Accepted Manuscript

Nonlinear exchange rate dynamics and behavioral expectations: New evidence from selected emerging and advanced economies

Lorenzo Biondi and Christian R. Proaño

First published in *Journal of Evolutionary Economics*, volume 36, article 52 (2026), by Springer Nature.

This version of the article has been accepted for publication, after peer review (when applicable) but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: https://doi.org/10.1007/s00191-026-00970-9.

Use of this Accepted Version is subject to the publisher's Accepted Manuscript terms of use.

Nonlinear Exchange Rate Dynamics and Behavioral Expectations: New Evidence from Selected Emerging and Advanced Economies

Lorenzo Biondi[1*] and Christian R. Proaño[2,3]

[1*]University of Siena, Siena, Italy.
[2]Otto-Friedrich-Universität Bamberg, Bamberg, Germany.
[3]Centre for Applied Macroeconomic Analysis, Australian National University, Canberra, Australia.

*Corresponding author(s). E-mail(s): lorenzobiondi.econ@gmail.com;

Abstract

Uncovered interest parity (UIP) is routinely rejected at short horizons, motivating departures toward heterogeneous-expectations frameworks. We estimate a behavioral UIP (BUIP) law of motion for monthly exchange-rate returns, where expectations form as a profitability-weighted average of a chartist (trend-following) rule and a fundamentalist rule anchored to absolute purchasing power parity (PPP). Building on Proaño (2011, 2013), rule shares evolve through a discrete-choice mechanism. We compare the structural BUIP against a reduced-form logistic smooth-transition regression across fourteen bilateral USD exchange rates (seven advanced, seven emerging) from the early 1980s to 2025. We reject linearity for 13 of 14 currencies, with nonlinear models reducing in-sample fit. However, the Akaike information criterion frequently prefers the linear benchmark in advanced economies, suggesting nonlinearity is most rewarded in emerging markets. The structural BUIP identifies stabilizing PPP-consistent mean reversion in seven countries. While embedding heterogeneous beliefs into UIP partly reconciles the parity condition with nonlinear dynamics, the mixed evidence on PPP correction highlights the limitations of a "one-model-fits-all" approach. Ultimately, the choice of fundamental anchor appears to be highly market- and time-dependent.

Keywords: Exchange rates, Uncovered interest parity, Purchasing power parity, Heterogeneous expectations, Discrete-choice switching, Smooth transition regression

JEL Classification: F31 , F32 , G15 , D84 , C22 , C52

1 Introduction

Nearly half a century after the seminal contribution by Meese and Rogoff (1983), nominal exchange rates remain notoriously hard to explain and forecast. Canonical parity conditions such as uncovered interest parity (UIP) and purchasing power parity (PPP) are routinely rejected, giving rise to the forward-premium anomaly (often called the "Fama puzzle") and the PPP puzzle (Rogoff 1996; Engel 2014). A growing body of evidence attributes these anomalies to the interaction of heterogeneous market participants, most prominently, fundamentalists who anchor expectations to macroeconomic values and chartists who extrapolate recent trends (de Grauwe and Grimaldi 2006; Proaño 2013; Bassi et al. 2023; Mignot and Westerhoff 2025). Such behavioral models can reproduce excess volatility, fat-tailed returns, and regime-like swings, but robust empirical validation across currencies and long samples remains limited.

Understanding exchange-rate dynamics is central for monetary policy and macroprudential design. Persistent misalignments can erode competitiveness, while sudden depreciations can threaten financial stability, especially in economies with sizable foreign-currency liabilities. Linear parity regressions are not well suited to these features because adjustment pressures plausibly depend on the market state: arbitrage and mean reversion may be weak close to equilibrium but strengthen when deviations become large (Sarno et al. 2006). Nonlinear time-series frameworks, such as smooth-transition models, offer a reduced-form way to capture state dependence, while heterogeneous-agent models provide a structural mechanism (endogenous switching between forecasting heuristics) capable of generating nonlinear mean dynamics (Taylor et al. 2001; de Grauwe and Grimaldi 2006). Empirically documenting how far such mechanisms can explain exchange-rate returns across a diverse set of currencies is therefore of direct interest to both international macroeconomics and policy.

Against this background, this paper tests UIP through the lens of a heterogeneous agent model with behavioral expectations developed by Proaño (2011, 2013). We estimate a behavioral UIP (BUIP) law of motion in which monthly exchange-rate returns depend on the interest rate differential and on market expectations formed by two forecasting rules: (i) a chartist trend-following rule and (ii) a fundamentalist rule that anchors expectations to the absolute PPP value. The market share of each rule evolves endogenously via a discrete-choice mechanism driven by relative performance, implying time variation in the strength of momentum versus PPP-consistent mean reversion (Brock and Hommes 1997). We include a cost parameter k in the fundamentalist utility, which we estimate without imposing a sign restriction. To disentangle reduced-form state dependence from the structural switching implied by the theory, we estimate both (a) a logistic smooth-transition regression (STR), where regime changes are driven by observable transition variables (misalignment or volatility proxies), and (b) the structural BUIP specification, where the transition is endogenized through the profitability-based utilities of the two rules. Specifically, we investigate (i) whether the estimated transition function captures the hypothesized switch between chartist and fundamentalist dominance; (ii) whether the resulting regime-dependent coefficients possess the expected signs and statistical significance; and (iii) whether systematic differences emerge between the economies under study, with a particular focus on whether the switching intensity $\hat{\gamma}$ behaves as a market-specific diagnostic

of behavioral fragility (smooth vs. herding-type step transitions) and whether such fragility maps onto the advanced/emerging divide.

We use monthly data for fourteen bilateral USD exchange rates—seven advanced economies (Australia, Canada, Euro Area, Japan, New Zealand, Switzerland, United Kingdom) and seven emerging economies (Brazil, Indonesia, Korea, Mexico, Philippines, Thailand, Türkiye)—with unbalanced samples extending back to the early 1980s and ending in 2025. Our results can be summarized as follows. First, the reduced-form STR indicates pervasive state dependence: linearity is rejected for 13 out of 14 currencies, and nonlinear specifications improve in-sample fit relative to a linear benchmark in terms of RMSE for 13 out of 14 countries. Nevertheless, once complexity is penalized, the evidence is weaker in several advanced economies, where the AIC frequently prefers the linear model, suggesting that nonlinearity is rewarded in emerging markets. Second, the STR delivers limited identification of the regime-dependent coefficients: significant momentum is found in few currencies, and PPP-consistent mean reversion in only a few cases, sometimes with the wrong sign. Third, the structural BUIP specification yields a clearer decomposition into fundamentalist and chartist forces and identifies PPP-consistent mean reversion more frequently: once we restrict the count to currencies for which the switching point itself is identified, the fundamentalist adjustment parameter is negative and statistically significant in seven currencies, while it is imprecisely estimated in several others and turns significantly positive in Mexico and Türkiye, suggesting that PPP is not an equally valid anchor across markets. Fourth, the chartist component is heterogeneous: some currencies display significant momentum, whereas others exhibit significant contrarian behavior at the monthly horizon. Fifth, the estimated transition function takes a discontinuous, threshold-like shape more frequently than a smooth one—read, following Granger and Teräsvirta (1993a), as a threshold (TR) transition rather than as non-identification, since identification is diagnosed via the empirical support of $\hat{k}$ in the BUIP and $\hat{c}$ in the STR—and the prior that emerging markets exhibit more step-like switching than advanced ones is not confirmed in the data: both groups contain smooth and step-like specifications. A pre-COVID robustness exercise corroborates the main findings while clarifying which identification failures are structural (the United Kingdom under STR, Türkiye under BUIP) and which are partly an artefact of the pandemic period (Japan and Korea under BUIP). Overall, allowing heterogeneous expectations to enter UIP helps to rationalize nonlinearities in exchange-rate dynamics, but the empirical support for structural switching and for a unique PPP-based fundamental anchor remains uneven across currency pairs, and a common monthly evaluation horizon for both rules is a useful tractability assumption rather than a supported empirical feature.

This work enters into the literature on nonlinear exchange-rate dynamics and the increasing literature on empirical validation of heterogeneous-agent models (see Lux and Zwinkels (2018) for a first overview of the literature). Recent works have contributed in this area by proposing different methods to estimate the parameters of heterogeneous-agent models, such as simulated maximum likelihood (Kukacka and Sacht 2023), state-space approaches (Gusella and Ricchiuti 2024), machine learning techniques (Di Francesco and Hommes 2025), and likelihood-based methods (Frijns

et al. 2025). From a methodological point of view, Lux (2024) sheds light on the limits of identifiability of behavioral parameters, suggesting more caution in interpreting previous empirical findings. In our work, we contribute by estimating the reduced form and structural model on a vast and heterogeneous amount of data (14 countries over long samples), thereby providing a comprehensive cross-country comparison of the evidence for nonlinearities and heterogeneous expectations in exchange-rate dynamics, reducing the risk of overfitting due to small samples or specific market conditions. Our findings suggest that the tradeoff between maintaining model parsimony and capturing nonlinearities is present, and a "one model fits all" approach may not be appropriate for exchange-rate dynamics across different economies.

The remainder of this paper is organized as follows. Section 2 reviews theoretical and empirical work on exchange-rate puzzles, heterogeneous-agent models, and nonlinear econometric approaches. Section 3 presents the behavioral UIP framework and its econometric mapping and the estimation strategy. Section 4 presents the data and the empirical results. Section 5 concludes and discusses limitations and possible extensions.

2 Literature Review

As is widely known, two benchmark parity conditions organize much of international finance: uncovered interest parity (UIP) and purchasing power parity (PPP), in absolute or relative form. Persistent empirical failures of these conditions, often summarized as "puzzles" (e.g., the exchange-rate disconnect, excess volatility, and the forward-premium/UIP anomaly), have motivated a large literature that enriches the canonical macro view with nonlinear dynamics (Taylor et al. 2001), heterogeneous trading behavior (de Grauwe and Grimaldi 2006) and market microstructure mechanisms (Lyons 2001).

Classic macroeconomic approaches. The traditional open-economy literature treats the exchange rate as an asset price driven by monetary fundamentals, nominal rigidities, and imperfect asset substitutability. Flexible-price monetary models (Frenkel 1977; Bilson 1978) recast the exchange rate as the relative price of monies, while sticky-price formulations deliver the canonical overshooting dynamics of Dornbusch (1976) and the real-rate synthesis of Frankel (1979). Portfolio-balance extensions allow for a time-varying risk premium driven by relative asset supplies (Branson and Henderson 1985; Frenkel and Mussa 1985), though later evidence points to signaling and coordination rather than to lasting portfolio effects (Sarno and Taylor 2001). The "new open-economy macroeconomics" following Obstfeld and Rogoff (1995) embeds these mechanisms in micro-founded DSGE environments with optimizing agents and nominal rigidities (Lane 2001; Galí and Monacelli 2005), and the present-value approach of Engel and West (2005) shows that exchange rates forecast future fundamentals more effectively than fundamentals forecast exchange rates, consistent with the persistent "disconnect" observed at short horizons. Despite their theoretical appeal, these frameworks struggled to reproduce the high-frequency volatility and the weak short-run connection to fundamentals documented in the data, motivating the behavioral and microstructural perspectives reviewed below.

Market microstructure and information aggregation. Macroeconomic models typically treat the exchange rate as an asset price determined by macro fundamentals. Market microstructure opens the trading "black box" of over-the-counter dealer markets where private information and inventory management matter. Order flow—the net of buyer-over seller-initiated trades—emerges as a proximate driver of dealers' quotes and transaction prices (Evans 2011). Using high-frequency data, Evans and Lyons (2002) show that order flow explains a substantial share of daily exchange-rate variation, consistent with dealers updating quotes in response to information conveyed by trades. This view explains seemingly "newsless" exchange-rate movements: the relevant information is in order flow rather than in public announcements (Lyons 2001). Survey evidence aligns with this interpretation: Cheung and Chinn (2001) report that FX professionals view fundamentals as important at long horizons, while technical trading and bandwagon effects matter at short horizons.

2.1 Heterogeneous-agent and behavioral exchange-rate models

Chartists vs. fundamentalists and endogenous switching. Building on standard NOEM frameworks, a growing literature introduces heterogeneous expectations to capture the coexistence of distinct trading rules in FX markets. In the canonical setup, traders switch between a *chartist* (trend-following, technical) rule and a *fundamentalist* rule that anchors expectations to macroeconomic fundamentals. Following the discrete-choice switching mechanism of Brock and Hommes (1997), strategy weights evolve with relative past performance, yielding nonlinear aggregate dynamics with persistent swings and prolonged parity deviations, punctuated by abrupt corrections—an idea that goes back to Frankel and Froot (1990). Early contributions already emphasized that nonlinear interactions between chartists and fundamentalists can generate complex exchange-rate dynamics. de Grauwe and Dewachter (1993) embed these trading rules in a monetary model and show that the resulting nonlinear speculative dynamics can exhibit chaotic behavior, reproducing unit-root-like properties and biased forward premia even in the absence of explicit risk premia. Related work emphasizes "animal spirits" and sentiment dynamics: De Grauwe and Rovira Kaltwasser (2012) develop a model with evolutionary switching among optimistic and pessimistic beliefs (and, in extensions, fundamentalists and chartists), generating endogenous volatility, fat tails, and potentially chaotic dynamics. These models interpret excess volatility and "disconnect" as emergent outcomes of bounded rationality and strategy interaction rather than as failures of arbitrage per se.

Structural implementations and policy-relevant open-economy extensions. A first structural implementation on a DSGE skeleton is de Grauwe and Grimaldi (2006), who embed evolutionary switching in a dynamic macro model and reproduce momentum, excess volatility, bubbles, and crashes while preserving a long-run fundamentals anchor. In the same vein, Proaño (2011, 2013) augment a New-Keynesian DSGE core with a behavioral modification of UIP, combining IS, Phillips curve, and Taylor rule with an exchange-rate equation driven by heterogeneous expectations and switching. Recent work extends heterogeneous-trader exchange-rate mechanisms into richer macro agent-based environments, explicitly linking FX dynamics to real activity, trade, and policy interventions. Gori and Ricchiuti (2018) propose a dynamic

exchange-rate model with heterogeneous agents where the *fundamental* exchange rate is endogenously determined by balance-of-payments conditions, thereby integrating heterogeneous speculation with real-side adjustment. Going further, Delli Gatti et al. (2025) embed a currency market with chartists and fundamentalists into a two-country macro ABM with firm heterogeneity, innovation (R&D), and international trade. In their framework, the exchange rate is jointly shaped by trade flows (real channel) and speculative demand for foreign currency (financial channel), and policy experiments with central-bank "leaning against the wind" highlight non-trivial trade-offs: aggressive interventions can compress speculative activity and cycle duration, but may also affect volatility in unintended ways.

Empirical evidence of heterogeneous agent models. Empirical validation of heterogeneous-agent models increasingly combines time-series estimation, micro-level expectation data, and methodological advances aimed at improving identification. Using monthly major USD exchange rates (1974–1998), Manzan and Westerhoff (2007) estimate a structural chartist–fundamentalist model and report strong in-sample fit for most currencies, with chartist behavior becoming stabilizing when deviations from fundamentals are sufficiently large. Micro-level evidence from survey expectations reinforces the case for persistent heterogeneity: Goldbaum and Zwinkels (2014), using Consensus Economics forecasts, classify institutions by forecasting rule and show that allowing for dynamic switching materially improves fit relative to static classifications, with switching related to both period-specific and forecaster-specific factors. At the same time, several contributions highlight that the strength of evolutionary selection can be weak or unstable, raising robustness and finite-sample identification concerns; Kukacka and Barunik (2017) propose nonparametric simulated maximum likelihood for financial agent-based models and find that the switching coefficient can be statistically insignificant even when trend-following behavior is robust. These concerns have motivated richer empirical implementations that explicitly confront nonlinear dynamics and potential overparameterization. For instance, Franke and Westerhoff (2011) estimate a two-type (fundamentalist/chartist) agent-based asset-pricing model with structural stochastic volatility via the method of simulated moments, obtaining a good fit for daily S&P 500 returns and evidence, in an FX application to USD–DEM, of parameter changes consistent with a weaker role for chartism. Complementing simulation-based structural estimation, Westerhoff and Reitz (2003) estimate a nonlinear chartist–fundamentalist exchange-rate model using a STAR–GARCH specification on daily USD spot rates (against the DM, BP, and JPY), providing empirical support for regime-dependent dynamics consistent with heterogeneous trading and nonlinear mean reversion, while also questioning the presumption that fundamentalists are always stabilizing. Recent work extends this empirical frontier across data and methods. ter Ellen et al. (2021) provide a comprehensive cross-asset comparison of behavioral heterogeneity, showing that it is larger and more persistent in macroeconomic variables than in financial assets. Gusella and Ricchiuti (2024) develop a state-space approach to identify endogenous cycles in the S&P 500, treating agents' unobservable beliefs as hidden state components. They find evidence of endogenous damped fluctuations with a higher probability of chartist behavior in the short-term horizon. Kukacka and Sacht (2023) estimate a heuristic switching New Keynesian

model via a multivariate version of Simulated Maximum Likelihood, providing evidence of the switching mechanism in the US data. On the methodological front, Lux (2024) investigate the identification of structural parameters in a behavioral New Keynesian model using Hamiltonian Markov Chain Monte Carlo and find that behavioral parameters are poorly identified at typical macroeconomic sample sizes, suggesting that many previous empirical findings may be inconclusive. Galanis et al. (2025) generalize the BH discrete-choice framework to the probit specification, showing that standard logit remains a good approximation in two-rule settings but that its independence of irrelevant alternatives restriction becomes empirically binding once more than two alternatives are considered. In FX markets specifically, Bassi et al. (2023) illustrates an alternative agent-based route to reproduce the main stylized facts of exchange-rate dynamics, without endogenous strategy switching. They do not explicitly estimate the parameters of a switching model, but they provide evidence that the coexistence of different trading rules can generate realistic exchange rate dynamics. Frijns et al. (2025) introduce sentiment into a heterogeneous agent model and estimate it on S&P 500 data via MLE showing that high sentiment leads to more extreme beliefs and reduces switching behavior. Di Francesco and Hommes (2025) propose a model of sentiment-driven speculation for the Bitcoin market, using a neural network to approximate rational expectations and finding that while fundamentalist speculation can amplify volatility, rational speculation has a stabilizing effect.

