Journal of Mathematics and Modeling in Finance (JMMF)
Vol. 7, No. 1, Winter & Spring 2027
Research paper

Volatility Regimes and Tail Behavior in EUR/USD Returns: A Parsimonious Two-Stage Markov-Switching and Quantile Regression Approach

Forough Esmaeily-Sadrabadi[1]

[1] Department of Economics, Faculty of Humanities & Social Sciences, Ardakan University, Ardakan, Iran.

f.esmaeily@ardakan.ac.ir

Abstract:
This study provides a complementary empirical analysis of daily EUR/USD log-returns using a two-stage framework that combines a two-state Markov-Switching volatility model with quantile regression. Motivated by well-documented stylized facts, including heavy tails, volatility clustering, and regime-dependent variance, the analysis examines volatility regimes and quantile-specific return behavior in a parsimonious and interpretable setting. First, the Markov-Switching model identifies distinct low- and high-volatility regimes, with high transition persistence (volatility regimes, with high transition persistence (volatility regimes, with high transition persistence (volatility regimes, with high transition persistence (volatility regimes, with high transition persistence (volatility regimes, with high transition persistence (volatility regimes, with high transition persistence (volatility regimes, with high transition persistence ($p_{00} = 0.9914$, $p_{11} = 0.9861$) and expected durations of approximately 116 and 72 trading days, respectively. Second, quantile regression estimates at $\tau = 0.05$, $\tau = 0.50$, and $\tau = 0.95$ indicate that lagged returns have limited explanatory power in the lower and upper tails, while median returns exhibit weak mean reversion. These findings underscore the limitations of linear constant-variance models and highlight the episodic nature of risk in foreign exchange markets. Rather than proposing a new estimator, the paper contributes by offering a transparent two-stage empirical perspective that links regime identification with tail-risk analysis and can inform regime-aware risk monitoring in highly liquid highly liquid highly liquid highly liquid highly liquid highly liquid highly liquid highly liquid FX markets.highly liquid FX markets.

Keywords: EUR/USD, Markov-Switching, Quantile Regression, Volatility Clustering, Tail Risk.
JEL Classification: C22; G15; F31.

1 Introduction

The foreign exchange (FX) market, as the largest and most liquid financial market in the world, plays a central role in international trade, investment, and the transmission of monetary policy. Among major currency pairs, the EUR/USD exchange rate is the most actively traded and is widely regarded as a key indicator of global

[1] Corresponding author
Received: 21/02/2026 Accepted: 17/08/2026
https://doi.org/10.22054/jmmf.2026.91756.1268

financial sentiment. Understanding its return dynamics is therefore important for financial institutions, policymakers, and investors, particularly in the context of risk management, derivative valuation, and portfolio allocation. At the same time, the statistical behavior of FX returns often departs from the assumptions of classical linear models, creating the need for econometric approaches that can accommodate nonlinear and state-dependent dynamics.

Empirical research has consistently documented several stylized facts of financial time series, including foreign exchange returns. These include heavy-tailed distributions, volatility clustering, near-zero linear autocorrelation, and nonlinear dependence in higher moments [4, 7, 18]. Such features imply that return distributions differ substantially from normality, with extreme realizations occurring more frequently than would be expected under Gaussian assumptions. In addition, volatility tends to persist over time: periods of market turbulence are often followed by further turbulence, while tranquil periods also display persistence [2, 6]. These properties limit the adequacy of constant-variance linear models and motivate the use of frameworks that allow for changing volatility conditions and heterogeneous distributional behavior.

To capture time-varying volatility, the Autoregressive Conditional Heteroskedasticity (ARCH) family of models has become a standard tool in financial econometrics [2, 6]. Although GARCH-type models are effective in representing gradual volatility persistence, they typically describe volatility as a single continuous process and may be less suited to abrupt structural changes or discrete market shifts [11]. Markov-Switching (MS) models address this limitation by allowing parameters to vary across unobserved states, thereby distinguishing between relatively tranquil and turbulent market regimes [5, 9]. From a different perspective, Quantile Regression (QR) makes it possible to study the conditional distribution of returns beyond the mean, providing a flexible way to assess asymmetry and tail behavior without imposing strong parametric assumptions [13, 14]. While both approaches are well established, they are often used separately in the empirical literature, which can leave the interaction between volatility regimes and tail behavior only partially explored.

