Finance Research Letters 86 (2025) 108941

Contents lists available at ScienceDirect

Finance Research Letters

journal homepage: www.elsevier.com/locate/frl

When timing matters: Regime-dependent delays in exchange rate fundamentals

Lucas M. Oliveira *, Airlane P. Alencar

Institute of Mathematics and Statistics – University of São Paulo, Rua do Matão 1010, São Paulo, SP 05508-090, Brazil

ARTICLE INFO

JEL classification:
F31
C32
C53

Keywords:
Exchange rates
Markov switching
Regime-dependent lags
Prediction exercise

ABSTRACT

The literature suggests that Markov-switching (MS) models can address the exchange rate disconnect puzzle, yet they assume a fixed transmission lag for fundamentals. We test whether this timing restriction is the key limitation and introduce a regime-switching regression with AR errors and state-dependent lags. Regimes are defined by volatility and reaction times vary across low- and high-volatility states. Each regime selects its own lag length from an economically plausible grid and selected lags range from 0 to 12 months. Allowing timing to vary across regimes improves medium- and long-horizon forecasts, and the dominant transmission channels remain currency specific.

1. Introduction

Exchange rate (ER) forecasting remains one of the most challenging problems in international finance. The classic disconnect between fundamentals and exchange rates (Meese and Rogoff, 1983) persists despite decades of research. Nonlinear relationships help explain this gap, and Markov-switching (MS) models (Hamilton, 1989) have become a natural tool because they improve the fit for many currencies (Engel, 1994), sometimes even forecasts (Dewachter, 2001), and show that the real interest differential (RID) relation (Frankel, 1979) changes across market states (Frömmel et al., 2005). Recent work develops a fully Bayesian MS-VECM with time-varying transition probabilities for BRICS–USD exchange rates, enabling clear regime identification and forecasting gains (Kumar et al., 2024).

Most MS applications fix fundamentals' lags, limiting the capture of temporal heterogeneity. Recent MS–VAR work allows state-dependent autoregressive orders (Li and Kwok, 2021) but not state-dependent exogenous lags, and many ER studies still keep fixed lags in MS settings (Stillwagon and Sullivan, 2020). We argue that transmission timing is itself state dependent, and recent external evidence supports this across complementary fronts.

Bayesian dynamic learning across 512 VARs yields sizable FX portfolio gains despite modest point-forecast gains over a random walk (Beckmann et al., 2020). Fundamentals-based expectations matter mainly at longer horizons (Beckmann and Czudaj, 2025), and policy uncertainty comoves with volatility regimes—in a GARCH–MIDAS, the Sino–US EPU ratio improves long-run CNY/USD volatility forecasts vs. standard GARCH (Zhou et al., 2020). DL results are mixed: LSTM improves statistical accuracy over vanilla RNN yet offers weak economic performance (Dautel et al., 2020); during COVID-19, stacked Bi-LSTMs with a Bagging–Ridge meta-learner beat classical ML on errors across 21 USD pairs, but effects are heterogeneous (Abedin et al., 2025), and loss-metric gains rarely translate into robust economic gains, favoring parsimonious, interpretable timing models.

We address these restrictions with a regime-dependent regression in which each fundamental's lag is jointly selected by BIC. Unlike the TVP-VAR of Primiceri (2005) and the smooth-transition VAR of Lundbergh et al. (2003), where lag orders are fixed and

* Corresponding author.
E-mail addresses: lucmoleconomista@gmail.com, lucmol@usp.br (L.M. Oliveira), lanealencar@usp.br (A.P. Alencar).

https://doi.org/10.1016/j.frl.2025.108941
Received 26 September 2025; Received in revised form 31 October 2025; Accepted 12 November 2025
Available online 12 November 2025

time variation lies in smoothly evolving coefficients with no discrete regimes, our approach places time variation in the lags of the fundamentals rather than in the coefficients. This yields explicit state-level lag vectors and regime durations, and it makes the priced information set visible in each regime.

Transmission timing is regime specific and dominant channels are currency specific, and forecasts beat a random walk and a fixed-lag MS across some horizons, with gains validated by Diebold and Mariano (2002) and Clark and West (2007) tests. Excess returns are positive and economically meaningful by the Sharpe ratio, the Lo-adjusted Sharpe ratio (Lo, 2002), and a high probabilistic Sharpe ratio (Bailey and Lopez de Prado, 2012), confirming that state-dependent timing matters for ER dynamics.

2. Research methods

2.1. Proposed model

We ask whether the timing of fundamentals depends on the state. Do the relevant lags change when the regime changes? We allow each fundamental to have a regime-specific delay and to capture the time-varying nature of market reaction speeds, we specify an MS model where the transmission lag of each macroeconomic fundamental is state-dependent. The exchange rate return, Δe_t, is driven by a set of monetary fundamentals, including differentials in industrial production (Δy_t), short-term interest rates ($\Delta i_{s,t}$), long-term interest rates ($\Delta i_{l,t}$), and money supply (Δm_t). The model is specified as:

$$\Delta e_t = \alpha_{s_t} + \beta^{(y)}_{s_t}\Delta y_{t-\ell^{(y)}_{s_t}} + \beta^{(is)}_{s_t}\Delta i_{s,t-\ell^{(is)}_{s_t}} + \beta^{(il)}_{s_t}\Delta i_{l,t-\ell^{(il)}_{s_t}} + \beta^{(m)}_{s_t}\Delta m_{t-\ell^{(m)}_{s_t}} + u_t,$$
$$u_t = \phi_{s_t}u_{t-1} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma^2_{s_t}),$$

where the unobserved state variable, $s_t \in \{1, 2\}$, follows a first-order Markov chain and governs all parameters, including the lag $\ell^{(x)}_{s_t}$ for each fundamental. Errors follow an AR(1) process and $|\phi_k| < 1$ ensures stationarity in each regime.