2.2 Nonlinear dynamics and empirical evidence beyond linear parity models

Nonlinearity from limits to arbitrage and state dependence. A key implication of heterogeneous-agent and behavioral models is *state-dependent* dynamics. Near equilibrium, chartist and fundamentalist pressures may coexist and small shocks can be amplified if technical trading becomes temporarily more profitable; as misalignment grows, mean-reversion strategies become relatively more profitable, inducing a regime shift and accelerating correction. This mechanism yields nonlinear adjustment consistent with long stretches of near-random-walk behavior interspersed with sharp corrections. Empirical support for nonlinear exchange-rate dynamics also comes from smooth transition and regime-switching time-series models. Foundational treatments of STR/STAR modeling include Granger and Teräsvirta (1993b); Franses and Van Dijk (2000); Teräsvirta (1994); van Dijk et al. (2002); Granger and Teräsvirta (1993a); Tay and Chen (2018). For PPP, a central result is nonlinear mean reversion in real exchange rates: adjustment is weak near equilibrium but accelerates when misalignment is large. In an ESTAR model, Michael et al. (1997) attribute this to transaction costs creating a band of inaction; Taylor et al. (2001) similarly document nonlinear mean reversion and reject the random walk for real exchange rates. Markov-switching models provide an alternative reduced-form representation of persistent exchange-rate "swings." Engel and Hamilton (1990) show that regime-dependent specifications can outperform the random-walk benchmark for the US dollar, consistent with latent-state shifts in conditional drifts. Continuous-time heterogeneous-beliefs models can in principle generate chaotic dynamics (de Grauwe and Dewachter 1993), but empirical tests typically reject strict chaos. Federici and Gandolfo (2012) fit such a

model to the EUR/USD and reject chaotic behavior using standard diagnostics (phase plots, correlation dimension, Lyapunov exponents). The broader point survives: economically motivated heterogeneity can induce intrinsic nonlinearities observationally close to stochastic processes. Mignot and Westerhoff (2025) illustrate this with an agent-based version of de Grauwe's model that reproduces bubbles, crashes, excess volatility, fat-tailed returns, and volatility clustering for major currency pairs at daily frequency (1999–2022).

2.3 Fundamentals, misalignment, and the horizon-dependent role of macro variables

The heterogeneous-agent models reviewed in Section 2.1 rationalize a well-documented empirical regularity: fundamentals exert greater influence at long horizons than at short ones. Survey evidence consistently maps this horizon dependence onto strategy type. Cheung and Chinn (2001) report that professional forecasters weight fundamentals heavily at horizons of a year or more, whereas technical and bandwagon considerations dominate at shorter horizons; Menkhoff and Taylor (2007) find that more than ninety percent of FX dealers employ technical analysis, predominantly at short horizons, with the weight placed on fundamentals rising monotonically as the forecast horizon lengthens. Microstructure evidence reinforces this asymmetry. Public macroeconomic announcements account for only a small share of exchange-rate movements because the relevant information enters through order flow rather than news releases; fundamentals are thus revealed gradually as private beliefs are aggregated through trading (Lyons 2001; Evans and Lyons 2002, 2008).

Within the HAM framework, the switching mechanism itself furnishes a structural account of this pattern. When chartists dominate, the exchange rate drifts away from its fundamental anchor with limited correction; as misalignment widens, fundamentalist strategies become relatively more profitable and mean reversion strengthens (Manzan and Westerhoff 2007; Goldbaum and Zwinkels 2014). The effective horizon at which fundamentals exert influence is therefore endogenous and contingent on market conditions rather than fixed ex ante. Macroeconomic ABM extensions further emphasize that the fundamental anchor is itself endogenously shaped by trade flows and real activity (Gori and Ricchiuti 2018; Delli Gatti et al. 2025), implying that the horizon over which it operates may vary systematically across economies.

The survey evidence also raises a methodological tension pertinent to our study (see (Bassi et al. 2023) for a discussion): if technical strategies operate at intraday-to-weekly horizons while fundamental analysis operates at horizons of six to twelve months, evaluating both rule types at a common monthly frequency entails a simplification. Two considerations mitigate this tension. First, the horizon-strategy mapping is not rigid in practice: Dick and Menkhoff (2013) show that even fundamentalist-oriented forecasters incorporate technical signals at shorter horizons, and Gehrig and Menkhoff (2005) document that fund managers routinely blend both approaches at intermediate horizons. Second, the monthly frequency is itself an intermediate horizon—long enough to smooth the microstructure noise most relevant to pure chartism, yet short enough to remain responsive to accumulating PPP misalignment. This common-frequency specification follows the behavioral DSGE literature in which the BUIP is embedded

(Proaño 2011, 2013); we examine the common-horizon assumption and its implications more formally in Section 3.2.

2.4 Uncovered interest parity and the evolving evidence on the "Fama puzzle"

The horizon and switching evidence reviewed above also bears directly on the empirical content of UIP itself, since deviations from UIP are precisely the carry-trade profitability that drives behavioral switching in the BUIP framework. We turn next to how that evidence has evolved.

UIP asserts that the expected depreciation equals the interest rate differential (under rational expectations and a zero risk premium), implying a unit slope in the Fama regression. Empirically, short-horizon tests typically reject UIP and often produce negative slope estimates (the forward-premium anomaly). More recent evidence stresses that the relationship is *time-varying* and sensitive to monetary policy regimes and the level of interest rates. Bussiere et al. (2018) document that the Fama coefficient changed sign around the global financial crisis and then flipped again, and argue that these reversals are not well explained by standard risk proxies such as the VIX. Importantly, when replacing ex post depreciation with survey-based expectations, UIP regressions exhibit fewer apparent breaks and expected depreciation is much more consistently positively correlated with interest differentials, pointing to a central role for expectation errors and learning rather than a stable risk-premium story alone. Focusing on the US dollar, Engel et al. (2022) show that the predictability of excess returns from interest differentials becomes weak when interest rates are near the zero lower bound (ZLB) and that the relationship is unstable over time. They also report that inflation differentials provide more stable predictive content for future excess returns than interest differentials, consistent with delayed market adjustment to predictable monetary-policy changes. At the same time, there is renewed evidence that UIP may hold in the long-run even if it fails at short horizons. Baillie et al. (2026) test long-horizon UIP using Durbin regression methods designed to improve efficiency in the presence of serial correlation and dynamics, and conclude that while UIP is rejected at short horizons, the long-run restriction implied by UIP (and the long-run International Fisher relation) is not rejected for major currencies over extended samples. Taken together, the recent literature suggests a nuanced view: UIP is a poor short-run description of exchange-rate dynamics, but its long-run content may be closer to the data once time variation, policy regimes, and expectation formation are properly accounted for—a perspective that aligns naturally with heterogeneous-expectations models in which carry-trade profitability and deviations from UIP are state-dependent.

The literature shows a clear structural narrative. Different expectations and shifting conditions can create nonlinear exchange-rate dynamics. They can also explain the varying roles of fundamentals and reshape the empirical content of UIP. However, identifying these effects is challenging. Reduced-form nonlinear models may show state dependence but don't clarify the economic mechanisms. Structural switching models can struggle if parameters are weakly identified or if the chosen fundamental proxy is incorrect. This need drives our empirical strategy. Next, we derive the behavioral UIP law of motion. We map it into (i) a reduced-form LSTR specification and (ii)

a structural BUIP model with profitability-based switching. We then estimate both frameworks across fourteen USD currency pairs using long monthly samples that end in 2025.

3 Methodology

3.1 The Underlying Theoretical Framework

The underlying theoretical framework for our empirical analysis builds on Proaño (2011, 2013), who considered the exchange rate determination in a small open economy with perfect capital mobility and boundedly rational expectations.

Agents' expectations and switching mechanism. Let s_t denote the log nominal exchange rate (domestic price of foreign currency), p_t and p_t^* the log domestic and foreign price levels and $\pi_t \equiv p_t - p_{t-1}$ and $\pi_t^* \equiv p_t^* - p_{t-1}^*$ the respective inflation rates. Short-term nominal interest rates in the domestic and foreign economies are i_t and i_t^* respectively.

FX traders are assumed to choose between two forecasting heuristics for the future nominal exchange rate: a fundamentalist rule and a chartist (trend-following) rule. Following Proaño (2011, 2013), their decisions – as well as those of the central bank and the real sector – are based on the information set available up to period $t-1$.

Fundamentalists expect mean reversion of the nominal exchange rate toward its long-run value f, which is assumed to be consistent with the PPP in absolute form, i.e.

$$f_t \equiv p_t - p_t^*.$$

Accordingly, the fundamentalists' expected nominal exchange is given by

$$\begin{aligned} E_t^f[s_{t+1}] &= s_{t-1} + \beta_s^f \left(f_{t-1} - s_{t-1}\right) \\ &= s_{t-1} + \beta_s^f \left(p_{t-1} - p_{t-1}^* - s_{t-1}\right), \\ &= s_{t-1} - \beta_s^f \, \eta_{t-1}, \end{aligned} \tag{1}$$

where β_s^f is a positive parameter measuring the speed of adjustment of the exchange rate to the fundamental value assumed by the fundamentalist traders, and

$$\eta_t \equiv s_t + p_t^* - p_t = s_t - f_t,$$

is the log real exchange rate, which is also a measure of misalignment from the PPP long-run value of the nominal rate.

By contrast, chartists extrapolate last period's change to form their expectation:

$$E_t^c[s_{t+1}] = s_{t-1} + \beta_s^c \, \Delta s_{t-1}, \tag{2}$$

where $\Delta s_{t-1} \equiv s_{t-1} - s_{t-2}$ and β_s^c is a positive parameter measuring the intensity of trend-following behavior of chartists[1].

[1] In Proaño (2011, 2013) both β_s^f and β_s^c are taken to be strictly positive on theoretical grounds: $\beta_s^f > 0$ ensures that the fundamentalist rule pulls the nominal rate toward its PPP value, and $\beta_s^c > 0$ ensures that

Following the discrete-choice framework of Brock and Hommes (1997, 1998), the market share ("market mood") of fundamentalists, $\omega_t \in (0,1)$, depends on the last-period relative profitability of the two rules. Let ψ^j_{t-1} denote the period $t-1$ earnings from rule $j \in \{f,c\}$ and $\sigma^2_{j,t-1}$ its squared forecast error. In Proaño (2011, 2013) the weight is

$$\omega_t = \frac{\exp\left[\gamma\left(\psi^f_{t-1} - \sigma^2_{f,t-1}\right)\right]}{\exp\left[\gamma\left(\psi^f_{t-1} - \sigma^2_{f,t-1}\right)\right] + \exp\left[\gamma\left(\psi^c_{t-1} - \sigma^2_{c,t-1}\right)\right]}, \qquad \gamma > 0, \tag{3}$$

so that $\omega_t \to 1$ when the fundamentalist rule dominates and $\omega_t \to 0$ when technical trading dominates.

Last-period earnings of investing one unit of domestic currency in the foreign currency, ψ^j_{t-1}, depend on the accuracy implied by the rule and the realized excess return:[2]

$$\psi^j_{t-1} = \left[i^*_{t-1} - i_{t-1} + \Delta s_{t-1}\right] \operatorname{sgn}\left[E^j_{t-2}\Delta s_{t-1}\right] \tag{4}$$

where $S_t \equiv e^{s_t}$ is the nominal exchange rate in levels, and $\operatorname{sgn}(E^j_{t-2}\Delta s_{t-1}) \in \{-1,0,1\}$ represents the anticipated appreciation/depreciation suggested by rule j at $t-2$. Each rule captures that profits are positive (negative) when the rule correctly (incorrectly) predicts appreciation/depreciation. Traders choose by comparing the rules' utilities associated with each trading rule defined as

$$U_{f,t-1} = \psi_{f,t-1} - \sigma^2_{f,t-1},$$

$$U_{c,t-1} = \psi_{c,t-1} - \sigma^2_{c,t-1},$$

with σ_j being the squared error of the logarithmic forecast[3] given by rule j:

$$\sigma^2_{j,t-1} = (E^j_{t-2}s_{t-1} - s_{t-1})^2. \tag{5}$$

This formulation mirrors the behavioral heterogeneous-agents literature (de Grauwe and Grimaldi 2006; Hommes 2013).

Further, fundamentalist strategies require access to information and the associated costs of processing it, and may be subject to additional structural frictions (e.g. institutional, regulatory) that systematically shift their relative attractiveness. To capture this, we introduce a wedge parameter k that enters negatively in the fundamentalist utility:

$$U_{f,t-1} = \psi_{f,t-1} - \sigma^2_{f,t-1} - k,$$

---

the chartist rule extrapolates the prevailing trend. In our empirical implementation we do not impose either sign restriction and let the data speak.

[2] Here we deviate from the original formulation of Proaño (2011, 2013) and replace the realized profit component $(S_{t-1}(1+i^*) - S_{t-2}(1+i))$ used there with the foreign-exchange excess return – also interpretable as the UIP deviation or FX risk premium — following Engel (2014).

[3] While Proaño (2011, 2013) used the squared forecast error of the exchange rate level $\left(E^j_{t-2}S_{t-1} - S_{t-1}\right)^2$, this formulation facilitates identification by consistently working in logarithmic terms, thereby avoiding scale inconsistencies across variables, which is an important consideration given the sensitivity of the logistic function to the magnitude of its argument.

so that the market share function takes the form

$$\omega_t = \frac{1}{1 + \exp[-\gamma((\psi_{f,t-1} - \sigma^2_{f,t-1}) - (\psi_{c,t-1} - \sigma^2_{c,t-1}) - k)]}. \tag{6}$$

A positive k is the canonical reading: an information-processing cost that delays fundamentalist takeovers until the profitability advantage is sufficient. We allow k to take either sign in estimation: a negative $\hat{k}$ is then naturally interpreted as a structural tilt toward the fundamentalist rule, capturing institutional or regulatory features that lower its effective cost relative to chartism. This formulation resembles a smooth transition regression (STR) with an estimable threshold, and the threshold has a clear economic interpretation: the utility differential must exceed k before fundamentalist traders become the majority and potentially dominate market dynamics. When $k > 0$ such episodes occur less frequently, so the stabilizing, mean-reverting influence of fundamentalist strategies is correspondingly reduced, thereby amplifying destabilizing dynamics; when $k < 0$ the opposite occurs. Among the specifications considered, this adjustment proved to be the most effective in terms of identification, yielding empirically reasonable results.

Smooth versus discontinuous transition and the interpretation of γ. The intensity of choice γ in Equation (6) controls how sensitive the market composition ω_t is to differences in relative profitability, and admits distinct behavioral readings at its two limits. When γ is finite, the transition between regimes is smooth and the market always hosts a mix of fundamentalist and chartist agents. This region can be interpreted as gradual social learning or heterogeneous switching costs across investors (Frijns et al. 2010). The latter aspect can be partly tested by the introduction of the cost parameter k in the structural model, which, if significantly greater than zero, can confirm the second interpretation. In the opposite limit, as $\gamma \to \infty$, the logistic function collapses into a step indicator and the entire population switches at once as soon as the utility differential crosses the threshold. In this regime, market behavior mimics synchronized herding and informational cascades, with periods of pure chartist or fundamentalist dominance (Brock and Hommes 1998; de Grauwe and Grimaldi 2006; De Grauwe and Rovira Kaltwasser 2012). This interpretation departs partly from the more "psychological" reading, where $\gamma \to \infty$ represents the "fully rational" neoclassical limit in which all agents instantaneously switch to the most profitable rule, and instead emphasizes the "sociological" aspect of switching, where γ captures the speed of social learning and the degree of synchronization in strategy choice across the population. The two limits carry different implications for market efficiency and systemic risk. At finite γ, the market composition adjusts gradually and price reactions to news are smoother, in a way broadly consistent with the stochastic nature of informed trading posited by efficient-market arguments. Discontinuous switching, by contrast, implies abrupt population-wide reallocations that amplify tail risk, shorten the typical distance between booms and crashes, and sustain extended deviations from fundamentals, constituting a stronger violation of the efficient-market hypothesis than simple time-varying risk premia would imply (de Grauwe and Grimaldi 2006; Bolt et al. 2019). Our estimates of γ are therefore informative not only about the speed of regime switching but also about the structural fragility of the currency market in

question: values close to a step transition signal markets exposed to abrupt herding-type reallocations, while lower values suggest a more stable and resilient environment in which heterogeneous traders coexist persistently. This reading is robust to the standard logit's independence of irrelevant alternatives restriction in our setting: Galanis et al. (2025) show that, for two heuristics with homogeneous γ, the probit-based generalization of the discrete-choice aggregator coincides closely with standard logit, both theoretically and in fits to experimental data.

This interpretation is directly operational in our empirical analysis: by estimating γ across currency markets that differ systematically in macroeconomic and institutional stability, the parameter can be read as a market-specific diagnostic of susceptibility to herding-driven reallocations. Our prior is that emerging markets, and particularly those characterised by weaker institutions, thinner liquidity, or persistent macroeconomic volatility, should exhibit estimates of γ closer to the step-transition limit, reflecting a higher propensity for synchronised switching between chartist and fundamentalist strategies. Advanced and stable markets should instead display moderate, finite values of γ, consistent with the persistent coexistence of heterogeneous traders and the more gradual composition dynamics typical of deep and well-anchored financial systems. Comparing $\hat{\gamma}$ across markets with different institutional and macroeconomic profiles therefore offers a first, descriptive assessment of the view that behavioural fragility is shaped by the environment in which a currency trades, rather than being a universal feature of FX markets.

Behavioral UIP and the exchange-rate law of motion. With perfect capital mobility, uncovered interest parity (UIP) holds *ex ante* at the market level, but with behavioral expectations. Define the market expectation as the weight-averaged forecast,

$$E_t^m[s_{t+1}] \equiv \omega_t E_t^f[s_{t+1}] + (1-\omega_t) E_t^c[s_{t+1}]. \tag{7}$$

Thus, the exchange rate updates based on the expectations of traders made at the beginning of t for the following period and on the contemporaneous interest rate differential. The equation thus takes the form of a Behavioral UIP (Proaño 2013):

$$s_t = i_t^* - i_t + E_t^m[s_{t+1}], \tag{8}$$

which, after subtracting s_{t-1} and substituting (1)–(2), yields the behaviorally founded law of motion of the exchange rate:

$$\Delta s_t = (i_t^* - i_t) - \omega_t \, \beta_s^f \, \eta_{t-1} + (1-\omega_t) \, \beta_s^c \, \Delta s_{t-1}. \tag{9}$$

Equation (9) shows that exchange-rate returns depend on the interest rate differential and on two state-dependent behavioral forces: a PPP-correcting term whose strength grows with ω_t and a momentum term whose strength grows with $1-\omega_t$. Because ω_t is a nonlinear function of lagged profitability (3), (9) is a nonlinear difference equation capable of generating persistent swings and regime-like behavior consistent with empirical evidence on nonlinear mean reversion and deviations from parity (Taylor et al. 2001; Sarno et al. 2006; Menkhoff and Taylor 2007).