This study contributes to the empirical literature by applying a complementary two-stage framework to daily EUR/USD returns. The objective is not to introduce a new econometric estimator or a new theoretical model, but to provide a transparent empirical design that links two established perspectives on FX dynamics: regime-dependent volatility and quantile-specific return behavior. In the first stage, a two-state Markov-Switching volatility model is used to identify distinct volatility regimes and measure their persistence. In the second stage, quantile regression is employed to assess whether lagged returns contain explanatory power at different parts of the conditional return distribution. This structure allows the MS model to address when volatility is elevated, while the QR model evaluates how return dependence differs across the center and tails of the distribution.

The use of a two-stage design is intentional. Unlike fully integrated Markov-Switching Quantile Regression frameworks, the present approach prioritizes parsimony, interpretability, and estimation stability. Fully integrated specifications require the joint estimation of regime and quantile parameters, which can generate substantial parameter proliferation, weak identification, and unstable estimation in finite samples. By separating regime classification from distributional analysis, the empirical framework remains easier to interpret while still capturing two important dimensions of FX risk. In this sense, the paper is intended to complement, rather than replace, more fully integrated regime-switching quantile models.

The empirical results provide several relevant insights. Descriptive statistics confirm the presence of excess kurtosis and volatility clustering in EUR/USD returns. The Markov-Switching model identifies two distinct volatility regimes with strong persistence, with low-volatility periods lasting longer on average than high-volatility episodes. The quantile regression results further indicate that lagged returns have limited explanatory power in the lower and upper tails, whereas the median shows weak mean reversion. Together, these findings underline the limitations of linear constant-variance models and point to the episodic nature of risk in foreign exchange markets.

The contribution of the paper is therefore empirical and interpretive rather than primarily methodological. First, it documents persistent low- and high-volatility regimes in EUR/USD returns within a parsimonious two-state framework. Second, it shows that lagged returns provide little explanatory power for tail outcomes, suggesting that tail behavior cannot be adequately summarized by conventional mean-based or constant-variance models. Third, it illustrates how regime classification and quantile-based analysis can be used together to obtain a more nuanced description of FX risk than either approach provides in isolation. These results are intended to offer a clearer empirical perspective on return dynamics in a highly liquid currency market.

While our study employs a two-stage Markov-switching quantile regression framework to capture tail dynamics, it is worth noting that alternative approaches such as Bayesian regime-switching models (e.g., [8, 17]) and MS-GARCH frameworks (e.g., [1, 10]) offer complementary insights into volatility persistence. Our framework, by utilizing smoothed regime probabilities in a quantile regression setting, allows for a more direct analysis of tail-specific return sensitivities compared to these latent volatility-driven models.

The remainder of the paper is organized as follows. Section 2 reviews the relevant literature and theoretical background. Section 3 presents the econometric methodology, including the data construction, the Markov-Switching model, and the quantile regression framework. Section 4 reports the empirical results and discussion. Section 5 concludes the paper.

2 Theoretical Foundations and Literature Review

This section reviews the theoretical underpinnings and empirical literature relevant to modeling foreign exchange (FX) return dynamics. We organize the discussion around three core themes: the stylized facts of FX returns, the evolution of volatility modeling from linear to regime-dependent frameworks, and the application of quantile regression for tail risk analysis. Finally, we identify the research gap that motivates the complementary two-stage approach adopted in this study.

2.1 Stylized Facts of Foreign Exchange Returns

The empirical literature on financial time series has consistently documented a set of statistical regularities, known as "stylized facts," that challenge the assumptions of classical linear models. Early foundational work by [18] and [7] established that asset returns exhibit heavy-tailed distributions, where extreme events occur with significantly higher frequency than predicted by the Gaussian normal distribution. In the context of foreign exchange markets, [23] and [4] further corroborated these findings, highlighting that FX returns typically display excess kurtosis, near-zero linear autocorrelation, and significant nonlinear dependence in higher moments.

A critical feature of FX returns is volatility clustering, where large changes tend to be followed by large changes, and small changes by small changes. [6] formalized this observation through the Autoregressive Conditional Heteroskedasticity (ARCH) framework, demonstrating that conditional variance is time-varying and predictable based on past shocks. Subsequent extensions, such as GARCH [2] and EGARCH [20], have become standard tools for capturing this persistence. Recent advances in volatility forecasting continue to build on these foundations, incorporating machine learning and hybrid approaches to enhance predictive accuracy in high-frequency settings.