For each regime $j \in \{1, 2\}$, we allow for regime-specific transmission lags by selecting the optimal pair (ℓ_1, ℓ_2) over a grid of economically plausible lag lengths and the selection was based on BIC. The model is estimated by maximum conditional likelihood using the expectation–maximization (EM) algorithm with the Hamilton and Kim filters (Kim and Nelson, 1999). The estimation procedure is detailed in Appendix A.

To address identification and economic plausibility, we fix the labels by ordering variances at every EM step. Regime 1 is the low variance state and regime 2 is the high variance state, with $\sigma_1 \leq \sigma_2$. Both regimes can choose long or short lags, so the turbulent state can be fast or slow depending on the data.

Delayed reaction under turbulence is economically plausible: noise trading and multiple equilibria can amplify volatility (Jeanne and Rose, 2002), while balance-sheet constraints on intermediaries slow price discovery when risk bearing capacity is strained (Gabaix and Maggiori, 2015). Liquidity premia co-move with exchange rates across G10, reinforcing this mechanism (Engel and Wu, 2023). Our specification does not impose contemporaneous transmission in either regime. If a regime chooses zero lags for all fundamentals the model collapses to the classic contemporaneous monetary specification in Dornbusch (1976), Frankel (1979).

2.2. Data

We analyze GBP/USD, CAD/USD, and JPY/USD, the only USD pairs from G7 economies with their own currencies. We choose them for market depth, data quality, and comparability. All series are monthly. The sample periods vary due to data availability: September 1987 to November 2023 for GBP/USD (435 observations), January 1972 to October 2023 for CAD/USD (622 observations), and January 1990 to June 2021 for JPY/USD (378 observations). For our forecast exercise, the period from January 2019 onward is used as the out-of-sample window for all series. Most are sourced from the Federal Reserve Economic Data (FRED), except the U.S. broad money (M3) series, which comes from the OECD Data Explorer.

Following Frömmel et al. (2005), we construct four monthly differentials with respect to the United States. For industrial production and money (M3), we used the annual log-difference:

$$\Delta_{12}\ln\left(\frac{x_t}{x^*_t}\right) = \ln\left(\frac{x_t}{x^*_t}\right) - \ln\left(\frac{x_{t-12}}{x^*_{t-12}}\right), \tag{1}$$

where x_t and x^*_t denote, respectively, the foreign and U.S. series.

For the 3-month (short) and 10-year (long) rates, we defined the interest-rate regressors as the 12-month percentage-change differentials:

$$\Delta i_{s,t} \equiv \frac{i_{s,t} - i_{s,t-12}}{i_{s,t-12}} - \frac{i^*_{s,t} - i^*_{s,t-12}}{i^*_{s,t-12}}, \qquad \Delta i_{l,t} \equiv \frac{i_{l,t} - i_{l,t-12}}{i_{l,t-12}} - \frac{i^*_{l,t} - i^*_{l,t-12}}{i^*_{l,t-12}},$$

where the symbol * denotes the U.S.

Most differential variables[1] rejected a unit root at the 5% level in Augmented Dickey–Fuller tests (ADF) (Said and Dickey, 1984). Full results are in Table S1 in Supplementary Material (SM).

[1] Two exceptions are GBP money supply and GBP long-term interest but while ADF does not reject a unit root for these series, the Phillips–Perron test rejects at the one percent level, indicating stationarity once serial correlation and heteroskedasticity are accommodated.

Table 1
Estimation results for RID and Two-State MS Fixed Lags with AR(1) errors.

	GBP/USD		CAD/USD		JPY/USD	
	State 1	*State 2*	*State 1*	*State 2*	*State 1*	*State 2*
RID						
Constant	0.0037		0.0093		−0.0241	
Industrial production	0.5226***		0.5698***		−0.0926	
Short-term rate	−0.0046**		−0.0029*		0.0056***	
M3	0.1588*		−0.2070***		−0.3856	
Long-term rate	−0.1539***		−0.0483*		−0.0219***	
σ	0.0865		0.0626		0.1040	
MS-FL-AR(1)						
Constant	−0.0054	0.0149	0.0274	0.1072	0.1143	−0.0314
Industrial production	−0.0137	−0.0708	−0.0590	0.2985**	0.0473	−0.2903***
Short-term rate	0.0009	−0.0007	0.0160*	0.0003	0.0006	−0.0050
M3	0.0956	0.5705**	−0.0240	0.6497***	2.3807***	−0.5607
Long-term rate	−0.1303***	−0.0233	0.0527*	−0.1215***	0.0007	−0.0138
ϕ	0.9391	0.9343	0.9399	0.9890	0.9651	0.9348
σ	0.0224	0.0412	0.0141	0.0245	0.0283	0.0361

Asterisks refer to significance at *10%, **5%, ***1%.