To connect FX dynamics to macroeconomic stabilization, Proaño (2011, 2013) embeds (9) in a parsimonious real-side block featuring an output gap equation, a

Phillips-Curve relationship and a different Taylor-like interest rate rule. A higher trend-following intensity β_s^c raises persistence in Δs_t and can amplify the real effects of shocks indirectly via η_{t-1} in the IS equation of Proaño (2013). The sensitivity parameter γ controls how sharply the market shifts between regimes; larger γ increases the amplitude of fluctuations in ω_t but does not by itself imply macro instability if policy satisfies the Taylor principle and real-side parameters remain in the stable region (Proaño 2013). These mechanisms echo the broader heterogeneous-agents literature, in which rule switching generates bubbles, crashes, and prolonged deviations from parity while maintaining a fundamentals-based anchor in the long-run (de Grauwe and Grimaldi 2006; Manzan and Westerhoff 2007). Empirically, nonlinear mean reversion in real exchange rates (Taylor et al. 2001; Michael et al. 1997) and the widespread use of technical analysis (Menkhoff and Taylor 2007; Park and Irwin 2007; Cheung and Chinn 2001) motivate the coexistence of rules and the profitability-based switching assumed here. The model is also consistent with the observed weak link between macro news and higher-frequency FX moves (Evans and Lyons 2008, 2002; Evans 2011).

3.2 Scope and limitations of the common-horizon assumption

A separate and more substantive question concerns the empirical plausibility of the common-horizon profitability switching embedded in Equations (3)–(6)[4]. Survey evidence on FX professionals consistently shows that technical strategies dominate at short horizons, from intraday to a few weeks, while fundamental analysis is typically assigned longer evaluation horizons of six months to one year (Cheung and Chinn 2001; Menkhoff and Taylor 2007). Taken literally, this evidence sits in tension with our specification in which both rule types are evaluated monthly against the same realized return, and raises the question of why a fundamentalist with a multi-month horizon should revise her strategy on the basis of a one-period outcome. We address this concern on three levels.

First, we reinterpret the market share ω_t not as the literal fraction of traders who change role in period t, but as the effective weight of fundamentalist strategies in aggregate market pricing. Under this reading, variation in ω_t can arise from (i) marginal reallocation of risk capital across strategies within institutions that pursue hybrid approaches, consistent with evidence that even investors with fundamentalist mandates routinely adjust the weight they assign to technical signals in response to recent profitability (Gehrig and Menkhoff 2005; Dick and Menkhoff 2013; Schmitt 2021; Bassi et al. 2023); (ii) entry and exit of capital toward recently profitable strategies without individual traders changing role; and (iii) asynchronous-updating variants of the heuristic-switching framework in which only a marginal fraction of agents reconsider their strategy at each date, generating slow and persistent shifts in market dominance rather than abrupt population jumps (Boswijk et al. 2007; Franke and Westerhoff 2018; Bolt et al. 2019; Cornea-Madeira et al. 2019). Each of these readings is compatible with the survey evidence that strategy choice reflects persistent institutional and time-horizon heterogeneity rather than high-frequency individual role-changing.

Second, we explicitly acknowledge an important alternative developed in Bassi et al. (2023), who abandon profitability-based switching altogether: the shares of trader

[4] We thank an anonymous referee for raising this point.

types are held fixed and interpreted as a structural feature of market microstructure reflecting institutional segmentation between fund managers and private traders, while fundamentalists themselves weight both fundamental and technical signals at the individual level. Bassi et al. (2023) show that this fixed-share, weighted-signals formulation is sufficient to reproduce the main FX stylized facts and they argue that it aligns more closely with survey evidence than standard Brock–Hommes switching. Our specification differs in that we retain profitability-based switching for continuity with the structural behavioral DSGE literature in which the BUIP is embedded (Proaño 2011, 2013), and because the switching parameter γ carries economic content (as noted in Section 3.1) that we are specifically interested in estimating across currencies with different institutional profiles. The two approaches should be read as complementary rather than competing: Bassi et al. (2023) emphasise the structural and population-level sources of heterogeneity, while our specification captures the aggregate state dependence that such heterogeneity induces in prices, conditional on a parsimonious profitability-based aggregator.

Third, our empirical results can themselves be read as partial support for the concern raised by the referee. The identification difficulties we document for several currencies in Section 4.2.2 indicate that the common-horizon switching mechanism is not equally informative across all markets, and that a "one model fits all" approach is unlikely to be appropriate. The monthly frequency at which we evaluate the utility differential is an intermediate horizon that smooths some of the high-frequency noise most relevant for pure chartism while preserving the responsiveness required to detect fundamental misalignment, but it remains a compromise rather than a micro-founded choice. A natural extension, which we flag in the concluding section, is to allow horizon-asymmetric evaluation in which fundamentalist profitability is assessed over a longer window than chartist profitability, thereby preserving the switching structure while respecting the horizon heterogeneity emphasised by survey evidence.

3.3 Econometric Methodology

Equation (9) fits in the Smooth Transition Regression class of econometric models, first developed by Granger and Teräsvirta (1993b); Teräsvirta (1994), and popularized in modelling nonlinear time series and exchange rate dynamics with the works of Taylor et al. (2001); van Dijk et al. (2002); Granger and Teräsvirta (1993a). These statistical models allow two or more states of the world and enable dynamic switching among them through a transition function. Most of the works in the literature rely on the Logistic (LSTR) or Exponential (ESTR) functions to define the switching behavior. The present work first estimates the LSTR model as defined below to understand the relevance of such a modeling strategy. Then the full behavioral law of motion with the utilities of the two types of traders is estimated via NLS to understand (i) whether model adds something to the "plain" STR specification, and (ii) whether the bilateral exchange rates dynamics in question can be resolved in part through this modelling choice. The reduced form LSTR model is also useful to understand which variables are relevant to define the switching mechanism, providing useful insights about market expectations.

Our aim is to capture the departures of the exchange rate returns from the UIP, so we define the dependent variable y as the difference between the monthly log return and the interest rate differential, i.e., $y = \Delta s_t - (i^* - i)$. This formulation preserves the UIP restriction while attributing short-horizon departures to expectation formation—an interpretation aligned with the microstructure framework (Evans 2011) and consistent with recent evidence supporting long-horizon UIP validity (Baillie et al. 2026). Two regressors capture, respectively, trend-following and fundamental mean reversion trading strategies: the lagged return Δs_{t-1} and the lagged real exchange rate $\eta_{t-1} = s_{t-1} + p^*_{t-1} - p_{t-1}$. Coefficients evolve smoothly between two regimes through the logistic function

$$G(z_{t-d}; \gamma, c) = \Big[1 + \exp\big\{-\gamma\,(z_{t-d} - c)\big\}\Big]^{-1},$$

where $\gamma > 0$ measures the switching intensity and c is the threshold value of the transition variable $z_{c,t-d}$, the driver of the regime-changes. To allow for regime switching, the candidate transition variables are measures of exchange-rate misalignment from fundamentals, all expressed in absolute value so that the switching mechanism depends on magnitude rather than sign and avoids one-sided thresholds, which would have no clear economic interpretation[5].

The candidates are, all in absolute value: (i) $|\eta_t|$, the real exchange rate; (ii) $|\Delta s_t|$, the nominal return; (iii) $|i^*_t - i_t|$, the interest rate differential; (iv) $|f_t|$ the PPP fundamental value; (v) $|\Delta f_t|$ the fundamental value PPP change; (vi) $|\Delta f_t - \Delta s_t|$, the real exchange rate returns; and (vii) $|\Delta s_t - (i^*_t - i_t)|$ the excess returns. Each series is considered up to lag $t-5$, and the transition variable is selected via an LM test for linearity as described below.

The estimated parameter c represents the threshold value beyond which fundamentalist traders dominate market dynamics. Nevertheless, given the smooth nature of the transition function, both types of traders are almost always present, with pure dominance occurring only in the limiting cases where the transition function $G(\cdot)$ approaches 0 or 1. The threshold thus denotes the point beyond which the proportion of fundamentalist traders exceeds 50% of the market.

Stacking these elements yields, for every economy,

$$y_t = \alpha_0 + \beta^c \Delta s_{t-1} + G_t\big(\alpha_1 + \beta^f \eta_{t-1} - \beta^c \Delta s_{t-1}\big) + \varepsilon_t, \tag{10}$$

with α_0 and α_1 the intercepts in the two regimes and $\varepsilon_t \sim$ i.i.d.$(0, \sigma^2)$. This formulation facilitates a comprehensive understanding of the distinctions between the two regimes. Within the chartist regime, where $G = 0$, the model is defined as:

$$y_t = \alpha_0 + \beta^c \Delta s_{t-1}$$

[5] Another possibility would have been estimating a LSTR2 model, with two threshold values c_1, c_2, such that for values inside the band $[c_1, c_2]$ the chartist regime dominates, while for values outside the band the fundamentalist regime takes over. We chose the simpler LSTR specification for parsimony and comparison with the structural BUIP model.

Conversely, in the fundamentalist regime, where $G = 1$, the model transitions to:

$$y_t = \alpha_0 + \alpha_1 + \beta^f \eta_{t-1}.$$

A negative β^f should indicate mean reversion towards the PPP fundamental value, while a positive β^c should capture momentum. The chartist coefficient could capture the destabilizing effect of technical trading if the estimated value is greater than 1.

A more structural approach estimates (9) while building ω_t from the discrete-choice mapping (3). The estimable system is

$$y_t = \alpha_{0,UIP} + \beta_{c,UIP}\Delta s_{t-1} + \omega_t(\alpha_{1,UIP} + \beta_{f,UIP}\eta_{t-1} - \beta_{c,UIP}\Delta s_{t-1}), \tag{11}$$

with ω_t defined as in (6):

$$\omega_t = \frac{1}{1 + \exp[-\gamma(U_{f,t-1} - U_{c,t-1})]}. \tag{12}$$

The structural approach tightens the link with theory and estimates directly endogenous regime switching, while LSTR offers a flexible reduced form that can capture state dependence and be interpreted through the lens of Proaño (2013). The empirical section implements both, beginning with LSTR for parsimony and moving to structural NLS for completeness and robustness.

The econometric estimation routine follows the methodology of Granger and Teräsvirta (1993a), encompassing model specification, estimation, and evaluation.

Specification. The specification phase entails testing the null hypothesis of linearity against the Smooth Transition Regression (STR) alternative. This procedure is implemented via a Lagrange Multiplier (LM) test derived from a Taylor expansion of the transition function around the parameter $\gamma = 0$ (Teräsvirta 1994; van Dijk et al. 2002). The test is conducted for each candidate transition variable and associated lag (up to a five-month horizon). The transition variable yielding the lowest p-value, which indicates a rejection of the null hypothesis at the highest significance level, is selected for the estimation stage. The standard procedure for model type selection (LSTR vs. LSTR2 vs. ESTR) is ignored in this analysis, as the functional form is dictated by the underlying theoretical model.

Estimation. Equations (10) and (11) are fitted by NonLinear Least Squares (NLS). The procedure is slightly different for the two models.

- **STR**: Because the likelihood surface can feature several local maxima, Granger and Teräsvirta (1993a) suggest beginning with a grid search over different combinations of (γ, c) for every candidate transition variable. At each grid point the linear coefficients conditional on (γ, c) are obtained by Ordinary Least Squares (OLS) and the best estimation is selected by minimizing the residual sum of squares (RSS). The resulting $(\hat{\gamma}, \hat{c})$ are used as starting values for the NLS regression, performed via the Trust-Region optimization algorithm (a variant of the Levenberg-Marquardt algorithm that allows to set bounds), that updates all parameters jointly, subject to the positivity constraint on γ. Convergence tolerances on parameters and objective are

set to 10^{-8}, and the iteration cap is fixed at 20000. Identification of γ is notoriously difficult as noted by (Granger and Teräsvirta 1993a), but a grid search over γ can help reach a local maximum. Adhering to the methodology proposed by Granger and Teräsvirta (1993a), γ is normalized by the standard deviation of the transition variable. This procedure reduces scale sensitivity and enhances the identification properties of the parameter.

- **BUIP**: The estimation strategy parallels that of the STR model; however, the conditional ordinary least squares step cannot be implemented due to the functional specification of the market share ω_t. Consequently, we initiate the procedure by fitting a linear benchmark specification ($y_t = const + \beta_c \Delta s_{t-1} + \beta_f \eta_{t-1}$) through ordinary least squares. The estimated coefficients serve to construct approximate utilities $U_{f,t}$ and $U_{c,t}$, under the assumption that the cost parameter equals zero ($k = 0$). A grid of candidate starting values for the parameter pair (γ, k) is constructed based on quantiles of the utility differential $(U_{f,t} - U_{c,t})$. For each grid point, the corresponding market share ω_t is computed. Given these market share proxies, the model becomes linear in the remaining parameters conditional on (γ, k), enabling estimation via ordinary least squares. These preliminary parameter estimates subsequently serve as initial values for the final nonlinear least squares optimization.

Reported in the tables are the confidence intervals and significance levels based on heteroskedasticity and autocorrelation consistent (HAC) standard errors computed via the Newey-West estimator with a lag length selected according to the Newey and West (1994) rule.

Evaluation. The first step in the evaluation phase involves checking whether the estimated threshold parameter ($\hat{c}$ in the STR and $\hat{k}$ in the BUIP model) lies within the empirical support of the transition variable and is statistically significant. This provides a first indication of correct identification of nonlinear models. Countries for which the estimated threshold lies outside that support are treated as having a non-identified switching point: their estimates are reported but interpreted with caution, and where the non-identification appears informative about market structure (as for Japan, Korea, and the United Kingdom in the BUIP case) it is itself the object of discussion. Non-identification of the threshold does not by itself disqualify the identification of the remaining coefficients, which we report and discuss whenever they are precisely estimated. Residual diagnostics include Ljung-Box tests for autocorrelation, ARCH Lagrange-multiplier tests for conditional heteroskedasticity, and Jarque-Bera tests for normality.

The next section presents the results of the empirical analysis.

4 Results

4.1 Data Description

We employ data obtained from the IMF International Financial Statistics (IFS) and OECD Main Economic Indicators. The dataset comprises monthly observations for fourteen economies: Australia (AU), Brazil (BRA), Canada (CAN), the Euro Area

(EA), Indonesia (ID), Japan (JPN), Korea (KOR), Mexico (MEX), New Zealand (NZ), the Philippines (PH), Switzerland (CH), Thailand (TH), Türkiye (TR), and the United Kingdom (GBR). We exploit the maximum available sample period for each economy, resulting in an unbalanced panel spanning from as early as the 1980s to the latest data available in 2025. Appendix A.1 provides a detailed overview of the sample periods for each country.

Table 1: Summary statistics for monthly log returns $\Delta s_t = s_t - s_{t-1}$ and log PPP fundamental $f_t = \ln(P_t^{\text{DOM}}/P_t^{\text{US}})$. Moments computed over country-specific samples (for data availability and sources see Table 10).

Country	$Mean_{\Delta s}$	$Var_{\Delta s}$	$Skew_{\Delta s}$	$Kurt_{\Delta s}$	$Mean_f$	Var_f
Australia	0.0008	0.0010	0.7045	2.5387	0.3314	0.1274
Brazil	0.0488	0.0104	1.6606	3.1150	3.3373	0.8037
Canada	0.0002	0.0004	0.4249	3.9327	-0.0051	0.1346
Euro Area	-0.0000	0.0007	0.1595	1.2867	-0.0919	0.2374
Indonesia	0.0057	0.0037	5.3678	69.4273	1.3997	0.1515
Japan	-0.0010	0.0009	-0.2226	1.3990	-2.0937	0.4672
Korea	0.0012	0.0009	0.7188	12.5920	0.5421	0.3243
Mexico	0.0116	0.0024	6.0905	68.4200	1.8292	0.6185
New Zealand	-0.0005	0.0012	0.3714	1.2110	-0.1009	0.0610
Philippines	0.0038	0.0008	3.5908	25.9394	1.0881	0.2147
Switzerland	-0.0019	0.0010	-0.1025	0.9630	-0.3797	0.4980
Thailand	0.0007	0.0008	0.2737	27.8384	0.2695	0.1638
Türkiye	0.0234	0.0025	2.4488	15.0845	2.5129	0.2761
United Kingdom	0.0005	0.0008	0.1674	2.4225	0.0776	0.0722

Three series are required per country: the nominal spot rate S_t expressed as the domestic price of USDs, the consumer-price index P_t, and the short-term interest rate i_t (which is almost equal to the monetary policy rate). US variables complete the construction of each country's dataset for real exchange rates and interest differentials construction. All variables – except for the interest rate – are used in natural logarithms ($s_t = \ln S_t$, $p_t = \ln P_t^{\text{DOM}}$, $p_t^* = \ln P_t^{\text{US}}$). Monthly log returns are computed as the differences of the log nominal exchange rate ($\Delta s_t = s_t - s_{t-1}$). Interest rates are converted from annual percentage rates to effective monthly decimals via $(1 + \text{APR})^{1/12} - 1$.

Table 1 shows marked heterogeneity across currencies: emerging markets (e.g., Brazil, Mexico, Indonesia) exhibit higher means, variances, and extreme skewness/kurtosis, while advanced economies display near-zero means, low dispersion, and thinner tails; the PPP fundamental f_t is modestly positive in mean for most emerging markets—consistent with the depreciation trend of emerging-markets currencies—and near zero or negative for advanced peers.

Figure 1 displays the log nominal exchange rate s_t (solid black line) and the PPP long-run fundamental value f_t (dashed gray line) for selected economies over the sample period. To visualize the deviations from the PPP long-run value, the variables

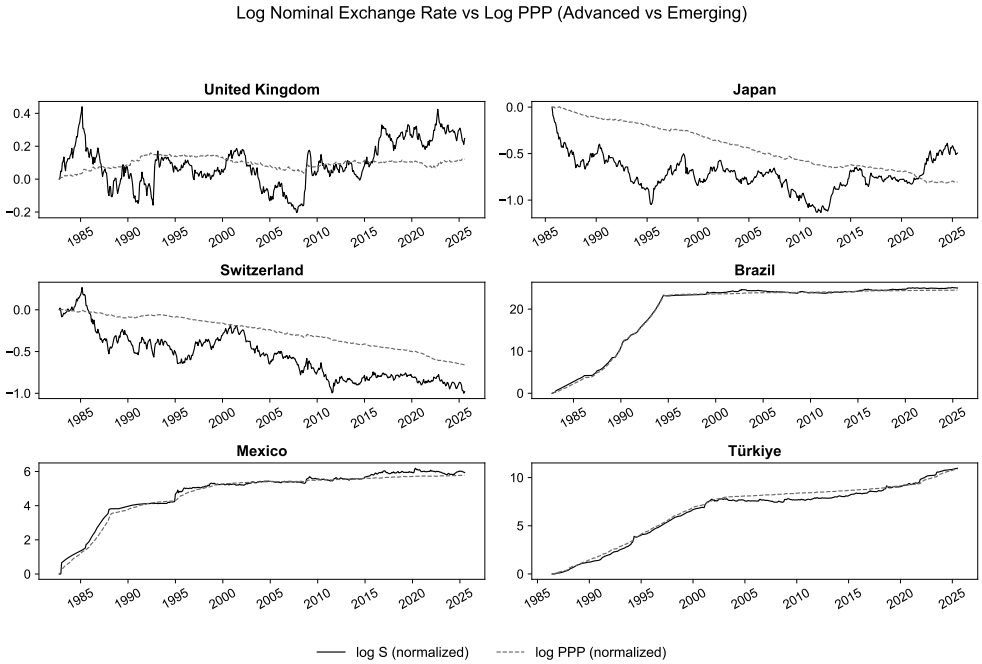


Fig. 1: Nominal exchange rate s_t (solid line) and long-run PPP fundamental f_t (dashed line) for selected economies.

have been normalized to the first observation, thus assuming that PPP holds at that date. The vertical distance between the nominal exchange rate and the PPP fundamental measures the misalignment from purchasing power parity. The presence of the Disconnect Puzzle (Obstfeld and Rogoff 2000) in each economy is immediately clear, where each series shows large deviations from the absolute PPP fundamental, that persist for long periods—sometimes more than 15 years—with short run movements always present. The difference between emerging and advanced economies is striking. The emerging countries show a historical trend of depreciation against the USD, while the advanced countries fluctuate around an almost stable mean. The deviations from the PPP long-run value are present in all countries.

