentile	12.7275	0.3206
Maximum	18.2300	12.2602

Note: level statistics correspond to the 462 dates with a return, from 23 July 2024. Source: author's elaboration.

The evolution of the exchange rate reveals four stages: an initial phase of relative stability around BOB 11.00 until early 2025; an accelerated depreciation process culminating in the peak of May 2025; an abrupt correction over the following months, particularly in August 2025; and a stabilisation stage around BOB 10.00 during 2026. Returns show marked volatility clustering, with the most intense movements concentrated at the transition points between these phases, a pattern incompatible with constant-variance models (Figure 1).

Figure 1. Evolution of the parallel exchange rate and log returns



Source: author's elaboration.

The preliminary diagnostic tests confirm this diagnosis. The ADF test rejects the presence of a unit root in returns, the Ljung-Box test with ten lags detects significant autocorrelation, the ARCH-LM test detects conditional heteroskedasticity effects and the Jarque-Bera test rejects normality (Table 2). This set of results (stationarity, linear dependence, volatility clustering and heavy tails) is precisely what characterises a series generated by a mixture of two regimes with differentiated volatility.

Table 2. Pre-estimation diagnostic tests

Test	Statistic	p-value	Result
ADF (stationarity)	−9.7131	< 0.0001	Stationary series
Ljung-Box (10 lags)	32.5186	0.000328	Autocorrelation present
ARCH-LM	46.7702	< 0.0001	ARCH effects detected
Jarque-Bera	4195.5103	< 0.0001	Non-normal distribution

Source: author's elaboration.

4. Methodology

4.1 Model specification

A Markov-switching autoregressive model is estimated in which returns follow an autoregressive process of order p whose mean and variance depend on an unobservable latent state that evolves according to a first-order Markov chain with two regimes. The transition probability between states is constant, and the resulting transition matrix determines both the persistence of each regime and its expected duration, obtained as the inverse of the probability of leaving the state. Estimation is carried out by maximum

likelihood using the Hamilton filter (Hamilton, 1990), a procedure equivalent to the forward algorithm of hidden-state models.

The adopted specification allows the variance to switch between regimes but keeps the autoregressive coefficients common to both states. This restriction reflects a parsimony criterion: the descriptive evidence suggests that the difference between states lies in the magnitude of volatility rather than in the structure of temporal dependence, and allowing regime-specific autoregressive coefficients would multiply the number of parameters without empirical justification. Estimation uses a quasi-maximum likelihood robust covariance matrix. To mitigate the risk of local optima, which is characteristic of the multimodal likelihood of this class of models, maximisation starts from the best of fifty sets of random initial values, each refined with five iterations of the EM algorithm. The likelihood assigns the variance of each observation to the regime prevailing on that same date. The statsmodels implementation (MarkovAutoregression, verified in versions 0.14.5 and 0.15.0), when the variance switches with the regime and the autoregressive order is greater than one, instead assigns it to the regime $p - 1$ observations earlier, so that function was corrected before estimation.

4.2 Selection of the autoregressive order

The autoregressive order was determined using information criteria over orders 1 to 12. Both criteria tend to decline with the order, although not monotonically (the BIC rises slightly between orders 1 and 3, and both rebound at order 8), and they reach their lowest value at order 12, so they show no interior minimum over the range evaluated. In the absence of an identified optimum, the selection was settled by a computational feasibility criterion: the cost of the Hamilton filter grows exponentially with the order, and the estimation of order 13 had to be stopped after two hours without converging. Order 10 was adopted as the final specification, given that orders 11 and 12 could only be estimated with a reduced number of restarts and therefore with less assurance that the optimum reached was robust. It is worth noting that the partial autocorrelation function suggests a low-order structure, but it assesses the individual contribution of each lag rather than the joint explanatory power of the autoregressive block, which is what the information criteria capture.

4.3 Out-of-sample evaluation protocol

The evaluation of generalisation ability uses a temporal split of the series that preserves chronological order. The training set runs from July 2024 to 31 October 2025 (322 observations, 69.7%) and the test set from November 2025 to June 2026 (140 observations, 30.3%). The cut-off point was set according to an estimability criterion: the test set must contain observations of both classes for the classification metrics to be computable, a condition that this cut-off meets by including both the December 2025 stress episode and the subsequent stabilisation phase.