2.3. Out-of-sample prediction

We performed an out-of-sample (OOS) forecasting exercise[2] that were generated with our two-state MS-SL-AR(1) model with both regimes using their BIC-selected lag vectors. All forecasts use ex post revised fundamentals in place of real-time data, which may overstate implementable predictability. The AR term enters as a function of the lagged residual, $\hat{\phi}_{s_t}(\Delta e_{t-1} - \hat{\mu}_{s_{t-1}}(t-1))$, which is why the lagged dependent variable, Δe_{t-1}, is not included as a separate regressor. For each OOS origin t, the model is re-estimated by EM on a rolling window ending at t, yielding filtered regime probabilities $\Pr(S_t = i \mid \mathcal{I}_t)$ and estimated transition matrix $\hat{P} = \{\hat{p}_{ij}\}$. The strict one-step forecast is

$$\widehat{\Delta e}_{t+1|\mathcal{I}_t} = \sum_{i=1}^{2}\sum_{j=1}^{2} \Pr(S_t = i \mid \mathcal{I}_t)\,\hat{p}_{ij}\Big\{\hat{\mu}_j(t+1) + \hat{\phi}_j \underbrace{\big[\Delta e_t - \hat{\mu}_i(t)\big]}_{u_t^{(i)}}\Big\},$$

$$\hat{\mu}_j(\tau) = \hat{\alpha}_j + \hat{\beta}_j^{(y)}\,\Delta y_{\tau-\ell_j^{(y)}} + \hat{\beta}_j^{(is)}\,\Delta i_{s,\tau-\ell_j^{(is)}} + \hat{\beta}_j^{(il)}\,\Delta i_{l,\tau-\ell_j^{(il)}} + \hat{\beta}_j^{(m)}\,\Delta m_{\tau-\ell_j^{(m)}}.$$

and forecast errors are assessed *ex post* using the Mean Error (ME), the Mean Absolute Error (MAE), the Root Mean Squared Error (RMSE), and the Mean Squared Prediction Error (MSPE).

3. Results

We fit a Real Interest Differential (RID) model and a two-state fixed-lag Markov-switching (MS-FL-AR(1)) with fixed contemporaneous transmission, then benchmark them against our switching lag model (MS-SL-AR(1)). The MS-FL-AR(1) redistributes average effects across states but still blurs timing, while the switching lag model separates channels. For GBP/USD the calm state is led by the long rate and turbulence is driven by money, with industrial production and the short rate largely secondary. Tables 1 and 2 report the corresponding estimates.

For CAD/USD the long rate flips sign across regimes and the switching lag model magnifies this asymmetry, with the calm state anchored by the long rate and turbulence loading more on real activity and money. For JPY/USD the linear model is weak, switching concentrates the signal in money, and the switching lag model delivers a clear calm state multifactor structure with money and rates jointly informative while the turbulent state collapses to a narrow money channel. AR persistence remains high for GBP and CAD and mean reversion is faster for JPY in calm periods.

Fig. 1 reports BIC-selected lags by regime. In GBP/USD the calm state sets industrial production to three months and the short rate to one while money and the long rate are contemporaneous, and in turbulence the delay shifts to the long rate at three with the others near zero. In CAD/USD the calm state is near contemporaneous with only one month for industrial production, and in turbulence money reaches six months with industrial production at three and the short rate at one. In JPY/USD the calm state places six months on money with one month on industrial production and the short rate and zero on the long rate, and in turbulence money falls to zero while industrial production and the long rate move to three and the short rate stays at one. Across pairs short rates adjust within zero to one month and money carries the longest state-dependent delays. Table 3 confirms the stability of these regime-specific lag choices across the Top-10 BIC models per currency and regime.

[2] We evaluate predictions, not operational forecasts. At horizon k, when the model sets $\ell = 0$, we use the observed fundamentals x_{t+k} to calculate the expected value $\hat{\mu}_j(t+k)$. The reported results measure conditional predictive content given future fundamentals, not real-time forecasting performance. Comparisons with the random walk (RW) follow the Meese–Rogoff tradition (Meese and Rogoff, 1983).

Table 2
Estimated coefficients for the switching lag MS model.

	GBP/USD	*CAD/USD*	*JPY/USD*
State 1			
Constant	−0.0092	0.0246	0.1809 ***
Industrial production	0.0591	0.0433	−0.6916 ***
Short-term rate	0.0023 *	0.0026 *	−0.1211 ***
M3	0.0833	0.0520	4.2175 ***
Long-term Rate	−0.1323 ***	0.0776 ***	0.1175 ***
State 2			
Constant	0.0025	−0.0031	0.0355
Industrial Production	−0.0871	0.5430 *	−0.0843
Short-term Rate	−0.0027	−0.0003	0.0017
M3	0.5420 ***	0.4938 **	0.9388 ***
Long-term Rate	0.0316	−0.2040 ***	0.0007
ϕ_1	0.9438 ***	0.9589 ***	−0.3894 **
ϕ_2	0.9364 ***	0.9462 ***	0.9542 ***
σ_1	0.0210	0.0152	0.0224
σ_2	0.0406	0.0311	0.0316
p_{11}	0.9817	0.9917	0.9676
p_{22}	0.9716	0.9706	0.9967
Expected duration of the regimes in months			
State 1	54.64	121.00	30.82
State 2	35.22	34.00	301.26

Switching Lag MS. Significance at *** 1%, ** 5%, * 10%.