Plots of the exchange-rate returns Δs_t and the change in the PPP fundamental Δf_t (not shown) confirm the higher volatility of the nominal exchange rate compared to the PPP changes (Engel and West 2005). The emerging markets show higher volatility levels, as confirmed by the results in Table 1.

4.2 Estimation Results

4.2.1 Empirical Results of the STR Estimation

The STR estimations provide evidence that exchange-rate dynamics can display state-dependent behavior, consistent with the nonlinear specification introduced in

Section 3. As a preliminary step, we test the null of linearity against the STR alternative and select, for each country, the transition variable that yields the strongest evidence of nonlinearity. Table 2 reports the selected transition variable and the corresponding LM linearity-test p-value. Linearity is rejected at conventional levels for 13 out of 14 economies. Australia is the only case in which the linearity null cannot be rejected (p-value=0.2426), suggesting that an STR-type regime mechanism is not empirically supported for the AUS/USD currency. This is confirmed by the identification results reported in Table 11 in Appendix A.2, where the estimated threshold lies outside the range of the transition variable, indicating a failure to identify the model. Along with Australia, also Mexico, Thailand, Türkiye and the UK show $\hat{c}$ estimates outside the range of the transition variable. While linearity is still rejected for these countries, this result suggests that the identification problem resides in the estimation algorithm. While Australia results won't be discussed further, these other countries' results are reported for completeness, but will be treated with caution.

Beyond the widespread rejection of linearity, the selected transition variables indicate heterogeneity in the variable that govern regime changes. For a large subset of countries (Brazil, Canada, Indonesia, Japan, Korea, Mexico), the preferred transition variable is a measure of (absolute) relative misalignment, consistent with switching driven by the magnitude of deviations from fundamentals. In other cases (Switzerland, Thailand, Türkiye, UK) the transition variable is the fundamental value change in absolute value. This cross-country heterogeneity is consistent with the view that nonlinear switching is a robust feature, driven mostly by fundamental misalignment, but that the precise mechanism varies across currency markets.

Table 3 reports the STR coefficient estimates (with 95% confidence intervals). Positive and significant β_c estimates are obtained for Brazil, Canada, and Japan only, indicating that the reduced form model fails to identify momentum. New Zealand instead features a significantly negative β_c, consistent with contrarian short-horizon dynamics within the sample under analysis. These results suggest that either the model specification is too simple to capture technical trading dynamics in FX markets at a monthly horizon, or that it is not a dominant force in the countries under consideration.

Turning to the mean-reversion component, β_f, the estimates are more selective. A negative β_f implies mean reversion towards the fundamental PPP value. Statistically significant mean reversion is found for Brazil, Korea, New Zealand, and Switzerland, while the Euro Area displays weaker evidence (significant at the 10% level). For Mexico, the estimated threshold $\hat{c}$ takes a negative value that is inconsistent with the non-negative support of the selected transition variable, so the STR fails to separate the data into distinct regimes and the regime-specific β_f cannot be read as a structural fundamentalist response. We treat this case as a negative result: rather than informing on the strength of PPP-based mean reversion, it indicates that the reduced-form STR with an absolute-PPP anchor is not well suited to the MXN/USD dynamics. This motivates either different reduced-form specifications—e.g. different fundamentals, more complex structures—or estimating the full structural behavioral UIP model, as done in Section 4.2.2.

The switching parameter provides additional insights but also highlight identification constraints that are common in STR applications. For the countries in which

Table 2: Selected transition variables and LM linearity test results for the STR model (Eq. (10)).

Country	Transition Variable	P-value
Australia	$\lvert\Delta f_{t-3} - \Delta s_{t-3}\rvert$	0.2426
Brazil	$\lvert\Delta f_{t-2} - \Delta s_{t-2}\rvert$	0.0000
Canada	$\lvert\Delta f_{t-2} - \Delta s_{t-2}\rvert$	0.0000
Euro Area	$\lvert\Delta s_{t-3}\rvert$	0.0112
Indonesia	$\lvert\Delta f_{t-4} - \Delta s_{t-4}\rvert$	0.0000
Japan	$\lvert\Delta f_{t-2} - \Delta s_{t-2}\rvert$	0.0071
Korea	$\lvert\Delta f_{t-4} - \Delta s_{t-4}\rvert$	0.0000
Mexico	$\lvert\Delta f_{t-4} - \Delta s_{t-4}\rvert$	0.0000
New Zealand	$\lvert(i^*_{t-4} - i_{t-4})\rvert$	0.0059
Philippines	$\lvert\Delta s_{t-3}\rvert$	0.0173
Switzerland	$\lvert\Delta f_{t-2}\rvert$	0.0000
Thailand	$\lvert\Delta f_{t-5}\rvert$	0.0000
Türkiye	$\lvert\Delta f_{t-4}\rvert$	0.0000
United Kingdom	$\lvert\Delta f_{t-3}\rvert$	0.0151

Note: The selected variable $z_{c,t-d}$ enters in absolute value; candidate variables are misalignment proxies and rate differentials. The p-value (last column) is from the Teräsvirta (1994) LM test; small values reject linearity in favour of the logistic STR.

the switching parameter is precisely estimated, the implied transition is economically plausible. Brazil and Korea display moderate γ estimates, consistent with gradual regime switching, as confirmed in Figure 2. For several other countries, γ is estimated at (or near) large values, indicating a threshold switching rather than a smooth one[6]. These results suggest that the hypothesis that advanced markets display smoother regime switching than emerging ones is not supported by the data. Before drawing conclusions on the matter, we evaluate and compare the BUIP model results.

The threshold parameter c is precisely estimated in many cases (e.g., Brazil, Canada, Euro Area, Indonesia, New Zealand, Philippines, Switzerland, Korea). Brazil, Indonesia and Korea estimates indicate that fundamental expectations start dominating the market when the real exchange rate returns—which measures the misalignment of the returns and the PPP change—get to 7-10%. However, Indonesia's β_f is imprecisely estimated. This result is consistent with the narrative that in emerging economies fundamentalists traders take more time to intervene and markets are more unstable, due to destabilizing forces of strong technical trading.

Residual diagnostics for the STR models are reported in Table 4. As typical for high-frequency FX returns, the Jarque–Bera test strongly rejects normality across all countries. The Ljung–Box results suggest that for several advanced economies (e.g., Australia, Euro Area, Switzerland, United Kingdom) serial correlation is not present,

[6] As stressed in Granger and Teräsvirta (1993a), the t-value of the gamma estimate does not have the usual interpretation. Large values suggest a TR (Threshold Regression) model where the transition function is discontinuous, and the model jumps from one regime to the other rather than switching gradually as with the STR.

Table 3: STR model estimates (Eq. (10)) with 95% Newey–West HAC confidence intervals. α_0: intercept; α_1: regime-2 intercept; β_c: chartist coefficient; β_f: fundamentalist (PPP mean-reversion) coefficient; γ: transition speed; c: threshold. Transition variables as in Table 2.

Country	α_0	α_1	β_c	β_f	γ	c
Australia	15.1129 [—]	-15.1073 [—]	294.2944 [—]	-0.0179 (-0.1297, 0.0938)	0.4071 (-1.6624, 2.4766)	-0.4191 [—]
Brazil	0.0139 (-0.0147, 0.0425)	**0.3476***** (0.1059, 0.5893)	**1.9925***** (1.3599, 2.6251)	**-0.2821**** (-0.5168, -0.0474)	**0.9070**** (0.1987, 1.6154)	**0.1037***** (0.0385, 0.1689)
Canada	-0.0002 (-0.0026, 0.0022)	0.0056 (-0.0011, 0.0123)	**0.1816**** (0.0273, 0.3359)	-0.0187* (-0.0402, 0.0028)	256.0000 [—]	**0.0113***** (0.0102, 0.0124)
Euro Area	-0.0059 (-0.0135, 0.0016)	0.0029 (-0.0058, 0.0116)	0.0840 (-0.1256, 0.2937)	-0.0224* (-0.0469, 0.0021)	256.0000 [—]	**0.0036***** (0.0023, 0.0049)
Indonesia	**0.0091***** (0.0064, 0.0118)	0.7043 (-0.5728, 1.9814)	0.4063 (-0.1821, 0.9948)	-0.0733 (-0.2074, 0.0608)	198.1245 [—]	**0.0722***** (0.0578, 0.0865)
Japan	-0.0020 (-0.0052, 0.0012)	-0.0408 (-0.1111, 0.0295)	**0.2274***** (0.1115, 0.3433)	0.0086 (-0.0068, 0.0239)	512.0001 [—]	0.0279 (-0.2503, 0.3061)
Korea	**0.0035**** (0.0007, 0.0063)	**3.0003***** (1.2617, 4.7389)	0.2684 (-0.0952, 0.6320)	**-0.4215***** (-0.6598, -0.1831)	1.2871 (-0.5169, 3.0912)	**0.0822**** (0.0129, 0.1514)
Mexico	72.8714 [—]	-73.0846 [—]	751.3892 [—]	**0.0814**** (0.0187, 0.1440)	0.9064 (-0.3385, 2.1512)	-0.3324 [—]
New Zealand	0.0005 (-0.0029, 0.0040)	**0.0351***** (0.0189, 0.0513)	**-0.1446**** (-0.2557, -0.0336)	**-0.0659***** (-0.0984, -0.0334)	256.0000 [—]	**0.0023***** (0.0023, 0.0024)
Philippines	**0.0295***** (0.0084, 0.0506)	0.0023 (-0.0471, 0.0517)	-0.0288 (-0.1729, 0.1154)	-0.0066 (-0.0185, 0.0052)	5544.9990 [—]	**0.0004***** (0.0003, 0.0004)
Switzerland	**-0.0038***** (-0.0065, -0.0012)	**0.0348***** (0.0119, 0.0576)	0.0769* (-0.0057, 0.1594)	**-0.2274***** (-0.3809, -0.0740)	63.8486 [—]	**0.0078***** (0.0077, 0.0079)
Thailand	-10.6922 [—]	23.9689 [—]	0.3122 [—]	-0.0660 [—]	0.0002 (-4.7013, 4.7018)	5.0255 [—]
Türkiye	-153.8334 [—]	153.9550 [—]	N.I.	-0.0009 (-0.0207, 0.0190)	0.3841 (-0.0750, 0.8433)	N.I.
United Kingdom	0.0092 (-0.0136, 0.0320)	-436.1348 [—]	0.0808 (-0.0767, 0.2382)	-527.8929 [—]	0.2400 (-0.5780, 1.0579)	0.1272 [—]

Note: 95% CIs in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. N.I. = not identified (non-convergence). [—] = CI suppressed for readability (non-informative due to optimizer divergence).

while most emerging markets still exhibits residual autocorrelation (notably Brazil, Indonesia, Korea, Mexico, Philippines, Thailand). ARCH effects remain widespread, indicating that the STR specification captures conditional mean nonlinearities but does not model conditional heteroskedasticity; an STR–GARCH extension is therefore a natural next step.

Table 4: STR residual diagnostics. Test statistic with p-value in parentheses.

Country	Ljung-Box Q	ARCH LM	Jarque-Bera
Australia	8.2801 (0.6015)	**14.9009** (0.0049)	**196.5628** (0.0000)
Brazil	**59.2758** (0.0000)	**39.6404** (0.0000)	**2545.6774** (0.0000)
Canada	**23.4180** (0.0093)	**15.2635** (0.0042)	**276.5862** (0.0000)
Euro Area	8.3538 (0.5943)	**48.8331** (0.0000)	**18.8333** (0.0001)
Indonesia	**44.2145** (0.0000)	**50.3805** (0.0000)	**73729.2951** (0.0000)
Japan	**20.1975** (0.0274)	5.3142 (0.2566)	**28.9678** (0.0000)
Korea	**38.7903** (0.0000)	**162.5248** (0.0000)	**1867.5550** (0.0000)
Mexico	**184.5590** (0.0000)	7.4513 (0.1139)	**11919.5325** (0.0000)
New Zealand	**19.1951** (0.0379)	**31.3239** (0.0000)	**73.7937** (0.0000)
Philippines	**25.1323** (0.0051)	8.9700 (0.0619)	**9863.3155** (0.0000)
Switzerland	10.3420 (0.4110)	6.0242 (0.1973)	**18.0328** (0.0001)
Thailand	**19.3667** (0.0358)	**52.8779** (0.0000)	**15587.9390** (0.0000)
Türkiye	14.5205 (0.1505)	**14.0027** (0.0073)	**8272.7336** (0.0000)
United Kingdom	4.6851 (0.9112)	**15.0545** (0.0046)	**90.7274** (0.0000)

Note: Ljung-Box Q uses 10 lags (null: no serial autocorrelation). ARCH LM uses 4 lags (null: no conditional heteroskedasticity). Jarque-Bera tests for normality of residuals. Bold entries indicate rejection of the null at the 5% level.

4.2.2 Estimation of the Behavioral UIP Model

Table 5 reports the estimates of the structural BUIP model of Proaño (2011, 2013) defined in Equations (11) and (12). Relative to the reduced-form STR, this specification imposes an explicit evolutionary switching mechanism through the utility differential. The structural model also improves interpretation of the results: coefficients can be read as the strength of chartist versus fundamentalist forces, while the switching parameters describe how strongly relative profitability translates into market-share dynamics, measuring the degree of social learning and herding behavior (or, with a Keynesian interpretation, "animal spirits" (De Grauwe and Rovira Kaltwasser 2012)).

We first note, as summarized in Table 11, that the estimated k cost parameter lies outside the empirical support of the utility differential for Japan, Korea and the United Kingdom, suggesting non-identification of the switching point for these countries. These three cases are discussed in more detail later, and their estimates are reported with caution rather than excluded from the discussion. For the remaining countries, $\hat{k}$ admits the structural reading developed in Section 3.1: a positive and significant $\hat{k}$ for Brazil, Canada, the Euro Area and Türkiye is consistent with a relatively

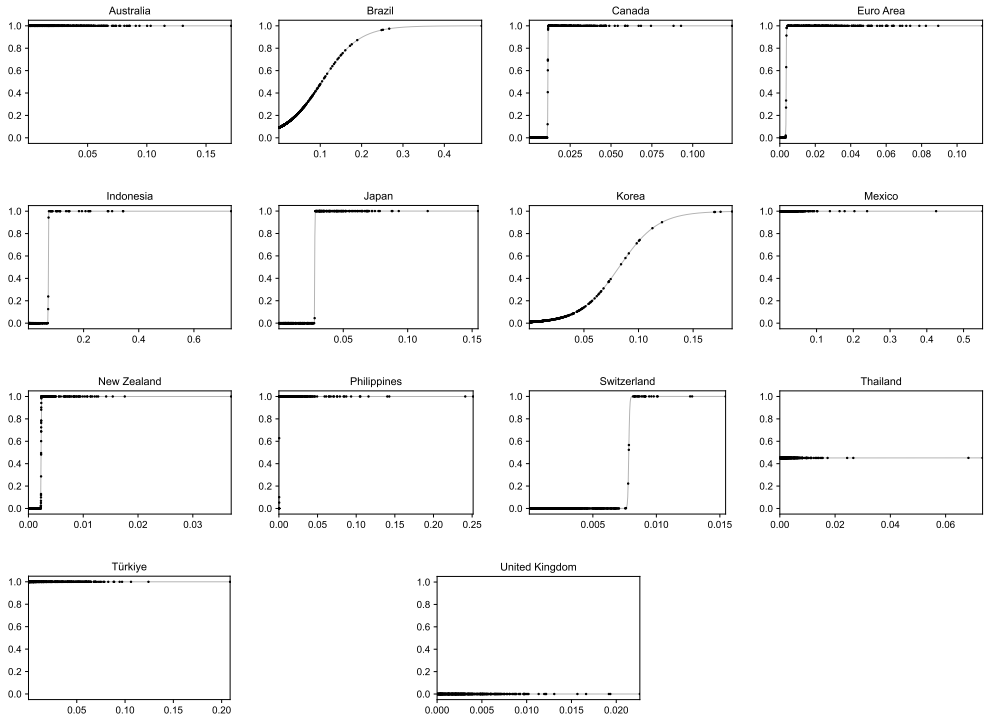


Fig. 2: STR model: scatter plots of transition function G_t (y-axis) against transition variable z_t (x-axis).

costly fundamentalist rule and a delayed fundamentalist takeover, while a negative and significant $\hat{k}$ for Indonesia, New Zealand and the Philippines is consistent with a structural tilt toward fundamentalism. $\hat{k}$ is imprecisely estimated for Australia, Mexico, Switzerland and Thailand. Both signs appear in advanced and in emerging economies, so the sign of $\hat{k}$ does not map onto the advanced-vs-emerging split in any clear way.

β_f is negative and significant in seven currencies (Australia, Euro Area, Indonesia, New Zealand, Philippines, Switzerland, Thailand), consistent with stabilizing mean reversion to the PPP anchor when fundamentalists dominate. Korea yields a precisely estimated negative $\hat{\beta}_f = -0.035$ ($p < 0.01$), but we exclude it because $\hat{k}$ falls outside the empirical support of the utility differential; without an identified switching point the coefficient cannot be read as regime-specific (we return to this case in Section 4.3). Elsewhere, $\hat{\beta}_f$ is insignificant for Brazil and Canada (both identified) and uninterpretable for Japan and the United Kingdom, where the model itself is not identified. Mexico and Türkiye display a positive and significant β_f: the "fundamentalist" term destabilizes rather than corrects. A plausible reading is that in economies with large inflation differentials or recurrent capital-account stress, the absolute-PPP anchor is dominated by interest-rate- or trade-balance-based fundamentals (see Section 5). The empirical relevance of a given anchor is thus market-dependent.