The model is re-estimated exclusively on the training set and its parameters are held fixed when the filter is applied to the test period, so that the regime probability assigned to each day uses only information available up to the previous day. In this subsample the random search converges to a local optimum (log-likelihood of −392.27), so the higher maximum (−391.46) reached by starting from the full-sample parameters is adopted; those values only set the starting point of an optimisation that uses exclusively the

training-set likelihood. This distinction is essential: smoothed probabilities, customary in the applied literature, incorporate information from the whole sample and are valid for historical classification but not for evaluating predictive ability.

Since the regimes are latent states without an observable label, the classification evaluation requires an explicitly stated proxy label: a day is considered a stress day if the five-day realised volatility exceeds the 75th percentile of that same measure in the training set, a threshold that approximates the proportion of speculative days identified in-sample. On that basis, the confusion matrix, precision, sensitivity, F1-Score, area under the ROC curve and Brier score are reported, together with a sensitivity analysis with respect to the labelling threshold. Binary classification uses a decision threshold calibrated on the training set as the one that maximises the F1-Score, without using information from the test period. Predictive performance is also compared against three benchmark models (naive with constant mean and variance, constant-variance AR(10) and GARCH(1,1)), using as the main metric the mean predictive log-score, a proper scoring rule that evaluates the entire forecast density (Gneiting & Raftery, 2007), complemented by the root mean squared error (RMSE) for the mean and the QLIKE loss for the variance (Patton, 2011). Finally, the one-day Value at Risk, derived from the predictive density as a mixture of the components of each regime, is subjected to the Kupiec (1995) unconditional coverage test and the Christoffersen (1998) independence test.

5. Results

5.1 Parameters by regime

The estimation identifies two regimes clearly differentiated by their volatility. The speculative regime has a conditional variance of 9.3229, approximately 68 times higher than that of the calm regime (0.1371), with both estimates significant at the 1% level. As regards central tendency, neither regime mean is statistically different from zero: the speculative regime mean is 0.1805% ($p = 0.591$) and the calm regime mean is −0.0357% ($p = 0.320$), confirming the absence of a systematic directional trend in both stress and stability periods (Table 3).

Table 3. Estimated parameters by regime

Regime	Parameter	Coefficient	Standard error	z	p-value
Speculative	μ	0.1805	0.336	0.537	0.591
Speculative	σ^2	9.3229	2.089	4.462	< 0.001
Calm	μ	−0.0357	0.036	−0.995	0.320
Calm	σ^2	0.1371	0.021	6.413	< 0.001

Note: μ is the mean return in each regime. Source: author's elaboration.

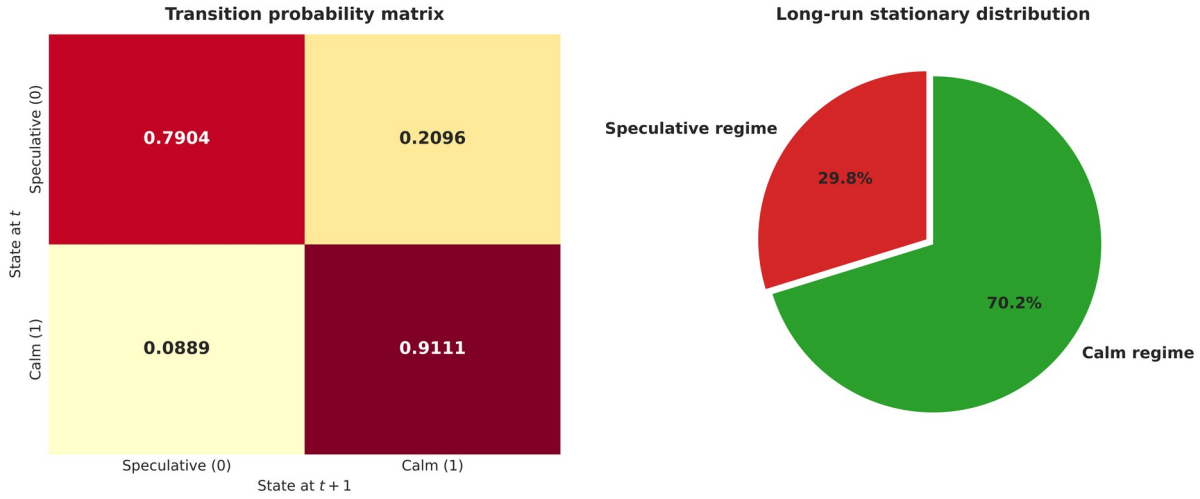
This result has a substantive implication: the differentiation between states is driven by the intensity of risk and not by differing depreciation expectations. The speculative regime mean, positive but not significant, reflects the fact that stress episodes include both abrupt rises and equally rapid corrections, so that its conditional mean remains contained. As for the autoregressive structure, five of the ten lags are significant at the 5% level (L1, L3, L5, L9 and L10), the first being the largest in magnitude (0.1148); the

significance of the more distant lags indicates that the incorporation of information into the market extends up to ten trading days, approximately two weeks.