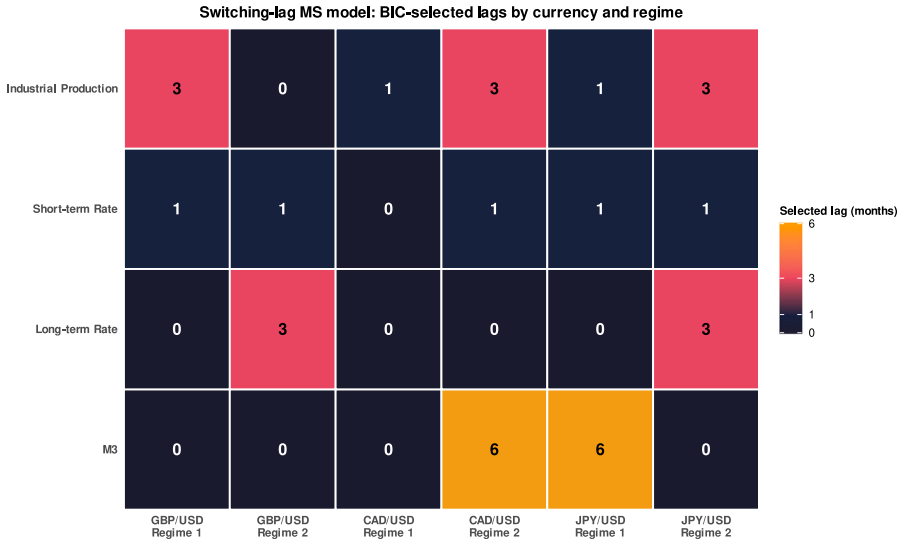

Fig. 1. Optimal lag structure from a Switching Lag Markov-Switching model. The heatmap displays selected lags for key covariates across two regimes for GBP/USD, CAD/USD, and JPY/USD.

The timing of the regimes aligns with well known episodes. For GBP/USD the high volatility state spikes around 1992 during the European Exchange Rate Mechanism crisis known as Black Wednesday and it reappears from 2007 to 2009 during the Global Financial Crisis and again from 2015 to 2016 around the United Kingdom referendum on European Union membership. Tranquil spans dominate from 1993 to 2006.

For the Canadian dollar against the U.S. dollar turbulence clusters in global and commodity shocks with a clear rise from 2007 to 2009 during the Global Financial Crisis and from 2014 to 2016 during the collapse in oil prices while the early to mid 2000s are calmer. This pattern is consistent with the change in sign of the long rate across states reported in Table 2.

For the yen we find a highly persistent high volatility state with larger variance and very high smoothed probability from the late nineteen nineties onward. Peaks align with the Asian financial crisis in 1997–1998 and with the global financial crisis in 2008–2009.

Out of sample the switching–lag model delivers gains that grow with horizon and vary by currency as shown in Table 4. Against the random walk it wins for GBP/USD at six and twelve months and for CAD/USD at twelve months while shorter horizons underperform. Against the MS-FL benchmark the strongest improvements arise for JPY/USD at three and twelve months and are statistically significant. For GBP/USD the twelve–month gain is large in magnitude yet not statistically different and for CAD/USD

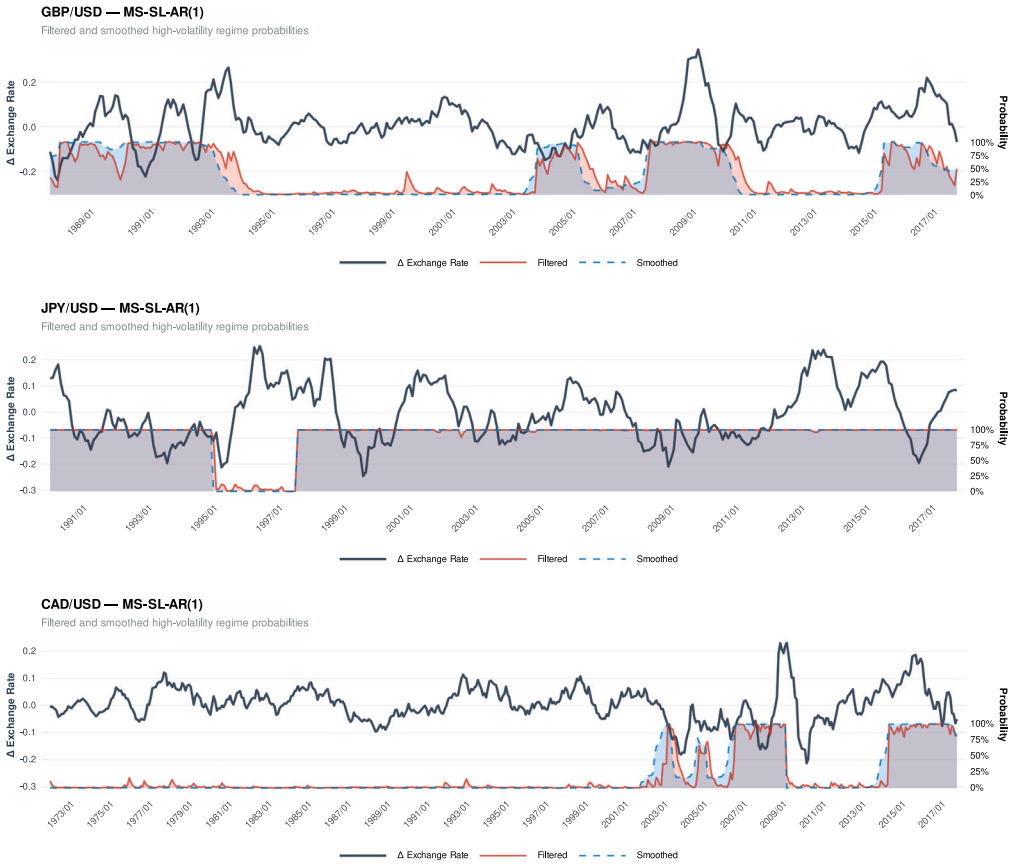

Fig. 2. Smoothed and filtered probabilities of Regime 2 (high volatility) for GBP/USD, CAD/USD, and JPY/USD under the MS-SL-AR(1) specification.