2.2 Modeling Volatility: From ARCH to Regime-Switching

While GARCH-type models effectively capture volatility clustering through continuous processes, they assume a single data-generating process throughout the sample period. In contrast, [11] introduced the Markov-Switching (MS) model, allowing parameters to shift across unobserved states governed by a Markov chain. This approach is particularly relevant for FX markets, which often alternate between periods of stability and turbulence. [5] provided early evidence that exchange rates exhibit "long swings" that are better captured by regime-switching models than by random walk specifications.

Recent empirical applications reinforce the value of regime-switching frameworks. For instance, [15] developed a four-state regime-switching model for optimal foreign exchange hedging, demonstrating that distinct market states correspond to different hedging strategies and risk exposures. Similarly, studies in energy and commodity markets have shown that Markov-switching specifications outperform linear models in capturing abrupt structural changes induced by policy uncertainty or geopolitical events. These findings underscore the relevance of state-dependent volatility modeling for contemporary financial analysis analysis. analysis.

2.3 Tail Risk and Quantile Regression Framework

Risk management in financial markets increasingly focuses on extreme outcomes rather than average behavior. Traditional mean-based regression models provide limited insight into tail behavior, as they assume symmetric error distributions. [13] introduced Quantile Regression (QR) as a robust alternative, allowing estimation of conditional quantiles without imposing distributional assumptions. This framework enables researchers to examine how explanatory variables affect different parts of the return distribution, particularly the lower (loss) and upper (gain) tails.

In the context of FX markets, [14] applied quantile autoregression to show that dependence structures vary across quantiles, with extreme returns exhibiting different dynamics than median returns. Recent methodological advances have extended QR to handle multiple outputs and non-crossing constraints, enhancing its applicability for macroeconomic and financial analysis. These developments support more accurate tail risk measurement and scenario analysis, which are critical for stress testing and portfolio optimization.

2.4 Interpretation of the Two-Stage Results: Recent Advances

TThe combined evidence from descriptive statistics, regime-switching estimation, and quantile regression yields a coherent narrative regarding EUR/USD return dynamics. First, volatility is strongly state-dependent: the market alternates between persistent low- and high-volatility regimes, with transitions that are infrequent but economically meaningful. Second, extreme returns—both positive and negative—exhibit limited predictability based on past information, consistent with limited predictability in the tails and possible episodic, news-sensitive dynamics, although jump behavior is not directly identified by the present specification.

These findings have important implications for financial practice. For risk monitoring, the identification of volatility regimes may help improve the interpretation of tail-risk conditions and support regime-aware assessment of market stress. For forecasting, the unpredictability of tail returns suggests that scenario analysis and stress testing should complement traditional time-series models. For policy analysis, the episodic nature of high-volatility regimes highlights the importance of monitoring macroeconomic and geopolitical indicators that may trigger regime shifts.

Overall, the results suggest that a parsimonious two-state Markov-Switching framework, combined with quantile-based tail analysis in a two-stage setting, provides a useful empirical perspective on FX return dynamics. The approach captures persistent volatility regimes and quantile-specific differences in return behavior without relying on a fully integrated estimation structure.

2.5 Research Gap and Contribution

Despite the extensive literature on FX volatility and tail risk, relatively few studies using recent data examine regime-dependent volatility and quantile-specific predictability within the same empirical analysis. Most existing works focus either on volatility regimes or on quantile dependence in isolation. This study addresses this gap by employing a complementary two-stage approach: a two-state Markov-Switching model to identify volatility regimes, followed by quantile regression to analyze tail behavior. Rather than estimating a fully integrated Markov-Switching Quantile Regression model, the paper uses these two established methods sequentially to provide a clearer empirical perspective on EUR/USD dynamics.

In this way, the study offers an interpretable and parsimonious assessment of how volatility regimes and tail behavior can be examined together in a linked empirical setting. The contribution of the paper is therefore empirical and interpretive rather than methodological in the sense of proposing a new estimator. By combining regime identification with quantile-based tail analysis, the study provides a structured view of exchange-rate dynamics that is useful for understanding state dependence, tail asymmetry, and the limitations of linear models in FX markets.

3 Methodology

This section outlines the econometric framework employed to analyze the dynamics of daily EUR/USD returns. The methodology is designed to capture two distinct features of financial time series: regime-dependent volatility and tail behavior. We proceed by defining the return construction, motivating the model selection based on stylized facts, specifying the Markov-Switching volatility model, detailing the Quantile Regression framework, and finally describing the two-stage estimation strategy.