5.2 Transition dynamics and persistence

The estimated transition matrix shows a marked asymmetry in the persistence of the states. The probability of remaining in the calm regime reaches 91.11%, compared with 79.04% for the speculative regime, which implies expected durations of 11.25 and 4.77 trading days, respectively. The stationary distribution assigns a probability of 70.2% to the calm regime and of 29.8% to the speculative regime (Figure 2).

Figure 2. Transition matrix and long-run stationary distribution

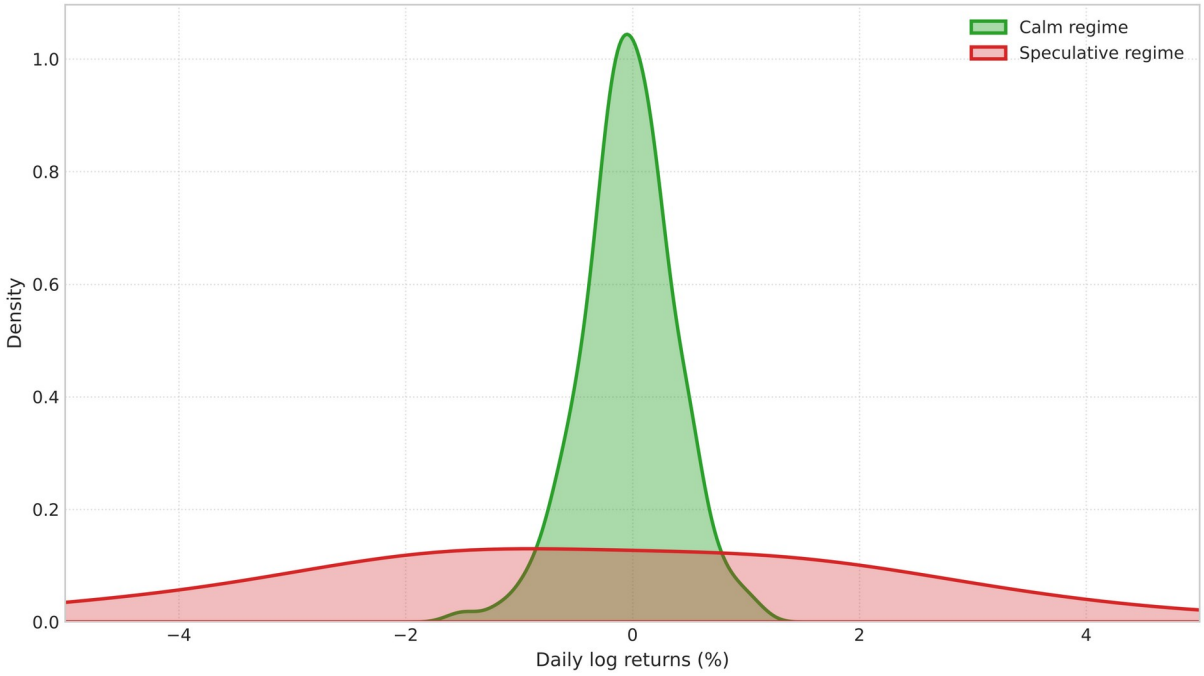


Source: author's elaboration.

This is the central result of the paper, and it contradicts the reading of a monotonic deterioration. Episodes of exchange rate stress are intense but transitory, with a mean duration of about one calendar week, and the market recurrently returns to a calm state that dominates the long-run dynamics. The evidence is consistent with a demand for digital dollars driven mainly by temporary hedging needs in the face of increases in uncertainty, rather than by an entrenched expectation of permanent depreciation (Bocola & Lorenzoni, 2020) (Basso, Calvo, & Jurgilas, 2011). The high persistence of both states, for its part, is in line with the international evidence for foreign exchange markets subject to state intervention (Frömmel, MacDonald, & Menkhoff, 2005).

The distribution of returns conditional on each regime confirms that the differentiation is one of dispersion rather than location: the density of the calm regime is tightly concentrated around zero, whereas that of the speculative regime has substantially wider tails (Figure 3).