Table 3
Most frequent lags by regime, computed over the Top-10 lowest-BIC models per currency. Entries are mode (share%).

Currency	Regime	Industrial production	Short-term rate	Long-term rate	Money (M3)
GBP/USD	Low Vol	3 (90%)	1 (100%)	0 (100%)	6 (40%)
GBP/USD	High Vol	0 (50%)	1 (60%)	3 (60%)	0 (100%)
CAD/USD	Low Vol	0 (50%)	0 (70%)	0 (50%)	12 (40%)
CAD/USD	High Vol	3 (80%)	1 (90%)	0 (90%)	6 (80%)
JPY/USD	Low Vol	1 (50%)	1 (60%)	0 (40%)	6 (50%)
JPY/USD	High Vol	3 (60%)	0 (50%)	0 (60%)	0 (50%)

the long–horizon gains are modest and not significant. These results indicate that allowing timing to vary across regimes yields forecast value mainly at medium and long horizons and that the payoff is currency specific. Our findings align with survey evidence that fundamentals weigh more at longer horizons, matching the longer lags in calm states and the 6–12 month gains (Figs. 1–2; Table 4) (Beckmann and Czudaj, 2025). In turbulence, frequent short-horizon revisions documented for chartists rationalize near contemporaneous lags and currency-specific timing (Dick and Menkhoff, 2013), a pattern consistent with investor learning and anchoring in calm states and with weaker policy credibility under stress, without altering identification, estimation, or tests.

3.1. Robustness check

We augment the fundamentals with a risk-sentiment proxy (Chow et al., 1997) given by the 12-month change in the Moody's Baa–Aaa corporate bond yield spread (FRED series BAA and AAA) and a capital-flows proxy 12-month change in U.S. TIC bond flows (Brooks et al., 2004) from the U.S. Treasury's TIC system. Because coverage of the Baa–Aaa spread and TIC flows is limited, only the CAD in-sample window shifts to 1978-01 while for GBP and JPY are unchanged.

Table 4
Out-of-sample predictive performance of the Switching Lag Model. Columns report MSPE ratios versus a random walk (RW) and versus MS-FL-AR(1); test statistics are from the Clark–West (CW, one-sided; H1: SL outperforms RW) and Diebold–Mariano (DM, two-sided) tests. MSPE ratio < 1 favors the Switching Lag model.

Currency	Horizon (h)	vs. Random walk		vs. MS-FL-AR(1)	
		MSPE ratio	CW test	MSPE ratio	DM test
GBP/USD	1	1.652	0.018	1.286	0.871
	3	1.609	−0.352	1.790	1.480
	6	0.950	2.070**	0.900	−0.285
	12	0.406	2.753***	0.368	−1.579
CAD/USD	1	2.010	1.001	1.653	1.512
	3	2.386	−0.262	1.737	1.510
	6	1.991	0.316	1.490	1.538
	12	0.840	2.909***	0.763	−0.819
JPY/USD	1	1.297	0.156	0.069	−1.113
	3	1.042	1.212	0.266	−2.008**
	6	1.605	1.062	0.354	−1.571
	12	6.408	0.079	0.308	−1.676*

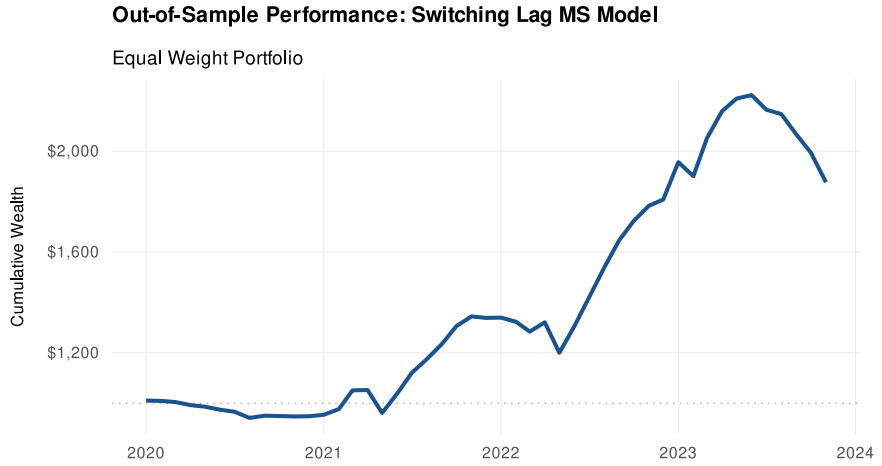

Fig. 3. Out-of-sample cumulative wealth from initial $1.000 of an equal-weight portfolio across GBP/USD, CAD/USD, and JPY/USD.

Adding risk and capital flow variables shifts GBP the most, with M3 moving to a 12-month lag in both regimes, short-rate lags dropping to 0 in Regime 1, and risk loading mainly in Regime 2 with BAA–AAA at lag 1 and TIC at lag 3. CAD changes are modest, with IP 1 to 0 and short rates 0 to 1 in Regime 1 and TIC at lag 3 in both, while JPY is largely stable except short rates in Regime 2 fall to 0 and BAA–AAA enters both regimes.