Table 5: BUIP model estimates (Eqs. (11)–(12)), NLS with 95% Newey–West HAC confidence intervals. α_0: intercept; α_1: regime-2 intercept; β_c: chartist coefficient; β_f: fundamentalist (PPP mean-reversion) coefficient; γ: Brock–Hommes switching intensity; k: cost/threshold parameter.

Country	α_0	α_1	β_c	β_f	γ	k
Australia	0.0455	-0.0382	0.0123	**-0.0159****	812.7878	-0.1377
	(-0.3391, 0.4302)	(-0.4220, 0.3457)	(-2.7092, 2.7339)	(-0.0316, -0.0001)	[—]	(-5.8102, 5.5347)
Brazil	**0.0460*****	-0.0698	**1.7350*****	0.0063	**2.3011****	**0.0998****
	(0.0301, 0.0618)	(-0.1987, 0.0592)	(1.5062, 1.9638)	(-0.0312, 0.0439)	(0.3223, 4.2798)	(0.0152, 0.1844)
Canada	0.0019*	-0.0284	**-0.1349****	-0.0297	2.3282	**0.0725*****
	(-0.0002, 0.0040)	(-0.0735, 0.0168)	(-0.2392, -0.0307)	(-0.1852, 0.1259)	(-0.4906, 5.1469)	(0.0400, 0.1049)
Euro Area	0.0012	**-0.0266*****	**0.0053****	**-0.0935*****	206.7932	**0.0278*****
	(-0.0022, 0.0045)	(-0.0360, -0.0172)	(0.0006, 0.0100)	(-0.1427, -0.0444)	[—]	(0.0155, 0.0401)
Indonesia	**1.8482*****	**-1.4450*****	**-4.3036*****	**-0.0429*****	**2.2241****	**-0.5137*****
	(1.6995, 1.9968)	(-1.6514, -1.2386)	(-4.8034, -3.8038)	(-0.0674, -0.0184)	(0.4633, 3.9848)	(-0.8090, -0.2184)
Japan	1.0760	-1.0686	56.7780	-0.0027	0.0164	-14.2687
	(-6.9135, 9.0655)	(-9.0535, 6.9162)	[—]	(-0.0136, 0.0082)	(-0.2089, 0.2417)	[—]
Korea	N.I.	N.I.	1.9776	**-0.0352*****	7.9629	N.I.
			[—]	(-0.0544, -0.0161)	[—]	
Mexico	**0.0234*****	**-0.0761*****	**0.7313*****	**0.0259*****	299.4982	0.0041
	(0.0179, 0.0290)	(-0.1279, -0.0243)	(0.4473, 1.0153)	(0.0070, 0.0449)	[—]	(-0.0019, 0.0101)
New Zealand	**0.0100*****	0.0100	**-0.2628*****	**-0.0393*****	619.0809	**-0.0014*****
	(0.0026, 0.0175)	(-0.0026, 0.0225)	(-0.3855, -0.1402)	(-0.0611, -0.0174)	[—]	(-0.0023, -0.0004)
Philippines	**0.0081*****	**0.6574*****	**0.0308*****	**-0.1746*****	33.2207	**-0.3544*****
	(0.0050, 0.0112)	(0.6124, 0.7023)	(0.0180, 0.0436)	(-0.1894, -0.1599)	[—]	(-0.4267, -0.2821)
Switzerland	-0.0022	0.0075	0.0240	**-0.0817*****	9.7171	0.0021
	(-0.0062, 0.0018)	(-0.0021, 0.0171)	(-0.0411, 0.0891)	(-0.1255, -0.0379)	(-14.1794, 33.6136)	(-0.0090, 0.0131)
Thailand	0.1490	-0.0328	-0.0044	**-0.0343*****	0.7657	-0.2108
	(-0.4401, 0.7380)	(-0.6886, 0.6230)	(-0.0324, 0.0235)	(-0.0547, -0.0140)	(-1.8699, 3.4013)	(-0.8474, 0.4257)
Türkiye	**0.0400*****	**-0.0505*****	**0.5654*****	**0.0361*****	1553.3028	**0.0016*****
	(0.0332, 0.0468)	(-0.0708, -0.0302)	(0.3731, 0.7578)	(0.0177, 0.0544)	[—]	(0.0009, 0.0023)
United Kingdom	0.0009	2.5795	**0.0006*****	-0.0595	3.8513	0.1443
	(-0.0015, 0.0032)	[—]	(0.0002, 0.0010)	(-0.4613, 0.3422)	(-2.1357, 9.8383)	(-0.8335, 1.1220)

Note: 95% CIs in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. N.I. = not identified (non-convergence). [—] = CI suppressed for readability (non-informative due to optimizer divergence).

The chartist coefficient β_c is precisely estimated in many cases and displays substantial cross-country variation in sign. Positive and significant trend-following force is found for Brazil, Euro Area, Mexico, Philippines and Türkiye, while negative and significant estimates characterise Canada, Indonesia, and New Zealand. Although the original specification of Proaño (2011, 2013) defines the chartist rule as trend extrapolation in eq. (2), the estimated $\hat{\beta}_c$ summarises the empirical autoregression in Δs_t once interest differentials and PPP misalignment are accounted for. A negative $\hat{\beta}_c$ should therefore be read as evidence that the short-horizon non-fundamental component of returns is not well represented by simple trend-following in those markets. From a behavioral perspective, a negative $\hat{\beta}_c$ suggests contrarian expectations (Hommes 2013). Cross-country variation in market structure and sample composition plausibly contributes to this heterogeneity.

It is also informative to translate the estimated coefficients into economic magnitudes, beyond their statistical significance, following the interpretation of Proaño (2013). The fundamentalist rule expects PPP deviations to be corrected at a monthly rate of $|\hat{\beta}_f|$, so that $1/|\hat{\beta}_f|$ is the implied correction horizon, ceteris paribus, when fundamentalists dominate the market. Among identified BUIP specifications, the implied horizons range from approximately 6 months for the Philippines ($\hat{\beta}_f = -0.175$) and 11 months for the Euro Area (-0.094) to about 23 months for Indonesia (-0.043), 29 months for Thailand (-0.034), and 63 months for Australia (-0.016), with Switzerland (-0.082) and New Zealand (-0.039) at 12 and 26 months respectively. The chartist coefficient $\hat{\beta}_c$, by analogy with Proaño (2013), indicates the expected scaling of the most recent return: values above one (e.g. $\hat{\beta}_c = 1.74$ for Brazil) imply that chartists expect overshooting of the prevailing trend, while values below one in absolute terms (e.g. 0.57 for Türkiye and 0.03 for the Philippines) imply mild trend extrapolation, and significantly negative estimates (e.g. -0.13 for Canada, -0.26 for New Zealand, -4.30 for Indonesia) point to a contrarian short-horizon component. Together, these magnitudes indicate that the implied fundamentalist correction horizons span roughly half a year to five years across identified currencies, while the chartist component is dominated by overshooting only in a handful of high-volatility emerging markets. We stress that these are simple readings of the static coefficients in the spirit of Proaño (2013), not properly state-dependent half-lives in the sense of Taylor et al. (2001); the latter would require simulating the impulse response of the nonlinear BUIP and is left to future work.

The switching intensity γ is statistically significant and economically plausible for Brazil and Indonesia. For other countries, γ is often estimated at large values with wide confidence intervals; as stressed in Granger and Teräsvirta (1993a), in such cases the usual t-statistic interpretation does not apply and large $\hat{\gamma}$ should be read as a step (threshold) transition rather than as a failure of identification, which we instead diagnose through the empirical support of $\hat{k}$ in the BUIP (and $\hat{c}$ in the STR). Figure 3 presents the scatter plots of the transition function against the utility differential. The plots show a smoother transition function for Brazil, Canada and Switzerland, and a more nearly threshold-like function for the remaining countries. The prior that

advanced markets display smoother regime switching than emerging ones is not supported by the data, and we defer the cross-country synthesis of the switching pattern to Section 4.2.4.

Possible sources of non-identification. As already noted, the cost/threshold parameter k is identified for all countries except Japan, Korea, and the United Kingdom (see Table 11), a finding corroborated by Figure 3, where all three display an invariant transition function. Two sources of non-identification are plausible: non-convergence of the optimisation algorithm, which we regard as the less plausible given the robustness of our estimation procedure, and country-specific structural features that hinder the identification for such a parsimonious economic model. The three cases are qualitatively distinct[7], and each points to the latter explanation.

For Japan, the mechanism emphasised by Engel et al. (2022) seems to apply: the predictive content of the interest-rate differential for exchange-rate returns weakens near the zero lower bound, thus in our model compressing the utility differential that drives ω_t. Japan is, however, a distinct case from other sample economies that experienced near-zero rates. The Bank of Japan has operated at or below 0.5% continuously since 1995 and below 0.1% since 2001, while yield-curve control (2016–2024) explicitly targeted the ten-year JGB yield, removing identifying variation at both ends of the curve for roughly a decade (Shiratsuka 2025). The point estimate $\hat{\gamma} = 0.02$, implying an essentially invariant ω_t, is consistent with this reading. That Canada, the Euro Area, Switzerland, and New Zealand retain BUIP identification despite shorter or less pervasive ZLB and consequent zero or negative interest rate policies (ZIRP/NIRP) episodes suggests that it is the depth and duration of the unconventional regime, rather than the mere presence of the lower bound, that is operative in the Japanese case. This reading receives partial support from the pre-COVID robustness exercise (Section 4.3): restricting the sample to the pre-2020 period yields $\hat{k}$ in the range of the utility differential, though most parameters remain imprecise, suggesting that the post-2020 period might have exacerbated the identification issues, but is not the sole driver of them.

For Korea, the ZLB reading does not apply: the Bank of Korea's minimum policy rate over our sample was 0.50% (2020–2021). A more plausible reading combines two structural features of the Korean FX market. First, Hur and Shin (2026) document that UIP, though frequently violated in Korea before the Global Financial Crisis, holds increasingly since, suggesting that exchange rates showed greater responsiveness to interest rate shocks. Thus, narrower UIP deviations leave less identifying variation for the switching mechanism, which relies on systematic departure from UIP. Second, the Bank of Korea operates a managed floating regime with de facto moving bands (Ilzetzki et al. 2017) and introduced a macroprudential toolkit in 2010: leverage caps on banks' FX forward positions, and a macroprudential stability levy on non-core FX liabilities (Bruno and Shin 2014; Ahn et al. 2022). Both features act directly on the carry-trade channel that drives our BUIP utility differential. Consistent with this, the fundamentalist coefficient $\hat{\beta}_f = -0.035$ (significant at 1%) survives BUIP estimation: what fails is identification of the profitability-based switching, not the underlying

[7] The following discussion, and the comparative reading of the different results relies also on the classifications of exchange rate regimes of (Ilzetzki et al. 2017), and of monetary policy frameworks of (Cobham 2024). For the latter, the data can be found at https://monetaryframeworks.org/.

mean-reversion force. Again, the pre-COVID robustness exercise (Section 4.3) provides a partial test of this reading: under the restricted sample, the model is identified and yields significant parameter estimates (except for γ). This robustness check suggests the plausibility of our interpretations of non-identification, given the Bank of Korea's interventions post-2020, especially the FX interventions during the second half of 2022 (Rhee 2026).

For the United Kingdom, the non-identification (with both STR and BUIP models, the only country) is likely driven primarily by structural breaks within the flexible-inflation-targeting (FIT) era.[8] Three overlapping shocks disrupt the parameter stability required by a constant-coefficient (constant γ) switching model: the Brexit referendum of 23 June 2016, which triggered the largest one-day depreciation of any major currency since the end of Bretton Woods[9] (Costa et al. 2024); the COVID-19 pandemic from March 2020; and the September 2022 liability-driven investment (LDI) episode following the UK's chancellor mini-budget announcement (Palacios and Patel 2023). These are isolated shocks rather than persistent switching between the fundamentalist and chartist regimes. We thus interpret the non-identification as a failure of the constant-coefficient switching model to capture the complex dynamics of the UK exchange rate during this period, rather than a failure of the underlying BUIP framework. A more complex specification of the model, allowing for time-varying switching parameters, could potentially capture these dynamics.

Diagnostic tests for the BUIP residuals (Table 6) mirror the STR diagnostics: non-normality is pervasive and ARCH effects remain common. The Ljung–Box test suggests that the BUIP specification reduces serial correlation for some developed markets (e.g., Australia, Canada, Euro Area, Switzerland, United Kingdom), while several emerging markets still display residual autocorrelation, indicating that additional dynamics (e.g., regime changes, risk premia, or time-varying volatility) may be needed for a fuller description.

4.2.3 Model Comparison

Table 7 summarizes the in-sample performance of the two models under study (STR and BUIP) against a linear benchmark for each country, defined as follows:

$$y_t = \phi_0 + \phi_c \Delta s_{t-1} + \phi_f \eta_{t-1} + \varepsilon_t,$$

with $\varepsilon_t \sim iid(0, \sigma^2)$. The STR estimation is not identified for Australia, Mexico, Thailand, Türkiye, and the UK because $\hat{c}$ lies outside the empirical support of the transition variable; for the BUIP, the same holds for Japan, Korea, and the UK. This means that the results for these countries should be interpreted with caution, or ignored. To compare the statistical performance of the three models, non-identification cases are indicated in parentheses in Table 7 and won't be considered.

[8] The United Kingdom has operated under flexible inflation targeting since the Bank of England obtained operational independence in 1997, with a 2% CPI inflation target pursued over the medium term alongside output and financial-stability considerations.

[9] Our exchange rate data show an 8.05% change from end of May to end of June, the largest monthly change among all countries.

Table 6: BUIP residual diagnostics. Test statistic with p-value in parentheses.

Country	Ljung-Box Q	ARCH LM	Jarque-Bera
Australia	6.9792 (0.7274)	**16.1039** (0.0029)	**95.8223** (0.0000)
Brazil	**116.5459** (0.0000)	**42.9387** (0.0000)	**1137.8206** (0.0000)
Canada	13.6909 (0.1876)	**15.7658** (0.0034)	**459.3778** (0.0000)
Euro Area	13.1940 (0.2130)	**44.5248** (0.0000)	**34.9023** (0.0000)
Indonesia	**18.4401** (0.0480)	**12.1655** (0.0162)	**14776.0672** (0.0000)
Japan	**18.9679** (0.0407)	5.8822 (0.2081)	**34.0992** (0.0000)
Korea	**23.6200** (0.0087)	**118.8522** (0.0000)	**2303.7557** (0.0000)
Mexico	**73.4924** (0.0000)	7.3842 (0.1169)	**76830.4361** (0.0000)
New Zealand	**19.0202** (0.0400)	**35.3253** (0.0000)	**57.6146** (0.0000)
Philippines	**44.5250** (0.0000)	8.7590 (0.0674)	**17646.0861** (0.0000)
Switzerland	9.9739 (0.4428)	**17.5657** (0.0015)	**12.9100** (0.0016)
Thailand	17.6522 (0.0611)	**58.8718** (0.0000)	**12733.6133** (0.0000)
Türkiye	**42.9295** (0.0000)	**43.7378** (0.0000)	**4739.0173** (0.0000)
United Kingdom	5.2295 (0.8753)	8.9690 (0.0619)	**101.2747** (0.0000)

Note: Ljung-Box Q uses 10 lags (null: no serial autocorrelation). ARCH LM uses 4 lags (null: no conditional heteroskedasticity). Jarque-Bera tests for normality of residuals. Bold entries indicate rejection of the null at the 5% level.

Table 7: In-sample fit comparison: linear benchmark, STR (Eq. (10)), and BUIP (Eqs. (11)–(12)).

	RMSE (lower is better)			AIC (lower is better)		
Country	Linear	STR	BUIP	Linear	STR	BUIP
Australia	0.0322	(0.0322)	**0.0319**	**-3506.13**	(-3485.40)	-3494.91
Brazil	0.1187	0.1145	**0.1060**	-2180.25	-2206.95	**-2281.42**
Canada	0.0205	0.0204	**0.0202**	**-3981.07**	-3971.79	-3975.31
Euro Area	0.0268	0.0267	**0.0261**	**-2289.66**	-2270.48	-2283.96
Indonesia	0.0609	0.0571	**0.0442**	-2608.09	-2644.93	**-2888.09**
Japan	0.0302	**0.0296**	(0.0297)	-3339.64	**-3347.17**	(-3335.19)
Korea	0.0290	**0.0275**	(0.0274)	-3618.03	**-3646.03**	(-3657.41)
Mexico	0.0535	(0.0433)	**0.0522**	-2999.03	(-3189.84)	**-3004.90**
New Zealand	0.0342	**0.0337**	0.0337	**-3247.73**	-3236.92	-3242.61
Philippines	0.0289	**0.0283**	0.0287	-3332.12	**-3332.67**	-3319.56
Switzerland	0.0307	0.0304	**0.0303**	**-3574.48**	-3573.40	-3567.17
Thailand	0.0273	(0.0274)	**0.0270**	**-3113.05**	(-3075.45)	-3102.66
Türkiye	0.0560	(0.0529)	**0.0537**	-2698.32	(-2727.91)	**-2720.02**
United Kingdom	**0.0278**	(0.0277)	(0.0272)	**-3677.91**	(-3660.45)	(-3677.27)

Note: RMSE = root mean squared error; AIC = Akaike Information Criterion. Bold values indicate the best (lowest) value for each metric within a country. Values in parentheses indicate models where the threshold/cost parameter (c, k) lies outside the empirical support of the transition variable, suggesting potential non-identification. For these cases, results should be interpreted with caution.