Figure 3. Conditional density of returns by inferred regime

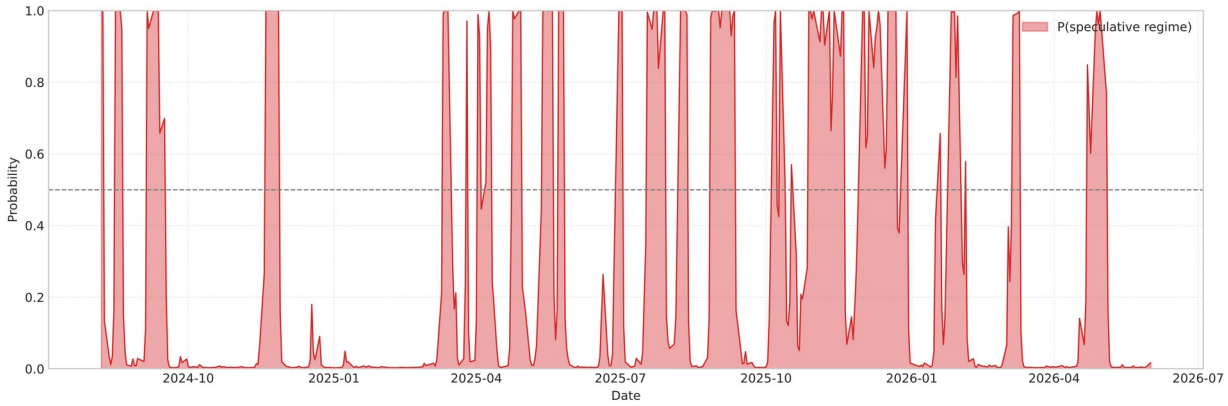


Source: author's elaboration.

5.3 Historical classification and validation

The smoothed probability of the speculative regime does not rise continuously but in brief, recurrent episodes: 26 episodes with a probability above 0.5 are identified, most of them lasting a few days. They are concentrated in the accelerated depreciation phase between March and May 2025, which culminated in the peak of BOB 18.23, in the correction from July to September 2025 and, most intensely, between late October and December 2025, when the speculative regime covers 61% of the observations in November and 81% of those in December, while the parallel exchange rate fell from BOB 11.46 to BOB 9.67; outside those periods, isolated episodes appear in the second half of 2024 and in 2026. The correspondence between the model's classification and the observed chronology of the market is evidence that stress episodes are identified using only the information contained in the dynamics of returns (Figure 4).

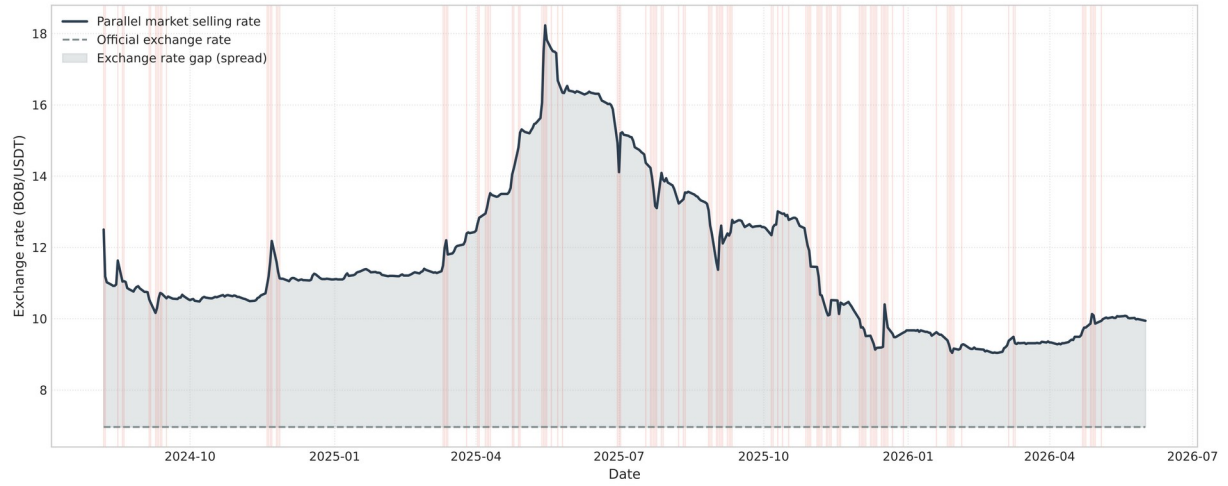
Figure 4. Smoothed probability of the speculative regime



Source: author's elaboration.