BIC worsens for all three pairs relative to the baseline, so the expanded models are not preferred in-sample (Table S4, S5 and S6). The full estimates are in Table 5 in Appendix C and lag choices appear in figure S5 in SM.

Adding risk and flow proxies preserves the core message and mostly reallocates effects across regimes. For GBP, the long rate stays negative in both states, M3 turns negative once risk is controlled, and the credit spread loads in calm states. For CAD, the long rate remains positive in calm and negative in turbulence, the short rate and output become clearly state-dependent, and M3 shifts toward calm. For JPY, the safe-haven pattern sharpens: the credit spread is positive in calm, TIC is negative in both states, production and the short rate stay negative, and M3 stays strongly positive. These proxies refine timing and leave signs and forecasts unchanged.

3.2. Portfolio allocation

Forecast gains matter if they convert into risk-adjusted returns. We therefore map the one-step forecasts directly into positions via a simple rule and an equal weight portfolio detailed in Appendix B. Out-of-sample windows begin in January 2019 for all pairs and the equal weight portfolio delivers an annualized excess Sharpe ratio of about 1.0042. The Lo-adjusted Sharpe ratio is about 0.7821 and the Probabilistic Sharpe Ratio equals 0.974. A Newey–West test with three lags for the mean excess return yields a t statistic of 1.57 and a p value of 0.12. Fig. 3 shows the gross cumulative wealth.

Allowing for monthly trading frictions of 20 bps proportional to average turnover $\bar{\tau}$, the annualized Sharpe declines from 1.0042 to $S_{\text{net}} \in [0.85, 0.94]$ for $\bar{\tau} \in [0.4, 1.0]$.

4. Conclusion

We test whether the timing of fundamentals is regime dependent with a two-state Markov switching regression with autoregressive errors and regime specific lags chosen by BIC. Timing differs across regimes and across currencies. Forecast gains are stronger at medium and long horizons, while short horizons are mixed. Timing matters, yet it does not fully resolve the disconnect.

The evidence comes from three liquid USD pairs with monthly data and a short out-of-sample window from 2019 to 2023. Forecasts rely on revised fundamentals and may condition on information that was not available in real time. Our portfolio analysis excludes costs and slippage, which compress realized P&L, so the reported results are an upper bound.

These results suggest that asset managers weight signals by regime and horizon, with most value at 6–12 months, and that risk teams allow effective lags to vary across states in stress tests and monitor liquidity and capital-flow conditions that can delay price discovery.

The mechanism of state-dependent lags is likely to extend within the G10, although dominant channels remain currency specific. Next steps include expanding to EUR USD and AUD USD and to major emerging pairs, bringing in fundamentals linked to risk and cross-border flows. Our results motivate a broader claim: timing is a state variable. If regimes proxy intermediation frictions or uncertainty, economically relevant lags should switch across asset classes, including equities, commodities, and sovereign bonds, especially at medium horizons. The methodology can be applied to these markets.

CRediT authorship contribution statement

Lucas M. Oliveira: Writing – original draft, Software, Methodology, Formal analysis, Conceptualization. **Airlane P. Alencar:** Writing – review & editing, Validation, Supervision.

Funding

Lucas M. Oliveira received a PhD scholarship from the Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES) – Brasil. Finance Code 001.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A

We estimate a two-regime MS-AR(1) by EM conditional on BIC-selected lag vectors (ℓ_1, ℓ_2) from the common grid for both regimes. The model is

$$\Delta e_t = \mu_{s_t}(t) + \phi_{s_t} r_{t-1}^{(s_{t-1})} + \varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, \sigma_{s_t}^2),$$

with $\mu_j(t) = \alpha_j + \boldsymbol{\beta}_j^\top \mathbf{X}_t^{(j)}$,

$$\mathbf{X}_t^{(j)} = \left[\Delta y_{t-\ell_j^{(y)}},\ \Delta i_{s,t-\ell_j^{(is)}},\ \Delta i_{l,t-\ell_j^{(il)}},\ \Delta m_{t-\ell_j^{(m)}}\right]^\top, \quad r_{t-1}^{(i)} = \Delta e_{t-1} - \mu_i(t-1).$$

the complete data log-likelihood is

$$\ell_c(\Theta) = \sum_{t=t_0}^{T} \sum_{i=1}^{2} \sum_{j=1}^{2} z_{t-1,t}(i,j) \left[\log p_{ij} - \frac{1}{2} \log \sigma_j^2 - \frac{\left(\Delta e_t - \mu_j(t) - \phi_j r_{t-1}^{(i)}\right)^2}{2\sigma_j^2} \right],$$

where $z_{t-1,t}^{(i,j)} = 1$ if $s_{t-1} = i$ and $s_t = j$, and 0 otherwise and $t_0 = L_{\max} + 1$. **E–step.** Set $P = \{p_{ij}\}_{i,j=1}^2$ and $t_0 = L_{\max} + 1$. For all $(i,j) \in \{1,2\}^2$, define

$$f_t(j \mid i) = \mathcal{N}\left(\Delta e_t;\ \mu_j(t) + \phi_j r_{t-1}^{(i)},\ \sigma_j^2\right).$$

Forward filter, for $t = t_0, \ldots, T$ and $j \in \{1,2\}$:

$$\tilde{f}_t(j) = \sum_{i=1}^{2} \xi_{t-1}(i)\, p_{ij}\, f_t(j \mid i), \qquad \xi_t(j) = \frac{\tilde{f}_t(j)}{\sum_{b=1}^{2} \tilde{f}_t(b)}.$$

Kim smoothing yields

$$\hat{\xi}_{t|T}(j) = \Pr(s_t = j \mid \mathcal{I}_T), \quad \hat{\Psi}_t(i,j) = \Pr(s_{t-1} = i, s_t = j \mid \mathcal{I}_T),$$

computed via standard forward–backward recursions.