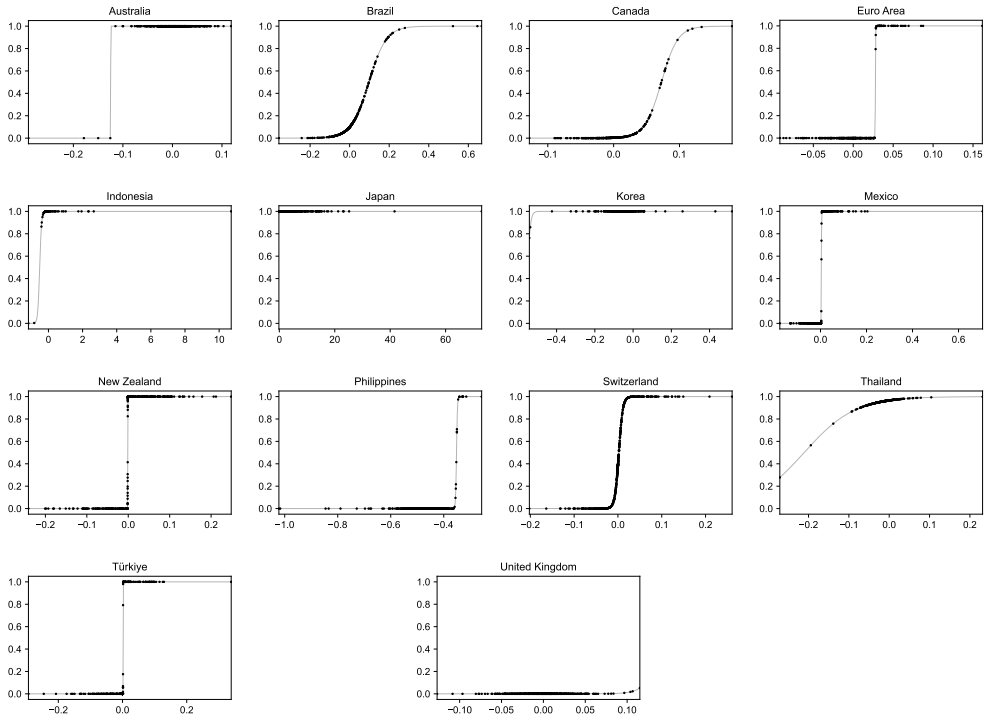


Fig. 3: BUIP model: scatter plots of fundamentalist market share ω_t (y-axis) against utility differential $U_{f,t} - U_{c,t}$ (x-axis).

According to RMSE, a nonlinear specification is preferred in 13 out of 14 countries: BUIP attains the lowest RMSE in nine economies (Australia, Brazil, Canada, Euro Area, Indonesia, Mexico, Switzerland, Thailand, Türkiye), while STR attains the lowest RMSE in four economies (Japan, Korea, New Zealand, Philippines), and the linear benchmark wins only for the United Kingdom. This confirms that, in terms of mean squared error, explicitly modeling regime switching dynamics improves in-sample fit for most of the currency pairs, with a notable improvement with the fully endogenous Behavioral UIP model. When penalizing for complexity via the AIC, linearity is rewarded in Advanced Economies. The linear benchmark is preferred in seven countries (Australia, Canada, Euro Area, New Zealand, Switzerland, Thailand, United Kingdom), illustrating that added complexity is not always rewarded once parsimony is accounted for. Among the nonlinear models, BUIP is preferred by AIC in Brazil, Indonesia, Mexico and Türkiye whereas STR is preferred in Japan, Korea and Philippines. This result suggests that nonlinear models (whether reduced form STR or the structural BUIP) matter more for emerging economies. Overall, these results suggest that nonlinearities matter for fit, but the degree to which they justify additional complexity is heterogeneous across markets and depends on their structure.

The non-overlapping pattern of identification failures across the two models is itself informative. The STR identifies but the BUIP fails for Japan and Korea, where

unconventional monetary policy suppressed the interest rate channel that drives the BUIP utility differential. The STR's data-driven selection of PPP return misalignment as transition variable provides identifying variation orthogonal to the rate differential, and is therefore unaffected by the ZLB environment. For Korea, the fundamentalist force survives BUIP estimation even when the switching point cannot be located, indicating that the failure is specific to the switching mechanism rather than to the underlying mean-reversion force.

The reverse cases, where the BUIP is identified but the STR is not, split along the sign of $\hat{\beta}_f$. For Australia and Thailand, the BUIP recovers a negative and significant $\hat{\beta}_f$, consistent with stabilizing PPP mean reversion; the endogenously constructed utility differential locates the transition where the reduced-form threshold cannot be pinned down. For Mexico and Türkiye, both specifications instead point to a positive $\hat{\beta}_f$, inconsistent with PPP-based mean reversion. The LM tests in Table 2 favor PPP-related transition variables for both economies (return misalignment for Mexico, fundamental change for Türkiye), so PPP deviations influence regime dynamics but not in a stabilizing direction. The BUIP achieves identification and improves on the linear benchmark by AIC, yet the estimated fundamentalist component lacks the intended economic content. These cases call for alternative specifications, as we discuss in Section 5.

Brazil and Indonesia provide the strongest internally consistent evidence for the behavioral hypothesis. Both STR specifications select PPP return misalignment as the transition variable, and the BUIP recovers precisely estimated switching intensities ($\hat{\gamma} = 2.30$ and 2.22) and is preferred by AIC over both the linear benchmark and the STR. The alignment between the STR's transition signal and the BUIP's profitability mechanism indicates that PPP misalignment serves both as the reduced-form trigger for regime change and as the structural driver of rule selection.

The Philippines is a partial exception: both models are identified, but the STR transition variable is return volatility ($|\Delta s_{t-3}|$) rather than a PPP-related measure, and the AIC marginally prefers the STR. The two models capture distinct aspects of the nonlinear dynamics in this case. The United Kingdom, where both fail in the full sample, is discussed in Section 4.3.

4.2.4 Advanced vs Emerging Economies

To synthesize the cross-country results, it is informative to contrast advanced and emerging economies along four dimensions: (i) evidence on fundamentalist forces, (ii) heterogeneity of technical trading, (iii) the extent to which complexity is justified by the data and (iv) the nature of the switching mechanism.

First, stabilizing fundamentalist forces are identified more frequently in the BUIP framework than in the STR framework, but they are not universal in either model. In the BUIP estimates, restricting attention to currencies for which $\hat{k}$ is identified, a negative and significant β_f emerges in several advanced economies (Australia, Euro Area, New Zealand, Switzerland) and in multiple emerging economies (Indonesia, Philippines, Thailand), suggesting that fundamental mean reversion is not confined to a particular group. Also, there is no clear distinction in terms of magnitude, so we cannot identify any particular regularity. Korea also yields a precisely estimated negative

$\hat{\beta}_f$, but is excluded from this count because $\hat{k}$ falls outside the empirical support of the utility differential and the regimes are not separated; we discussed this case in Section 4.2.2. At the same time, Mexico and Türkiye exhibit a positive and significant β_f in BUIP, and several countries display imprecisely estimated β_f, highlighting that the empirical relevance of the PPP fundamental might be market dependent, but not systematically so across advanced vs emerging economies.

Second, chartist behavior is heterogeneous, particularly among emerging markets. BUIP estimates show significant positive chartism in Brazil, Euro Area, Mexico, Philippines, and Türkiye, while other countries (Canada, Indonesia, New Zealand) display significant contrarian expectations. This heterogeneity is compatible with the idea that the composition of short-horizon strategies varies across markets and periods, and that "technical trading" need not imply trend following.

Third, parsimony is rewarded more often in advanced economies. The AIC favors the linear benchmark in all advanced and highly liquid markets, except Japan, while nonlinear models are more often preferred using AIC in emerging markets. In contrast, RMSE results indicate that nonlinear models frequently improve fit even when AIC does not select them, suggesting that nonlinearities can matter for fit but may be rare or insufficiently persistent to justify additional parameters.

Fourth, we focus the discussion of the switching mechanism on the BUIP estimates. The STR parameter γ governs the speed of transition around a fixed, exogenously chosen threshold and does not carry the evolutionary, profitability-based interpretation of the BUIP framework. Among the identified BUIP specifications ($\hat{k}$ in range, see Table 11), the intensity-of-choice parameter is precisely estimated only for Brazil ($\hat{\gamma} = 2.30$, $p < 0.05$) and Indonesia ($\hat{\gamma} = 2.22$, $p < 0.05$), both emerging economies. Both point estimates imply smooth, gradual switching in the sense of Section 3.1, with persistent coexistence of fundamentalists and chartists rather than synchronised herding. For the remaining identified specifications, confidence intervals for γ span several orders of magnitude, a well-known feature of logistic switching estimation (Teräsvirta 1994; Granger and Teräsvirta 1993a), so that point estimates are only suggestive of the shape of the transition and must be read alongside the scatter plots in Figure 3. The scatter plots reveal two qualitatively distinct patterns in the implied weight function $\omega_{f,t}$. Brazil, Canada, Indonesia, Philippines, Switzerland, and Thailand display a smoother, S-shaped $\omega_{f,t}$ that rises gradually as the fundamentalist payoff advantage increases, consistent with persistent heterogeneity and a relatively diffuse population of switching thresholds. Australia, Euro Area, Mexico, New Zealand, and Türkiye, by contrast, display a near-step function: $\omega_{f,t}$ remains close to zero or one across most of the observed utility differential, indicating that the population switches almost entirely in one direction once the profitability gap crosses the $\hat{k}$ threshold. Crucially, neither pattern maps cleanly onto the advanced/emerging divide: both groups contain smooth-switching and step-like specifications, and the only cases where γ is separately identified at conventional levels point toward smooth switching in two emerging markets. The BUIP evidence therefore does not support the prior that emerging-market currencies systematically exhibit the destabilizing step-function pattern of synchronized herding.

Overall, the evidence supports a view in which exchange rates can reflect nonlinear interactions between heterogeneous expectations, but the empirical support of these mechanisms—and the extent to which more complex models are needed—differs across currency pairs. Indeed, the advanced-vs-emerging dichotomy is a relevant distinguishing principle, but alternative dimensions, such as the monetary policy framework, could potentially provide other useful information for understanding the cross-country heterogeneity. This heterogeneity calls for further research into the structural features of each market that may drive the differences highlighted in this section, and for the development of more flexible models that can accommodate a wider range of nonlinear dynamics.

4.3 Robustness: Estimation without the COVID Period

The baseline estimation uses samples extending to 2025 and therefore includes the COVID-19 pandemic, a period of exceptional inflation dynamics, the subsequent central bank tightening cycle, supply-chain disruptions, and elevated exchange rate volatility, which may have disproportionately driven regime identification relative to the pre-2020 low-volatility period. To assess the sensitivity of the identification results to this episode, all models are re-estimated on samples ending in December 2019. Table 8 compares identification outcomes across the two sample periods.

For the STR, the results are qualitatively unchanged across most currencies: the selected transition variables, the sign and significance of the behavioral coefficients, and the AIC rankings are all preserved. The main exception is the United Kingdom, which achieves identification in the pre-COVID sample: the estimated $\hat{c} = 0.0004$ falls within the empirical support of $|\Delta f_{t-3}|$, and all parameters except $\hat{\gamma}$ are statistically significant with signs consistent with the hypothesis (Table 9, Panel A). This confirms that the full-sample identification failure for the United Kingdom is attributable to the structural disruption introduced by the pandemic rather than to a fundamental incompatibility between the model and GBP/USD dynamics. As a secondary result, the pre-COVID sample also delivers a negative and significant $\hat{\beta}_f = -0.034$ ($p < 0.05$) for Japan under the STR, consistent with PPP-based mean reversion; the full-sample estimate is positive and insignificant, suggesting that COVID-period dynamics partly obscured the fundamentalist component in the JPY/USD rate within the reduced-form specification.

For the BUIP, the pre-COVID results are more mixed (Table 9, Panel B). Japan and Korea, which fail to locate $\hat{k}$ in the full sample, both achieve identification in the pre-COVID estimation. For Korea, the improvement is substantial: all parameters reach statistical significance, with $\hat{\beta}_f = -0.058$ ($p < 0.01$), $\hat{\beta}_c = 3.196$ ($p < 0.01$), and $\hat{\gamma} = 0.274$ ($p < 0.01$) all precisely estimated. For Japan, the improvement is more modest: $\hat{k} = -0.005$ falls within the empirical support of the utility differential, but most parameters remain imprecise. These results suggest that pandemic-period disruptions, including modifications to the Bank of Japan's yield-curve control and amplified FX intervention in Korea, further compressed the utility differential relative to already-low pre-COVID levels, exacerbating identification problems that were present but less severe before 2020. The United Kingdom remains unidentified under the BUIP regardless of sample length, reinforcing the view that it might require a different fundamental

or a different structural specification. By contrast, Türkiye, identified in the full sample, loses BUIP identification in the pre-COVID estimation ($\hat{k} = 3.67$ far exceeds the empirical support maximum of 0.29), indicating that the identifying variation for the switching mechanism is concentrated in the post-2019 lira depreciation episodes following the "unconventional" monetary policy of the Turkish government.

Overall, the pre-COVID results confirm the main qualitative findings while revealing asymmetric sensitivity across specifications: the STR is largely robust to the pandemic period, with the UK as the principal exception, while BUIP identification is more sensitive, with improvements in Japan and Korea and a deterioration for Türkiye. The main conclusions, i.e. the superiority of the BUIP for Brazil and Indonesia, the preference for the linear benchmark in advanced economies, and the wrong-sign fundamentalist coefficient for Mexico, are preserved throughout.

Table 8: Parameter identification: full sample vs. pre-COVID sample.

	STR: $\hat{c}$ in range of z_t		BUIP: $\hat{k}$ in range of $(U_f - U_c)$	
Country	Full sample	Pre-2020	Full sample	Pre-2020
Australia	NO	NO	YES	YES
Brazil	YES	YES	YES	YES
Canada	YES	YES	YES	YES
Euro Area	YES	YES	YES	YES
Indonesia	YES	YES	YES	YES
Japan	YES	YES	NO	**YES**
Korea	YES	YES	NO	**YES**
Mexico	NO	NO	YES	YES
New Zealand	YES	YES	YES	YES
Philippines	YES	YES	YES	YES
Switzerland	YES	YES	YES	YES
Thailand	NO	NO	YES	YES
Türkiye	NO	NO	YES	**NO**
United Kingdom	NO	**YES**	NO	NO

Note: "YES"/"NO" = estimated threshold $\hat{c}$ (STR) or cost parameter $\hat{k}$ (BUIP) within/outside the empirical support of its driving variable (following Granger and Teräsvirta (1993a)). Bold entries mark changes relative to Table 11. The pre-2020 sample ends in December 2019.

5 Concluding Remarks

This paper tested UIP through the lens of a heterogeneous-expectations mechanism, with switching between fundamentalist and chartist forecasting rules, for fourteen currencies vis-à-vis the US dollar over long samples ending in 2025. We estimated (i) a reduced-form logistic smooth-transition regression (STR), where regime changes are driven by observable transition variables, and (ii) a structural behavioral UIP (BUIP) specification, where regime weights are endogenized through a profitability-based discrete-choice mechanism augmented by a wedge parameter k between fundamentalist and chartist utilities.

Table 9: Pre-COVID (pre-2020) estimates for countries with changed identification. Panel A: STR for the United Kingdom. Panel B: BUIP for Japan, Korea, and Türkiye. Full-sample counterparts in Tables 3–5. NLS with 95% Newey–West HAC confidence intervals.

Panel A: STR (pre-2020 sample)						
Country	α_0	α_1	β_c	β_f	γ	c
United Kingdom	**0.0090**** (0.0016, 0.0165)	**-0.0227***** (-0.0383, -0.0072)	**0.2883**** (0.0163, 0.5604)	**-0.0312**** (-0.0620, -0.0004)	512.0002 (-2539.6, 3563.6)	**0.0004***** (0.0003, 0.0004)
Panel B: BUIP (pre-2020 sample)						
Country	α_0	α_1	β_c	β_f	γ	k
Japan	0.0020 (-0.0028, 0.0067)	0.0220 (-0.0710, 0.1150)	0.0266* (-0.0016, 0.0548)	-0.0068 (-0.0271, 0.0136)	645.0441 (-37985, 39275)	-0.0048* (-0.0104, 0.0008)
Korea	**-0.3854***** (-0.6482, -0.1226)	**0.8317***** (0.4498, 1.2136)	**3.1962***** (1.2368, 5.1557)	**-0.0579***** (-0.0844, -0.0313)	**0.2736***** (0.0732, 0.4740)	**-0.7007***** (-1.0284, -0.3730)
Türkiye	**0.0600***** (0.0357, 0.0843)	5978.6591 [—]	**0.4866***** (0.3161, 0.6570)	0.8018 (-0.453, 2.057)	0.2420* (-0.044, 0.528)	3.6710 [—]

Note: 95% CIs in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$. N.I. = not identified. [—] = CI suppressed for readability (non-informative due to optimizer divergence). For Türkiye, $\hat{k} = 3.67$ lies outside the pre-2020 empirical support $[-1.21, 0.29]$, confirming non-identification.

Five conclusions follow from the empirical analysis. First, there is broad evidence of state dependence in exchange-rate returns. In the STR framework, linearity is rejected for 13 out of 14 currencies. Nonlinear models reduce RMSE relative to a linear benchmark in 13 out of 14 countries. At the same time, when penalizing complexity via the AIC, the linear benchmark is preferred in advanced economies, indicating that nonlinearities can be present but not always persistent or strong enough to justify additional parameters. To study emerging-economy FX markets, nonlinear models appear to be more relevant.

Second, the reduced-form STR provides only selective identification of the economic forces suggested by heterogeneous-agent interpretations. Evidence for trend-following behavior is limited, and PPP-consistent mean reversion is detected only for a subset of currencies; in some cases the estimated fundamentalist coefficient has a sign inconsistent with mean reversion to the PPP fundamental. These results suggest that a parsimonious STR may capture state dependence but can struggle to map it into the hypothesized behavioral channels.

Third, the structural BUIP specification of Proaño (2011, 2013) strengthens economic interpretation and more frequently identifies stabilizing PPP-consistent mean reversion. Restricting the count to currencies for which the BUIP switching point itself is identified, the fundamentalist adjustment parameter is negative and statistically significant in seven currencies (including both advanced and emerging markets); β_f is not universally identified and turns significantly positive for Mexico and Türkiye. This pattern indicates that the empirical relevance of PPP as the fundamental value is market-dependent, and that the meaning of the fundamentalist term is sensitive to the specification of the fundamental forecast.

Fourth, technical trading is heterogeneous across currency pairs. The BUIP estimates show significant positive chartist dynamics for several currencies, while other cases exhibit significant negative coefficients, consistent with contrarian short-horizon behavior rather than trend extrapolation. This heterogeneity is consistent with the view that the empirical content of β_c in our specification captures a broader class of short-horizon, non-fundamental dynamics than pure trend-following alone. Moreover, identification of the switching mechanism remains uneven: switching-intensity and threshold/cost parameters are precisely estimated only in a subset of cases, and in some currencies the estimated switching point (c or k) lies outside the empirical support of the corresponding driving variable, calling for caution in interpretation.