The classification of the 452 effective observations assigns 329 (72.8%) to the calm regime and 123 (27.2%) to the speculative regime. The average exchange rate gap with respect to the official rate is BOB 4.62 during calm periods and BOB 4.77 during stress periods. The difference between the two averages, BOB 0.15, is small relative to the total range of the gap over the period, which reinforces the previous diagnosis: the regimes are distinguished by the intensity of fluctuations rather than by the level of the gap (Figure 5).

Figure 5. Exchange rate gap and regime classification



Source: author's elaboration.

Post-estimation diagnostics validate the specification in terms of temporal dependence: the Ljung-Box test on the residuals does not reject the absence of autocorrelation ($p = 0.3207$), in contrast to the rejection obtained on the original returns, confirming that ten lags are sufficient to capture the linear dependence. The ARCH-LM test, by contrast, continues to reject homoskedasticity ($p = 0.0018$), indicating that the regime-dependent variance absorbed a considerable part, but not all, of the conditional heteroskedasticity. This is the main statistical limitation of the model and motivates the extension towards Markov-switching GARCH specifications.

5.4 Out-of-sample predictive evaluation

Re-estimating the model on the training set faithfully reproduces the full-sample parameters: the variance ratio between regimes is 83, compared with 68 in the full sample; the mean durations are 11.59 and 4.59 trading days, compared with 11.25 and 4.77; and the significance pattern of the autoregressive lags is fully preserved. This stability in the face of a reduced sample constitutes evidence of the structural robustness of the specification.

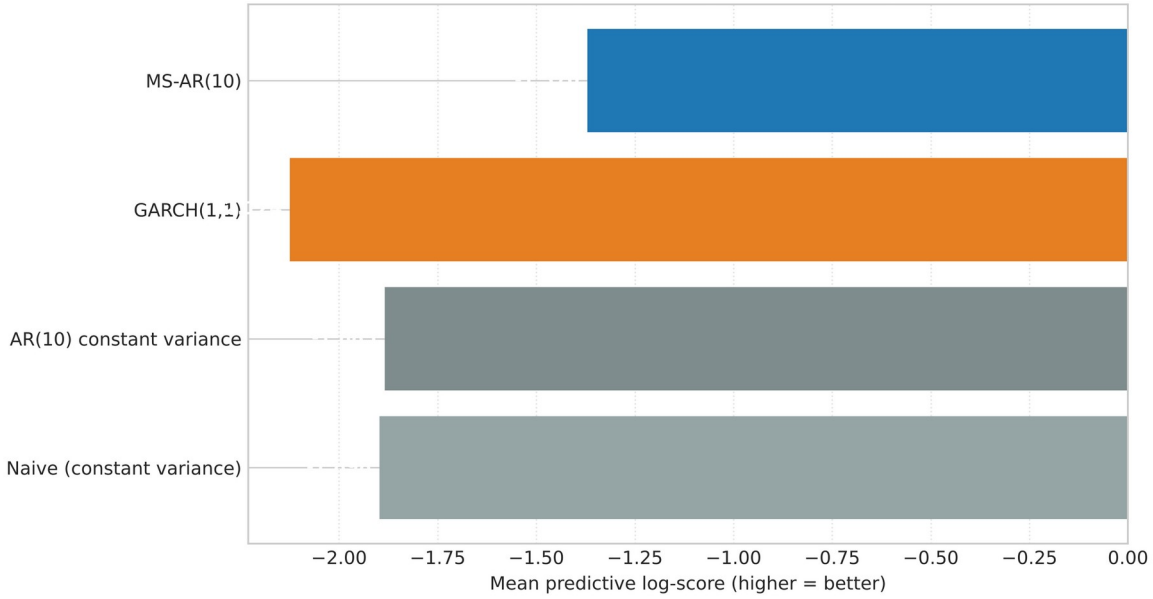
The comparison of out-of-sample predictive performance shows that the regime-switching model dominates the three benchmark models on the main metric by a wide margin, with a log-score of −1.3699 compared with −1.8846 for the best competitor. The GARCH(1,1) model records the worst performance, penalised by the overreaction of its forecast variance after the December 2025 episode. In terms of the RMSE for the mean, the four models are practically indistinguishable, an expected result for a series with a mean return close to zero, which confirms that the informational content of the model lies in the variance and not in the direction. The QLIKE loss confirms this: the regime-switching model also obtains the best score (1.5973 compared with 1.8748 for the constant-variance AR(10)), so its forecast variance tracks observed volatility more closely, even in a test period dominated by the calm phase (Table 4, Figure 6).