M–step. Set

$$\mathbf{w}_t^{(i,j)} = \begin{bmatrix} \mathbf{X}_t^{(j)} \\ r_{t-1}^{(i)} \end{bmatrix} \in \mathbb{R}^6, \qquad \boldsymbol{\theta}_j = \begin{bmatrix} \boldsymbol{\beta}_j \\ \boldsymbol{\phi}_j \end{bmatrix} \in \mathbb{R}^6.$$

At EM iteration k, use weights $\omega_t^{(i,j,k)} = \hat{\Psi}_t(i,j) \big/ \hat{\sigma}_j^{2\,(k)}$ and form

$$\mathbf{G}_j^{(k)} = \sum_{t=t_0}^{T} \sum_{i=1}^{2} \omega_t^{(i,j,k)}\, \mathbf{w}_t^{(i,j)} \mathbf{w}_t^{(i,j)\top}, \qquad \mathbf{g}_j^{(k)} = \sum_{t=t_0}^{T} \sum_{i=1}^{2} \omega_t^{(i,j,k)}\, \mathbf{w}_t^{(i,j)}\, \Delta e_t.$$

Update regression and AR parameters:

$$\hat{\boldsymbol{\theta}}_j^{(k+1)} = \left(\mathbf{G}_j^{(k)}\right)^{-1} \mathbf{g}_j^{(k)}.$$

Update variances and transition probabilities:

$$\hat{\sigma}_j^{2\,(k+1)} = \frac{\displaystyle\sum_{t=t_0}^{T} \sum_{i=1}^{2} \Psi_t(i,j) \left[\Delta e_t - \mathbf{w}_t^{(i,j)\top} \boldsymbol{\theta}_j^{(k+1)}\right]^2}{\displaystyle\sum_{t=t_0}^{T} \sum_{i=1}^{2} \Psi_t(i,j)}, \qquad \hat{p}_{ij}^{(k+1)} = \frac{\displaystyle\sum_{t=t_0}^{T} \hat{\Psi}_t(i,j)}{\displaystyle\sum_{t=t_0}^{T} \sum_{b=1}^{2} \hat{\Psi}_t(i,b)}.$$

We stop when the absolute change in log-likelihood falls below 10^{-4} or upon reaching 500 iterations, whichever comes first.

Model selection strategy

We jointly select the regime-specific lag vectors (ℓ_1, ℓ_2) by minimizing the BIC over a common grid used for both regimes:

$$\ell_j^{(y)} \in \{0,1,3\}, \quad \ell_j^{(is)} \in \{0,1\}, \quad \ell_j^{(il)} \in \{0,3,6\}, \quad \ell_j^{(m)} \in \{0,6,12\}, \qquad j \in \{1,2\}.$$

This yields 54^2 candidates for (ℓ_1, ℓ_2). For forecasts, the lag vectors of both regimes are fixed at the in-sample BIC winners for each currency.

Appendix B

For currency c, the directional weight at $t+1$ is the sign of the forecast:

$$\pi_{t+1}^{(c)} = \operatorname{sgn}\left(\widehat{\Delta e}_{t+1|I_t}^{(c)}\right) \in \{-1, +1\}.$$

With equal weights across the available set C_t, the portfolio excess return is

$$r_{t+1}^{p,e} = \frac{1}{|C_t|} \sum_{c \in C_t} w_{t+1}^{(c)}\, \Delta e_{t+1}^{(c)} - r_{t+1}^{f}.$$

Positions are set at the end of month t and applied to the return of month $t+1$. The portfolio remains fully invested with no leverage. We assign equal weights across the pairs available each month with shorting allowed and we rebalance every month.

Excess returns are computed by subtracting the risk-free rate measured by the three month Treasury bill from FRED with code TB3MS.

Appendix C

See Table 5.

Appendix D. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.frl.2025.108941.

Data availability

Data will be made available on request.

Table 5
Estimated Coefficients for the Switching Lag MS with augmented Risk sentiment and Capital Flow.

	GBP/USD	CAD/USD	JPY/USD
State 1			
Constant	−0.0133	−0.0307	0.1774***
Industrial production	0.0638*	−0.1979***	−0.5701***
Short-term rate	0.0006	−0.0367***	−0.0943***
M3	−0.1656*	0.3592***	3.6800***
Long-term rate	−0.1234***	0.1769***	0.1057***
Baa–Aaa	0.0548***	−0.0095**	0.1404***
TIC	−0.0001	−0.0003	−0.0018***
ϕ_1	0.9074***	0.9744***	−0.6177***
σ_1	0.0200	0.0141	0.0141
State 2			
Constant	0.0524	0.0054	0.0450*
Industrial production	−0.1093	0.1597*	−0.1059**
Short-term rate	−0.0015	−0.0001	0.0007
M3	−0.8066***	−0.1930	0.8604**
Long-term rate	−0.0607**	−0.1565***	0.0009
Baa–Aaa	0.0777	0.0899***	−0.0156
TIC	−0.0003	−0.0001	−0.0002*
ϕ_2	0.9292***	0.9505***	0.9578***
σ_2	0.0387	0.0200	0.0316
Expected duration of the regimes in months			
p_{11}	0.9790	0.9894	0.9671
p_{22}	0.9598	0.9893	0.9895
Duration state 1	47.6190	94.3396	30.3670
Duration state 2	24.8501	93.4579	95.2381

Switching lag Markov-switching with risk proxies BAA–AAA and TIC. Significance: *** 1%, ** 5%, * 10%.