Fifth, the shape of the transition function is more frequently threshold-like than smooth, but the cross-country pattern does not map cleanly onto the advanced/emerging divide. The BUIP scatter plots reveal smooth, S-shaped market-share dynamics in both advanced (Switzerland) and emerging (Brazil, Indonesia, Philippines) currencies, alongside near-step transitions in both groups; following Granger and Teräsvirta (1993a), large $\hat{\gamma}$ in this latter group is read as a threshold transition rather than as non-identification, since identification of the switching mechanism is diagnosed through the empirical support of $\hat{k}$. The prior that emerging-market currencies should systematically exhibit more aggressive profitability-based switching than their advanced-economy counterparts, motivated by differences in liquidity and institutional depth, is not supported by the data. A pre-COVID robustness

exercise corroborates the main findings: the structural readings for the United Kingdom (non-identification confirmed by the robustness exercise, calling for different model specifications) and for Japan and Korea (BUIP identification recovers when the post-2020 period is excluded) are consistent with non-identification reflecting market-specific structural breaks rather than a generic failure of the BUIP framework.

Taken together, the evidence supports a view in which exchange-rate returns reflect nonlinear interactions between heterogeneous expectations, and in which such heterogeneity can improve in-sample fit. However, the gains from additional complexity are not uniform across markets, and the results highlight the need for more complex specifications mainly for emerging-economy FX markets.

The empirical pattern documented above leaves three concrete avenues open for future work: the choice of fundamental anchor, the horizon at which the two rules are evaluated, and the econometric handling of time-varying volatility and parameters. We discuss each in turn.

Alternative fundamental anchors. Mexico and Türkiye in particular display a positive and significant β_f at odds with PPP-consistent mean reversion, so the choice of fundamental is itself a margin of investigation. A first alternative builds on the real-side channel: Gori and Ricchiuti (2018) and Delli Gatti et al. (2025) define the fundamental through the ratio of exchange-rate indices consistent with trade-balance equilibrium, embedding chartist–fundamentalist trading in a macro agent-based environment where the anchor is endogenous to international trade flows. A closely related stream, exemplified by Bassi et al. (2023), centres the fundamental on expected treasury-bond differentials, so that the fundamentalist force is driven by the interest-rate channel rather than by price levels. A further route treats the fundamental as an exogenous stochastic process, typically a random walk or an AR(1), following de Grauwe and Grimaldi (2006); De Grauwe and Rovira Kaltwasser (2012); Mignot and Westerhoff (2025); this route is parsimonious but forgoes the structural mapping to macroeconomic data. Finally, richer alternatives derived from open-economy DSGE equilibrium conditions or from monetary-policy reaction functions offer tighter macroeconomic content at the cost of additional parametric assumptions. Our choice of absolute PPP keeps the structural specification as close as possible to Proaño (2011, 2013) and provides a transparent benchmark against which the observed deviations and identification failures can be interpreted. Extending the estimation to a broader set of anchors, notably a trade-balance-consistent target for emerging economies and a policy-consistent rate for advanced economies, is a natural next step that we leave for future research.

Horizon-asymmetric evaluation. Our baseline specification evaluates both fundamentalist and chartist utilities at a common monthly horizon, for tractability and continuity with the behavioral exchange-rate literature. Survey evidence, however, suggests that fundamental analysis operates at six- to twelve-month horizons while technical trading operates at intraday-to-weekly horizons (Cheung and Chinn 2001; Menkhoff and Taylor 2007). Allowing horizon-asymmetric evaluation—for instance, computing the fundamentalist squared forecast error over a six-month window while

retaining a monthly window for the chartist rule—would reconcile the switching structure with this horizon heterogeneity and is a promising direction for future work, complementing the fixed-share approach of Bassi et al. (2023).

Discrete-choice generalization. The logit-based switching mechanism is a close approximation to the more general probit framework in our two-rule, homogeneous-γ setting (Galanis et al. 2025). This no longer holds once the model is enriched with additional rules (for example, horizon-specific strategies, or multiple chartist/fundamentalist rules) or with cross-sectional heterogeneity in the weights traders assign to observables, capturing differences in risk aversion, attention, or information-processing costs. Estimating such extensions of the BUIP within a probit or mixed-logit framework is therefore a natural complement to the horizon-asymmetric and alternative-anchor extensions discussed above.

Econometric extensions. On the econometric side, allowing for time-varying volatility (e.g., STR–GARCH as in Westerhoff and Reitz (2003); Reitz and Westerhoff (2007) or related heteroskedastic specifications) and exploring rolling or time-varying parameter estimation could clarify how crisis episodes affect nonlinear dynamics and switching. On the structural side, enriching the set of forecasting rules, introducing alternative fundamentals beyond absolute PPP, and allowing utilities to incorporate risk considerations or longer memory may improve identification and provide a tighter link between exchange-rate dynamics, policy regimes, and macro fundamentals.

Acknowledgements. We thank the anonymous referees and editors for constructive comments. All remaining errors are our own.

Declarations

- **Funding**: No funding is to be declared.
- **Competing Interests**: The authors have no competing interests to declare that are relevant to the content of this article.
- **Data Availability**: Data are sourced from publicly available databases (IMF International Financial Statistics and OECD Main Economic Indicators), accessed via the DBnomics API. No proprietary data are used.
- **Code Availability**: Replication code is available at https://github.com/lorenzobiondiecon/str-buip-estimation.
- **Author Contributions**: The underlying behavioral UIP framework that the empirical analysis builds on was developed in prior work by Christian R. Proaño (Proaño 2011, 2013). For the present paper, both authors jointly designed the empirical study and the methodological extensions to the original framework. Data collection, econometric implementation, and estimation were carried out by Lorenzo Biondi. The first draft was written by Lorenzo Biondi and critically revised by Christian R. Proaño. Both authors read and approved the final manuscript.

Appendix A Appendix

A.1 Data Sources and Transformations

Table 10 summarizes the data availability and sources for each country in the sample. The exchange rate series correspond to nominal bilateral rates vis-à-vis the USD, expressed as the domestic price of USD. Consumer price indices (CPI) are used to construct real exchange rates. Policy or money market rates serve as short-term interest rate proxies. All series are at monthly frequency except for the CPI of Australia and New Zealand, for which monthly observations are not available and we linearly interpolate the quarterly series. Because the resulting real exchange rate enters the STR transition variable and the BUIP fundamental misalignment for both currencies, the CPI interpolation introduces some artificial smoothness into the corresponding nonlinear dynamics; we have not tested sensitivity of the regime-switching estimates to alternative interpolation choices and flag this as a possible confound for the AUS/USD and NZD/USD results. The data are fetched via the DBnomics API directly into Python using the `dbnomics` package.

Table 10: Data Availability and Sources by Country.

Country	Start Date	End Date	Observations	Source
Australia	1982-09	2025-04	512	IMF/IFS (S, CPI interp., i)
Brazil	1982-09	2025-06	514	IMF/IFS (S, CPI, MM rate)
Canada	1982-09	2025-06	514	IMF/IFS (S, CPI), OECD (i)
Euro Area	1999-02	2025-07	318	IMF/IFS (S, CPI), OECD (i)
Indonesia	1986-06	2025-05	468	IMF/IFS (S, CPI, i combined)
Japan	1985-08	2025-06	479	IMF/IFS (S, CPI), OECD (i)
Korea	1982-09	2025-05	513	IMF/IFS (S, CPI, MM rate)
Mexico	1982-09	2025-06	514	IMF/IFS (S, CPI, MM rate)
New Zealand	1985-02	2025-04	483	IMF/IFS (S, CPI interp., MM rate)
Philippines	1982-09	2021-12	472	IMF/IFS (S, CPI, MM rate)
Switzerland	1982-09	2025-07	515	IMF/IFS (S, CPI), OECD (i)
Thailand	1989-02	2025-03	434	IMF/IFS (S, CPI, MM rate)
Türkiye	1986-05	2025-06	470	IMF/IFS (S, CPI, i)
United Kingdom	1982-09	2025-07	515	IMF/IFS (S, CPI), OECD (i)

Note: S=Exchange Rate, CPI=Consumer Price Index, i=Policy/Interest Rate, MM=Money Market Rate. All data accessed via DBnomics API

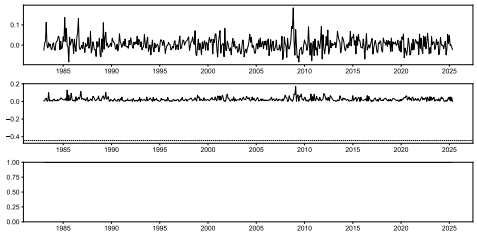
Table 11: Identification Summary: STR Threshold and BUIP Cost Parameter (In-Range Checks)

Country	STR: c in range of z_t	BUIP: k in range of $(U_f - U_c)$
Australia	NO	YES
Brazil	YES	YES
Canada	YES	YES
Euro Area	YES	YES
Indonesia	YES	YES
Japan	YES	NO
Korea	YES	NO
Mexico	NO	YES
New Zealand	YES	YES
Philippines	YES	YES
Switzerland	YES	YES
Thailand	NO	YES
Türkiye	NO	YES
United Kingdom	NO	NO

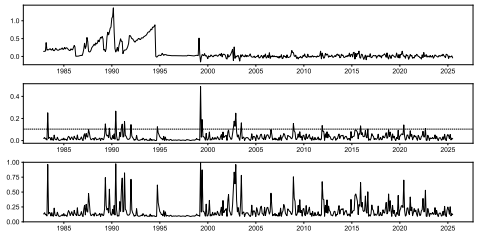
"In range" indicates whether the estimated threshold/cost parameter lies within the empirical support of the corresponding transition variable. "NO" flags potential identification concerns in the switching point.

A.2 Model Validation Summary

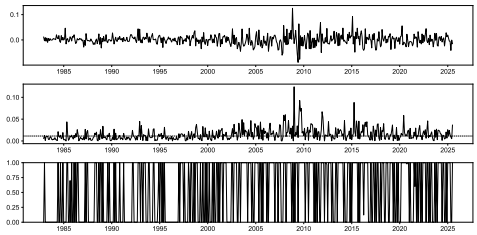
A.3 STR model additional figures



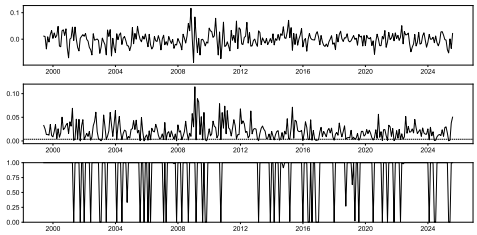
(a) Australia



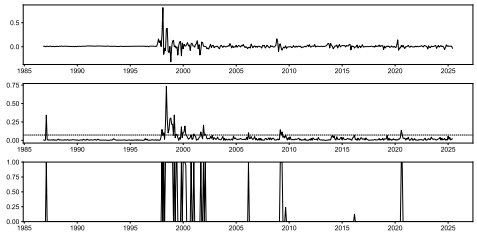
(b) Brazil



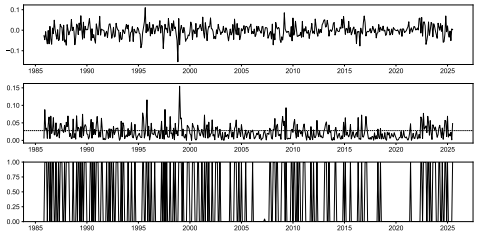
(c) Canada



(d) Euro Area

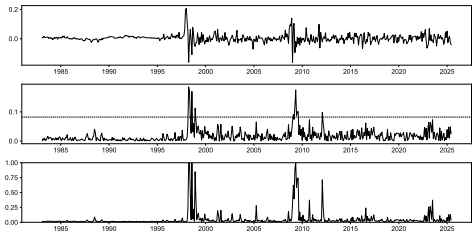


(e) Indonesia

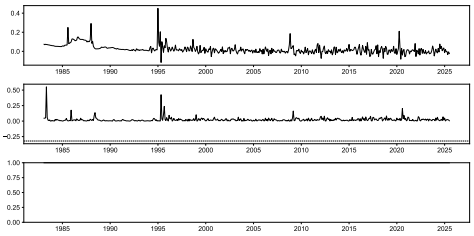


(f) Japan

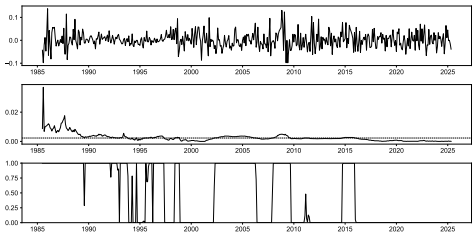
Fig. 4: STR model: time series of log nominal exchange rate returns (top panel), transition variable (middle panel), transition function G_t (bottom panel). The dashed line in the middle panel indicates the estimated threshold c. (Continued in the following pages.)



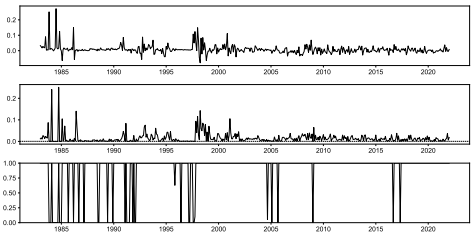
(g) Korea



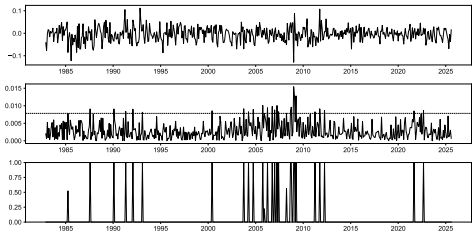
(h) Mexico



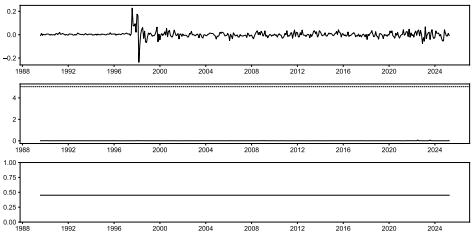
(i) New Zealand



(j) Philippines

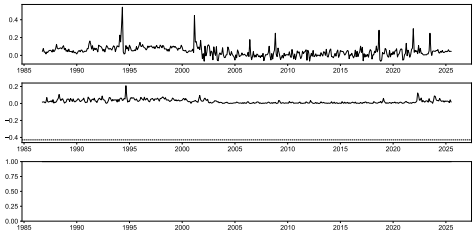


(k) Switzerland

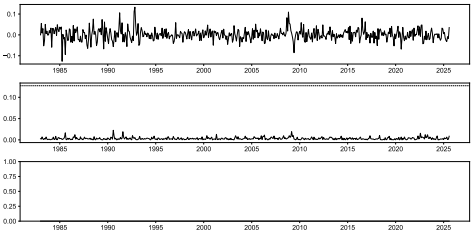


(l) Thailand

Figure 4 (continued).



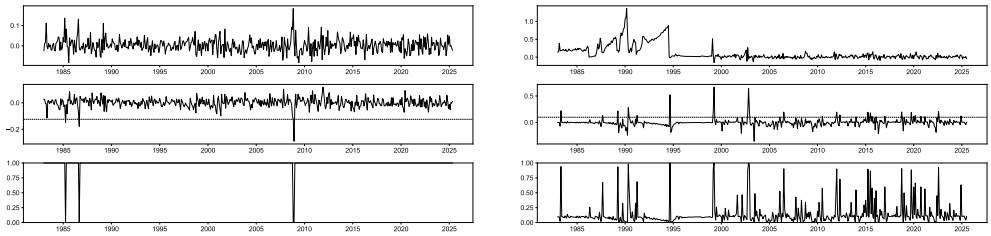
(m) Türkiye



(n) United Kingdom

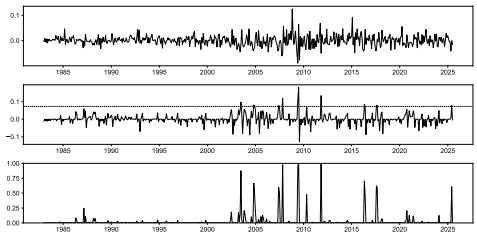
Figure 4 (continued).

A.4 BUIP model additional figures

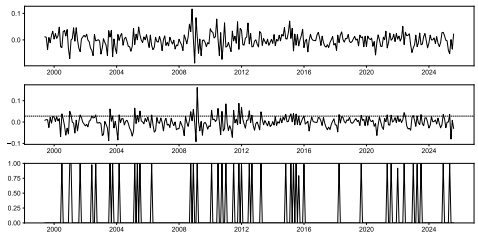


(a) Australia

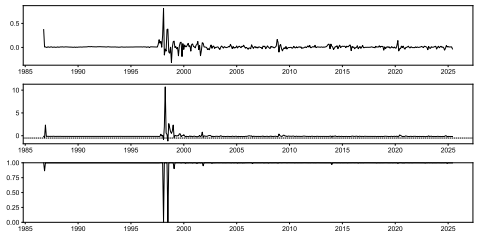
(b) Brazil



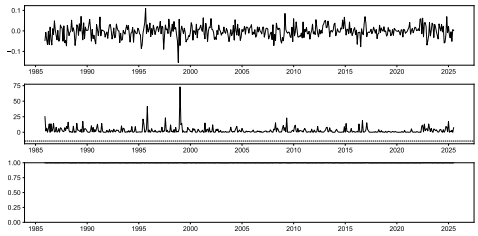
(c) Canada



(d) Euro Area

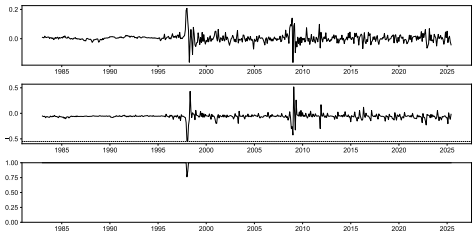


(e) Indonesia

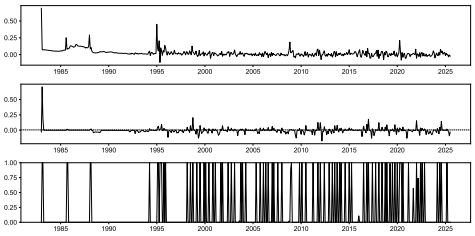


(f) Japan

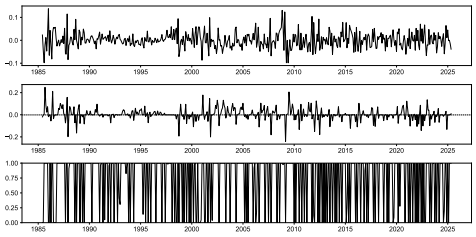
Fig. 5: BUIP model: time series of log nominal exchange rate returns (top panel), utility differential (middle panel), market share of fundamentalists ω_t (bottom panel). The dashed line in the middle panel indicates the estimated cost threshold k. (Continued in the following pages.)