Table 4. Out-of-sample comparison with benchmark models

Model	Predictive log-score	RMSE (mean)	QLIKE (variance)
Naive (constant variance)	−1.8977	1.5408	1.9552
AR(10) constant variance	−1.8846	1.5887	1.8748
GARCH(1,1)	−2.1243	1.5423	2.4190
MS-AR(10)	−1.3699	1.5619	1.5973

Note: a higher (less negative) log-score indicates better performance; lower RMSE and QLIKE indicate better performance.
Source: author's elaboration.

Figure 6. Out-of-sample predictive performance by model



Source: author's elaboration.

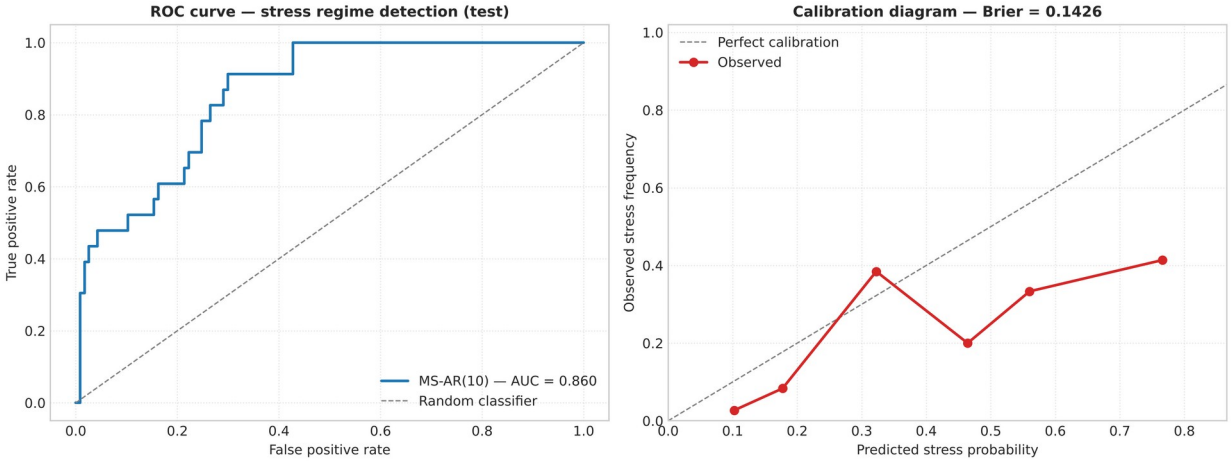
The one-step-ahead probability of stress also shows considerable discriminating power. With the threshold calibrated on the training set (0.20), the model identifies 21 of the 23 stress days in the test period, with a sensitivity of 91.3%, at the cost of 35 false alarms, which puts precision at 37.5% and the F1-Score at 0.532. The area under the ROC curve reaches 0.860, well above that of a random classifier (Table 5, Figure 7). The sensitivity analysis with respect to the labelling threshold, evaluated at the 70th, 75th and 80th percentiles of the training set, confirms the stability of the result, with AUC values between 0.860 and 0.890.

Table 5. Out-of-sample classification: anticipated versus contemporaneous

Approach	Precision	Sensitivity	F1-Score	ROC-AUC	Brier
Anticipated one step ahead	0.3750	0.9130	0.5316	0.8599	0.1426
Contemporaneous (filtered)	0.3281	0.9130	0.4828	0.8606	0.1844

Note: decision thresholds calibrated on the training set: 0.20 (anticipated) and 0.07 (contemporaneous). Source: author's elaboration.

Figure 7. ROC curve and calibration diagram



Source: author's elaboration.

A complementary exercise compares anticipation with contemporaneous classification, which uses the filtered probabilities with same-day information and a threshold calibrated in the same way (0.07): sensitivity remains at 91.3%, precision falls to 32.8% and the F1-Score to 0.483, with an AUC of 0.861 and a Brier score of 0.184. One-day-ahead anticipation loses no discriminating power relative to contemporaneous classification, so the predicted probability works as an early warning. Its limitation is the share of false alarms: in the test period, less volatile than the training period, isolated daily moves of around 1% or 2% are unlikely under the calm regime, whose standard deviation is only 0.36%, and they raise the next day's probability of stress even when they do not push the five-day realised volatility above the label threshold. The return on the day before a false alarm has a median of 1.16% in absolute value, compared with 0.22% on days correctly classified as calm.