References

Abedin, M.Z., Moon, M.H., Hassan, M.K., Hajek, P., 2025. Deep learning-based exchange rate prediction during the COVID-19 pandemic. Ann. Oper. Res. 345 (2), 1335–1386. http://dx.doi.org/10.1007/s10479-021-04420-6, First online: 26 Nov 2021.

Bailey, D.H., Lopez de Prado, M., 2012. The sharpe ratio efficient frontier. J. Risk 15 (2), 13.

Beckmann, J., Czudaj, R.L., 2025. Fundamental determinants of exchange rate expectations. Int. J. Forecast. 41 (3), 1003–1021. http://dx.doi.org/10.1016/j.ijforecast.2024.09.004.

Beckmann, J., Koop, G., Korobilis, D., Schüssler, R.A., 2020. Exchange rate predictability and dynamic Bayesian learning. J. Appl. Econometrics 35 (4), 410–421.

Brooks, R., Edison, H., Kumar, M.S., Sløk, T., 2004. Exchange rates and capital flows. Eur. Financ. Manag. 10 (3), 511–533.

Chow, E.H., Lee, W.Y., Solt, M.E., 1997. The exchange-rate risk exposure of asset returns. J. Bus. 105–123.

Clark, T.E., West, K.D., 2007. Approximately normal tests for equal predictive accuracy in nested models. J. Econometrics 138 (1), 291–311.

Dautel, A.J., Härdle, W.K., Lessmann, S., Seow, H.-V., 2020. Forex exchange rate forecasting using deep recurrent neural networks. Digit. Financ. 2 (1), 69–96.

Dewachter, H., 2001. Can Markov switching models replicate chartist profits in the foreign exchange market? J. Int. Money Financ. 20 (1), 25–41.

Dick, C.D., Menkhoff, L., 2013. Exchange rate expectations of chartists and fundamentalists. J. Econom. Dynam. Control 37 (7), 1362–1383.

Diebold, F.X., Mariano, R.S., 2002. Comparing predictive accuracy. J. Bus. Econom. Statist. 20 (1), 134–144.

Dornbusch, R., 1976. Expectations and exchange rate dynamics. J. Political Econ. 84 (6), 1161–1176.

Engel, C., 1994. Can the Markov switching model forecast exchange rates? J. Int. Econ. 36 (1–2), 151–165.

Engel, C., Wu, S.P.Y., 2023. Liquidity and exchange rates: An empirical investigation. Rev. Econ. Stud. 90 (5), 2395–2438.

Frankel, J.A., 1979. On the mark: A theory of floating exchange rates based on real interest differentials. Am. Econ. Rev. 69 (4), 610–622.

Frömmel, M., MacDonald, R., Menkhoff, L., 2005. Markov switching regimes in a monetary exchange rate model. Econ. Model. 22 (3), 485–502.

Gabaix, X., Maggiori, M., 2015. International liquidity and exchange rate dynamics. Q. J. Econ. 130 (3), 1369–1420.

Hamilton, J.D., 1989. A new approach to the economic analysis of nonstationary time series and the business cycle. Econ.: J. Econ. Soc. 357–384.

Jeanne, O., Rose, A.K., 2002. Noise trading and exchange rate regimes. Q. J. Econ. 117 (2), 537–569.

Kim, C.-J., Nelson, C.R., 1999. State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications. The MIT Press, Cambridge, MA.

Kumar, U., Ahmad, W., Uddin, G.S., 2024. Bayesian Markov switching model for BRICS currencies' exchange rates. J. Forecast. 43 (6), 2322–2340.

Li, N., Kwok, S.S., 2021. Jointly determining the state dimension and lag order for Markov-switching vector autoregressive models. J. Time Series Anal. 42 (4), 471–491.

Lo, A.W., 2002. The statistics of sharpe ratios. Financ. Anal. J. 58 (4), 36–52.

Lundbergh, S., Teräsvirta, T., Van Dijk, D., 2003. Time-varying smooth transition autoregressive models. J. Bus. Econom. Statist. 21 (1), 104–121.

Meese, R.A., Rogoff, K., 1983. Empirical exchange rate models of the seventies: Do they fit out of sample? J. Int. Econ. 14 (1–2), 3–24.

Primiceri, G.E., 2005. Time varying structural vector autoregressions and monetary policy. Rev. Econ. Stud. 72 (3), 821–852.

Said, S.E., Dickey, D.A., 1984. Testing for unit roots in autoregressive-moving average models of unknown order. Biometrika 71 (3), 599–607.

Stillwagon, J., Sullivan, P., 2020. Markov switching in exchange rate models: will more regimes help? Empir. Econ. 59 (1), 413–436.

Zhou, Z., Fu, Z., Jiang, Y., Zeng, X., Lin, L., 2020. Can economic policy uncertainty predict exchange rate volatility? New evidence from the GARCH-MIDAS model. Financ. Res. Lett. 34, http://dx.doi.org/10.1016/j.frl.2019.08.006.