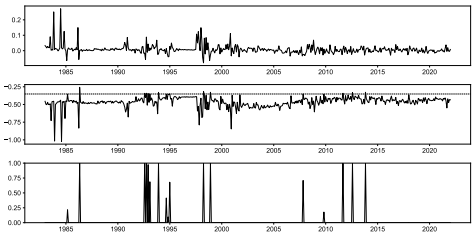
(g) Korea



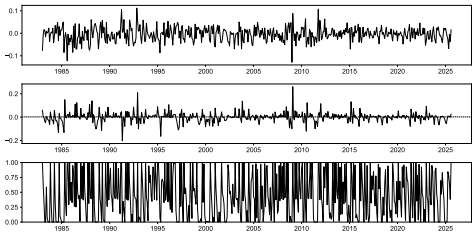
(h) Mexico



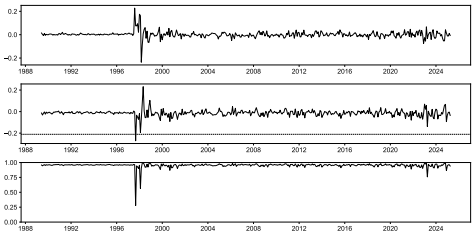
(i) New Zealand



(j) Philippines

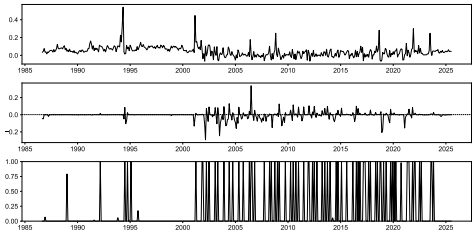


(k) Switzerland

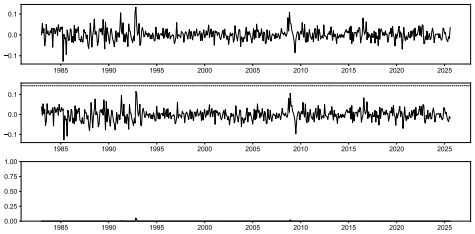


(l) Thailand

Figure 5 (continued).



(m) Türkiye



(n) United Kingdom

Figure 5 (continued).

References

Ahn J, Kim Y, Lim H (2022) For whom the levy tolls: The case of a macroprudential stability levy in south korea. IMF Economic Review 70(3):520–559

Baillie RT, Kapetanios G, Kim KH (2026) Yes! uncovered interest parity does hold in the long run. Journal of International Money and Finance 160:103455. https://doi.org/https://doi.org/10.1016/j.jimonfin.2025.103455

Bassi F, Ramos R, Lang D (2023) Bet against the trend and cash in profits: An agent-based model of endogenous fluctuations of exchange rates. Journal of Evolutionary Economics 33(2):429–472. https://doi.org/10.1007/s00191-023-00821-x

Bilson JFO (1978) The Monetary Approach to the Exchange Rate: Some Empirical Evidence (La théorie monétaire du taux de change: Preuves empiriques) (El enfoque monetario del tipo de cambio: Algunas pruebas empíricas). Staff Papers (International Monetary Fund) 25(1):48–75. https://doi.org/10.2307/3866655

Bolt W, Demertzis M, Diks C, et al (2019) Identifying booms and busts in house prices under heterogeneous expectations. Journal of Economic Dynamics and Control 103:234–259. https://doi.org/10.1016/j.jedc.2019.04.003

Boswijk HP, Hommes CH, Manzan S (2007) Behavioral heterogeneity in stock prices. Journal of Economic Dynamics and Control 31(6):1938–1970. https://doi.org/10.1016/j.jedc.2007.01.001

Branson WH, Henderson DW (1985) Chapter 15 The specification and influence of asset markets. Handbook of International Economics, vol 2. Elsevier, p 749–805, https://doi.org/10.1016/S1573-4404(85)02006-8

Brock WA, Hommes CH (1997) A Rational Route to Randomness. Econometrica 65(5):1059–1095. https://doi.org/10.2307/2171879, 2171879

Brock WA, Hommes CH (1998) Heterogeneous beliefs and routes to chaos in a simple asset pricing model. Journal of Economic Dynamics and Control 22(8):1235–1274. https://doi.org/10.1016/S0165-1889(98)00011-6

Bruno V, Shin HS (2014) Assessing macroprudential policies: Case of south korea. The Scandinavian Journal of Economics 116(1):128–157. https://doi.org/https://doi.org/10.1111/sjoe.12037

Bussiere M, Chinn M, Ferrara L, et al (2018) The New Fama Puzzle. Tech. Rep. w24342, National Bureau of Economic Research, Cambridge, MA, https://doi.org/10.3386/w24342

Cheung YW, Chinn MD (2001) Currency traders and exchange rate dynamics: A survey of the US market. Journal of International Money and Finance 20(4):439–471.

https://doi.org/10.1016/S0261-5606(01)00002-X

Cobham D (2024) Monetary policy frameworks since bretton woods, across the world and its regions. Review of World Economics 160(3):873–903. https://doi.org/10.1007/s10290-023-00517-1

Cornea-Madeira A, Hommes C, Massaro D (2019) Behavioral heterogeneity in u.s. inflation dynamics. Journal of Business & Economic Statistics 37(2):288–300. https://doi.org/10.1080/07350015.2017.1321548

Costa R, Dhingra S, Machin S (2024) New dawn fades: Trade, labour and the Brexit exchange rate depreciation. Journal of International Economics 152:103993. https://doi.org/10.1016/j.jinteco.2024.103993

de Grauwe P, Dewachter H (1993) A chaotic model of the exchange rate: The role of fundamentalists and chartists. Open Economies Review 4(4):351–379. https://doi.org/10.1007/BF01011136

De Grauwe P, Rovira Kaltwasser P (2012) Animal spirits in the foreign exchange market. Journal of Economic Dynamics and Control 36(8):1176–1192. https://doi.org/10.1016/j.jedc.2012.03.008

Delli Gatti D, Ferraresi T, Gusella F, et al (2025) The complex interplay between exchange rate and real markets: An agent-based model exploration. Journal of Economic Behavior & Organization 238:107252. https://doi.org/10.1016/j.jebo.2025.107252

Di Francesco T, Hommes C (2025) Sentiment-driven speculation in financial markets with heterogeneous beliefs: A machine learning approach. Journal of Economic Dynamics and Control 175:105092. https://doi.org/10.1016/j.jedc.2025.105092

Dick CD, Menkhoff L (2013) Exchange rate expectations of chartists and fundamentalists. Journal of Economic Dynamics and Control 37(7):1362–1383. https://doi.org/10.1016/j.jedc.2013.03.006

van Dijk D, Teräsvirta T, Franses PH (2002) SMOOTH TRANSITION AUTOREGRESSIVE MODELS — A SURVEY OF RECENT DEVELOPMENTS. Econometric Reviews 21(1):1–47. https://doi.org/10.1081/ETC-120008723

Dornbusch R (1976) Expectations and exchange rate dynamics. Journal of Political Economy 84(6):1161–1176. URL http://www.jstor.org/stable/1831272

Engel C (2014) Chapter 8 - exchange rates and interest parity. In: Gopinath G, Helpman E, Rogoff K (eds) Handbook of International Economics, Handbook of International Economics, vol 4. Elsevier, p 453–522, https://doi.org/10.1016/B978-0-444-54314-1.00008-2

Engel C, Hamilton JD (1990) Long swings in the dollar: Are they in the data and do markets know it? The American Economic Review 80(4):689–713. URL http://www.jstor.org/stable/2006703

Engel C, West KD (2005) Exchange rates and fundamentals. Journal of Political Economy 113(3):485–517. https://doi.org/https://doi.org/10.1086/429137

Engel C, Kazakova K, Wang M, et al (2022) A reconsideration of the failure of uncovered interest parity for the U.S. dollar. Journal of International Economics 136:103602. https://doi.org/10.1016/j.jinteco.2022.103602

Evans MD, Lyons RK (2008) How is macro news transmitted to exchange rates? Journal of Financial Economics 88(1):26–50. https://doi.org/10.1016/j.jfineco.2007.06.001

Evans MDD (2011) Exchange-Rate Dynamics. Princeton Series in International Economics, Princeton University Press, Princeton

Evans MDD, Lyons RK (2002) Order Flow and Exchange Rate Dynamics. Journal of Political Economy 110(1):170–180. https://doi.org/10.1086/324391, 10.1086/324391

Federici D, Gandolfo G (2012) The Euro/Dollar exchange rate: Chaotic or non-chaotic? A continuous time model with heterogeneous beliefs. Journal of Economic Dynamics and Control 36(4):670–681. https://doi.org/10.1016/j.jedc.2011.11.008

Franke R, Westerhoff F (2011) Estimation of a Structural Stochastic Volatility Model of Asset Pricing. Computational Economics 38(1):53–83. https://doi.org/10.1007/s10614-010-9238-7

Franke R, Westerhoff F (2018) Taking Stock: A Rigorous Modelling of Animal Spirits in Macroeconomics. In: Analytical Political Economy. p 5–38, https://doi.org/10.1002/9781119483328.ch2

Frankel JA (1979) On the mark: A theory of floating exchange rates based on real interest differentials. The American Economic Review 69(4):610–622. URL http://www.jstor.org/stable/1808707

Frankel JA, Froot KA (1990) Chartists, fundamentalists, and trading in the foreign exchange market. The American Economic Review 80(2):181–185. URL http://www.jstor.org/stable/2006566

Franses PH, Van Dijk D (2000) Nonlinear Time Series Models in Empirical Finance. Cambridge University Press, Cambridge, https://doi.org/10.1017/CBO9780511754067

Frenkel JA (1977) A Monetary Approach To The Exchange Rate: Doctrinal Aspects And Empirical Evidence. In: Herin J, Lindbeck A, Myhrman J (eds) Flexible Exchange Rates and Stabilization Policy. Palgrave Macmillan UK, London, p 68–92, https://doi.org/10.1007/978-1-349-03359-1_7

Frenkel JA, Mussa ML (1985) Chapter 14 Asset markets, exchange rates and the balance of payments. Handbook of International Economics, vol 2. Elsevier, p 679–747, https://doi.org/10.1016/S1573-4404(85)02005-6

Frijns B, Lehnert T, Zwinkels RC (2010) Behavioral heterogeneity in the option market. Journal of Economic Dynamics and Control 34(11):2273–2287. https://doi.org/10.1016/j.jedc.2010.05.009

Frijns B, Huynh T, Zwinkels RC (2025) Expectation formation in financial markets: Heterogeneity and sentiment. Journal of Economic Dynamics and Control 177:105133. https://doi.org/10.1016/j.jedc.2025.105133

Galanis G, Kollias I, Leventides I, et al (2025) Generalizing heuristic switching models and a (boundedly) rational route away from randomness. Journal of Economic Dynamics and Control 177:105125. https://doi.org/https://doi.org/10.1016/j.jedc.2025.105125

Galí J, Monacelli T (2005) Monetary policy and exchange rate volatility in a small open economy. The Review of Economic Studies 72(3):707–734. https://doi.org/10.1111/j.1467-937X.2005.00349.x

Gehrig T, Menkhoff L (2005) The rise of fund managers in foreign exchange:will fundamentals ultimately dominate? The World Economy 28(4):519–540. https://doi.org/https://doi.org/10.1111/j.1467-9701.2005.00690.x

Goldbaum D, Zwinkels RC (2014) An empirical examination of heterogeneity and switching in foreign exchange markets. Journal of Economic Behavior & Organization 107:667–684. https://doi.org/10.1016/j.jebo.2013.08.004

Gori M, Ricchiuti G (2018) A dynamic exchange rate model with heterogeneous agents. Journal of Evolutionary Economics 28(2):399–415. https://doi.org/10.1007/s00191-017-0513-9

Granger CWJ, Teräsvirta T (1993a) Modelling Nonlinear Economic Relationships. Oxford University Press, https://doi.org/10.1093/oso/9780198773191.001.0001

Granger CWJ, Teräsvirta T (1993b) Modelling Nonlinear Economic Relationships. Oxford University Press

de Grauwe P, Grimaldi M (2006) The Exchange Rate in a Behavioral Finance Framework. Princeton University Press, Princeton, N.J

Gusella F, Ricchiuti G (2024) Endogenous cycles in heterogeneous agent models: A state-space approach. Journal of Evolutionary Economics 34(4):739–782. https://doi.org/10.1007/s00191-024-00870-w

Hommes C (2013) Behavioral Rationality and Heterogeneous Expectations in Complex Economic Systems, 1st edn. Cambridge University Press, https://doi.org/10.1017/CBO9781139094276

Hur J, Shin K (2026) Does the uncovered interest parity hold better in korea? Journal of International Money and Finance 160:103470. https://doi.org/10.1016/j.jimonfin.2025.103470

Ilzetzki E, Reinhart CM, Rogoff KS (2017) The country chronologies to exchange rate arrangements into the 21st century: Will the anchor currency hold? Working Paper 23135, National Bureau of Economic Research, https://doi.org/10.3386/w23135

Kukacka J, Barunik J (2017) Estimation of financial agent-based models with simulated maximum likelihood. Journal of Economic Dynamics and Control 85:21–45. https://doi.org/10.1016/j.jedc.2017.09.006

Kukacka J, Sacht S (2023) Estimation of heuristic switching in behavioral macroeconomic models. Journal of Economic Dynamics and Control 146:104585. https://doi.org/10.1016/j.jedc.2022.104585

Lane PR (2001) The new open economy macroeconomics: A survey. Journal of International Economics 54(2):235–266. https://doi.org/10.1016/S0022-1996(00)00073-8

Lux T (2024) Lack of identification of parameters in a simple behavioral macroeconomic model. Journal of Economic Dynamics and Control 168:104972. https://doi.org/10.1016/j.jedc.2024.104972

Lux T, Zwinkels RC (2018) Chapter 8 - Empirical Validation of Agent-Based Models. In: Hommes C, LeBaron B (eds) Handbook of Computational Economics, Handbook of Computational Economics, vol 4. Elsevier, p 437–488, https://doi.org/10.1016/bs.hescom.2018.02.003

Lyons RK (2001) The Microstructure Approach to Exchange Rates. MIT Press, Cambridge, MA

Manzan S, Westerhoff FH (2007) Heterogeneous expectations, exchange rate dynamics and predictability. Journal of Economic Behavior & Organization 64(1):111–128. https://doi.org/10.1016/j.jebo.2006.08.005

Meese RA, Rogoff K (1983) Empirical exchange rate models of the seventies: Do they fit out of sample? Journal of International Economics 14(1):3–24. https://doi.org/10.1016/0022-1996(83)90017-X

Menkhoff L, Taylor MP (2007) The obstinate passion of foreign exchange professionals: Technical analysis. Journal of Economic Literature 45(4):936–972. https://doi.org/10.1257/jel.45.4.936

Michael P, Nobay AR, Peel DA (1997) Transactions costs and nonlinear adjustment in real exchange rates: An empirical investigation. Journal of Political Economy 105(4):862–879. https://doi.org/10.1086/262096

Mignot S, Westerhoff F (2025) Explaining the Stylized Facts of Foreign Exchange Markets with a Simple Agent-based Version of Paul de Grauwe's Chaotic Exchange Rate Model. Computational Economics 65(2):845–876. https://doi.org/10.1007/s10614-024-10546-z

Newey WK, West KD (1994) Automatic Lag Selection in Covariance Matrix Estimation. The Review of Economic Studies 61(4):631–653. https://doi.org/10.2307/2297912, 2297912

Obstfeld M, Rogoff K (1995) Exchange Rate Dynamics Redux. Journal of Political Economy 103(3):624–660. 2138701

Obstfeld M, Rogoff K (2000) The Six Major Puzzles in International Macroeconomics: Is There a Common Cause? NBER Macroeconomics Annual 15:339–390. https://doi.org/10.2307/3585403, 3585403

Palacios SIS, Patel KB (2023) Uk pension market stress in 2022-why it happened and implications for the us. Chicago Fed Letter

Park CH, Irwin SH (2007) What do we know about the profitability of technical analysis? Journal of Economic Surveys 21(4):786–826. https://doi.org/10.1111/j.1467-6419.2007.00519.x

Proaño CR (2011) Exchange rate determination, macroeconomic dynamics and stability under heterogeneous behavioral FX expectations. Journal of Economic Behavior & Organization 77(2):177–188. https://doi.org/10.1016/j.jebo.2010.10.005

Proaño CR (2013) Monetary policy rules and macroeconomic stabilization in small open economies under behavioral fx trading: Insights from numerical simulations. The Manchester School 81(6):992–1011. https://doi.org/10.1111/j.1467-9957.2012.02308.x

Reitz S, Westerhoff F (2007) Commodity price cycles and heterogeneous speculators: A STAR–GARCH model. Empirical Economics 33(2):231–244. https://doi.org/10.1007/s00181-006-0100-7

Rhee C (2026) Post-pandemic monetary policy in korea: Toward an integrated policy framework. Journal of International Money and Finance 161:103490. https://doi.org/10.1016/j.jimonfin.2025.103490

Rogoff K (1996) The Purchasing Power Parity Puzzle. Journal of Economic Literature 34(2):647–668. 2729217

Sarno L, Taylor MP (2001) Official intervention in the foreign exchange market: Is it effective and, if so, how does it work? Journal of Economic Literature 39(3):839–868. https://doi.org/10.1257/jel.39.3.839

Sarno L, Valente G, Leon H (2006) Nonlinearity in deviations from uncovered interest parity: An explanation of the forward bias puzzle. Review of Finance 10(3):443–482. https://doi.org/10.1007/s10679-006-9001-z

Schmitt N (2021) Heterogeneous expectations and asset price dynamics. Macroeconomic Dynamics 25(6):1538–1568. https://doi.org/10.1017/S1365100519000774

Shiratsuka S (2025) What did the Bank of Japan do under the yield curve control policy? Journal of the Japanese and International Economies 76:101352. https://doi.org/10.1016/j.jjie.2025.101352

Tay RS, Chen R (2018) Nonlinear Time Series Analysis, 1st edn. Wiley Series in Probability and Statistics, Wiley, https://doi.org/10.1002/9781119514312

Taylor MP, Peel DA, Sarno L (2001) Nonlinear Mean-Reversion in Real Exchange Rates: Toward a Solution to the Purchasing Power Parity Puzzles. International Economic Review 42(4):1015–1042. 826983

ter Ellen S, Hommes CH, Zwinkels RC (2021) Comparing behavioural heterogeneity across asset classes. Journal of Economic Behavior & Organization 185:747–769. https://doi.org/10.1016/j.jebo.2019.12.013

Teräsvirta T (1994) Specification, Estimation, and Evaluation of Smooth Transition Autoregressive Models. Journal of the American Statistical Association 89(425):208–218. https://doi.org/10.2307/2291217, 2291217

Westerhoff FH, Reitz S (2003) Nonlinearities and cyclical behavior: The role of chartists and fundamentalists. Studies in Nonlinear Dynamics & Econometrics 7(4). https://doi.org/doi:10.2202/1558-3708.1125