5.5 Value at Risk backtesting

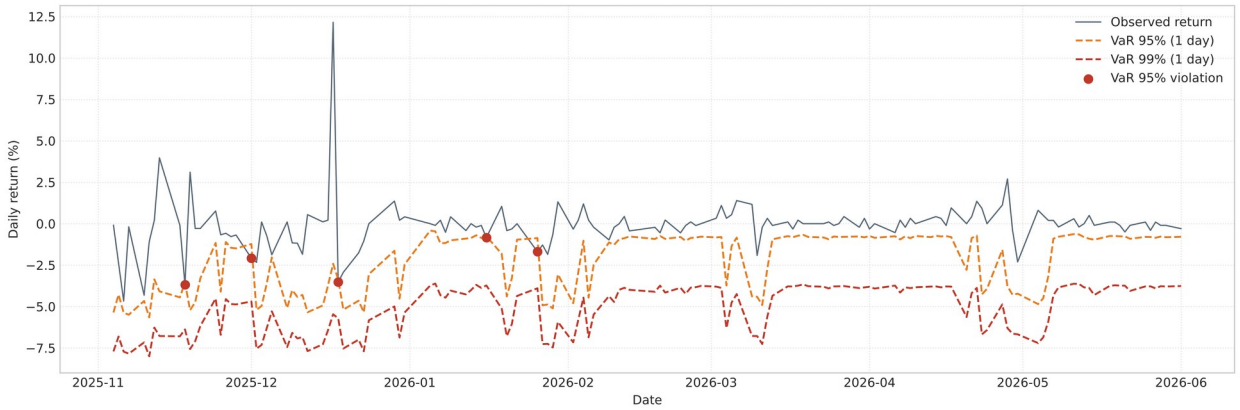
The one-day Value at Risk, calculated for each date of the test period from the model's predictive density, formally passes the standard adequacy criterion. The Kupiec test does not reject correct coverage at either of the two levels: at 95%, 5 violations are recorded against 7.0 expected (p = 0.4147), and at 99%, none against 1.4 expected (p = 0.0934). Nor does the Christoffersen test at 95% reject the independence of the violations (p = 0.5413), which rules out their clustering over time; at 99% it cannot be computed owing to the absence of violations. It should be noted that, with 140 observations, the power of the backtesting at 99% is limited (Table 6, Figure 8).

Table 6. One-day Value at Risk backtesting

Measure	Violations	Expected	Kupiec p-value	Christoffersen p-value
VaR 95%	5/140	7.0	0.4147	0.5413
VaR 99%	0/140	1.4	0.0934	n/a

Source: author's elaboration.

Figure 8. Observed returns versus estimated Value at Risk



Source: author's elaboration.

The lower-than-expected number of violations reveals a slightly conservative bias: the forecast quantiles moderately overestimate risk during calm phases, as a consequence of the positive weight that the transition dynamics permanently assign to the stress regime. From a risk management perspective, this bias is prudent, insofar as it leads to sufficient provisions.

6. Conclusions

The parallel dollar market in Bolivia, in its phase of digitalisation through stablecoins, exhibits dynamics governed by the alternation between two regimes that differ fundamentally in volatility rather than in trend. Stress episodes are intense but transitory, with a mean duration of 4.77 trading days, whereas calm phases last 11.25 trading days on average and account for 70.2% of the long-run stationary distribution.

This pattern differs from what the classical literature on currency crises in economies with managed regimes would anticipate. Instead of a continuous deterioration progressively leading to the collapse of the official parity, the results show cyclical dynamics in which episodes of pressure are followed by processes of normalisation. The interpretation consistent with this evidence is that the demand for digital dollars responds mainly to temporary hedging needs in the face of increases in uncertainty, and that the USDT/BOB market performs a function that goes beyond mere foreign exchange scarcity, acting as a hedging mechanism against confidence shocks. A direct practical implication is that an increase in volatility should not automatically be interpreted as the onset of a permanent depreciation, a distinction that linear models tend not to capture.

The out-of-sample evaluation supports the specification as a description of the data-generating process: the model substantially outperforms the constant-variance and Gaussian conditional volatility alternatives in predictive log-score and QLIKE loss, its Value at Risk is not rejected by either the coverage or the independence tests, and its parameters remain stable when the sample is reduced. The one-day-ahead probability of stress discriminates well between stress and calm days (ROC-AUC of 0.86) and flags almost all of those in the test period, with the same discriminating power as contemporaneous classification, albeit with a high share of false alarms. The scope of the model is thus precisely delimited: probabilistic identification of the state, characterisation of the transition dynamics, early warning and risk quantification, rather than exact prediction of the timing of crises.

Three extensions follow from these results. First, the persistence of ARCH effects in the residuals motivates Markov-switching GARCH specifications with heavy-tailed distributions, which would foreseeably improve the estimation of risk measures. Second, the evaluation of a third regime, supported by information criteria, could capture transition states or crisis episodes of greater intensity. Third, and most relevant in policy terms, the share of false alarms indicates that the history of returns is not enough to pinpoint the timing of regime breaks: incorporating macroeconomic variables (net international reserves, fiscal variables, exchange rate policy announcements) as determinants of time-varying transition probabilities is the natural route for linking regime changes to the structural factors that give rise to them.

References

Agénor, P. (1991). A Monetary Model of the Parallel Market for Foreign Exchange. Journal of Economic Studies, 18.

Amajaya, R. (2026). Dinámica del mercado digital de divisas USDT/BOB y su impacto en la economía boliviana. Boletín Económico Virtual.

Banco Central de Bolivia. (2011). Nota de Prensa NP 42/2011. La Paz.

Basso, H., Calvo, O., & Jurgilas, M. (2011). Financial dollarization: The role of foreign-owned banks and interest rates. Journal of Banking & Finance, 35, 794–806.

Bergman, M., & Hansson, J. (2005). Real exchange rates and switching regimes. Journal of International Money and Finance, 24, 121–138.

Bocola, L., & Lorenzoni, G. (2020). Financial Crises, Dollarization, and Lending of Last Resort in Open Economies. American Economic Review, 110, 2524–2557.

Caporale, G., & Spagnolo, N. (2004). Modelling East Asian exchange rates: a Markov-switching approach. Applied Financial Economics, 14, 233–242.

Caporale, G., & Zekokh, T. (2019). Modelling volatility of cryptocurrencies using Markov-Switching GARCH models. Research in International Business and Finance, 48, 143–155.

Christoffersen, P. (1998). Evaluating Interval Forecasts. International Economic Review, 39, 841–862.

Dornbusch, R. (1986). Special Exchange Rates for Capital Account Transactions. The World Bank Economic Review, 1, 3–33.

Figà, G., Focardi, S., & Patacca, M. (2019). Regime Switching Analysis of Cryptocurrencies. SSRN Electronic Journal.

Frömmel, M., MacDonald, R., & Menkhoff, L. (2005). Markov switching regimes in a monetary exchange rate model. Economic Modelling, 22, 485–502.

Giudici, P., & Hashish, I. (2020). A hidden Markov model to detect regime changes in cryptoasset markets. Quality and Reliability Engineering International, 36, 2057–2065.

Gneiting, T., & Raftery, A. (2007). Strictly Proper Scoring Rules, Prediction, and Estimation. Journal of the American Statistical Association, 102, 359–378.

Hamilton, J. (1990). Analysis of Time Series Subject to Changes in Regime. Journal of Econometrics, 45, 39–70.

Kiguel, M., & O'Connell, S. (1995). Parallel Exchange Rates in Developing Countries. The World Bank Research Observer, 10, 21–52.

Kupiec, P. (1995). Techniques for Verifying the Accuracy of Risk Measurement Models. The Journal of Derivatives, 3, 73–84.

Lee, H.-Y., & Chen, S.-L. (2006). Why use Markov-switching models in exchange rate prediction? Economic Modelling, 23, 662–668.

Ma, F., Liang, C., Ma, Y., & Wahab, M. (2020). Cryptocurrency volatility forecasting: A Markov regime-switching MIDAS approach. Journal of Forecasting, 39, 1277–1290.

Patton, A. (2011). Volatility forecast comparison using imperfect volatility proxies. Journal of Econometrics, 160, 246–256.

Stillwagon, J., & Sullivan, P. (2019). Markov switching in exchange rate models: will more regimes help? Empirical Economics, 59, 413–436.

Syllignakis, M., & Kouretas, G. (2011). Markov-switching regimes and the monetary model of exchange rate determination: Evidence from the Central and Eastern European markets. Journal of International Financial Markets, Institutions and Money, 21, 707–723.

Wang, G.-J., Ma, X.-y., & Wu, H.-y. (2020). Are stablecoins truly diversifiers, hedges, or safe havens against traditional cryptocurrencies as their name suggests? Research in International Business and Finance, 54.