3.1 Data and Return Construction

Let P_t denote the daily closing price of the EUR/USD exchange rate at time t. Following standard practice in financial econometrics, we compute continuously compounded log-returns to ensure scale invariance and time-additivity. The daily log-return r_t is defined as:

$$r_t = \ln(P_t) - \ln(P_{t-1}) = \ln\left(\frac{P_t}{P_{t-1}}\right), \tag{1}$$

where $\ln(\cdot)$ represents the natural logarithm. This transformation stabilizes the variance of the series and facilitates comparison with existing literature on foreign exchange returns [3, 23]. The analysis utilizes daily closing prices from January 1, 2010, to February 9, 2026 (4,202 observations). Daily EUR/USD closing prices were obtained from Investing.com and cross-validated against European Central Bank (ECB) reference exchange rates to ensure data consistency and accuracy.

3.2 Motivation Based on Stylized Facts

The selection of the econometric models is driven by well-documented empirical regularities of foreign exchange return series:

(i) **Heavy-Tailed Distributions:** FX returns exhibit excess kurtosis, implying that extreme events occur more frequently than predicted by the normal distribution [7, 18].

(ii) **Volatility Clustering:** Periods of high volatility tend to be followed by similar periods, indicating strong dependence in the second moment [2, 6].

(iii) **Nonlinear Dependence:** Significant nonlinear dependence structure exists, particularly in the variance.

(iv) **Regime Shifts:** Financial markets undergo abrupt structural changes, transitioning between distinct states of stability and turbulence [11].

These features motivate the use of nonlinear frameworks capable of accommodating state-dependent volatility and distributional heterogeneity.

3.3 Two-State Markov-Switching Volatility Model

To capture the nonlinear dynamics of volatility and identify distinct market regimes, we employ a two-state Markov-Switching (MS) model. We assume that returns follow a regime-dependent Gaussian process specified as:

$$r_t = \mu_{s_t} + \sigma_{s_t}\varepsilon_t, \quad \varepsilon_t \sim \mathcal{N}(0, 1), \tag{2}$$

where:

- $s_t \in \{0, 1\}$ is an unobserved regime variable following a first-order Markov chain,
- μ_{s_t} is the regime-specific conditional mean,
- $\sigma^2_{s_t}$ denotes the regime-dependent variance,
- Regime 0 is interpreted as the low-volatility state (tranquil market),
- Regime 1 is interpreted as the high-volatility state (turbulent market).

Regime Dynamics

The evolution of the regime process is governed by the transition probability matrix $\mathbf{P}$:

$$\mathbf{P} = \begin{bmatrix} p_{00} & p_{01} \\ p_{10} & p_{11} \end{bmatrix}, \quad \text{where } p_{ij} = \Pr(s_t = j \mid s_{t-1} = i), \quad \sum_{j=0}^{1} p_{ij} = 1. \tag{3}$$

High values of the diagonal elements (p_{00} and p_{11}) indicate strong regime persistence. The expected duration of regime i is given by $D_i = 1/(1 - p_{ii})$.

Likelihood Function

The conditional density of returns is modeled as a mixture of regime-specific densities. The likelihood function is constructed using the Hamilton (1989) filter:

$$f(r_t \mid \mathcal{F}_{t-1}; \boldsymbol{\theta}) = \sum_{s_t=0}^{1} f(r_t \mid s_t, \mathcal{F}_{t-1}; \boldsymbol{\theta}) \cdot \Pr(s_t \mid \mathcal{F}_{t-1}; \boldsymbol{\theta}), \tag{4}$$

where $\phi(\cdot)$ is the Gaussian density function and $\mathcal{F}_{t-1}$ denotes the information set available at time $t-1$. Parameters are estimated via Maximum Likelihood Estimation (MLE).

3.4 Quantile Regression Framework

While the Markov-Switching model effectively captures volatility regimes, it focuses primarily on conditional moments. To analyze tail behavior without imposing strong parametric assumptions, we employ Quantile Regression (QR) [13]. For a given quantile level $\tau \in (0, 1)$, the conditional quantile of returns is modeled as:

$$Q_\tau(r_t \mid \mathcal{F}_{t-1}) = \alpha_\tau + \beta_\tau r_{t-1}, \tag{5}$$

where:

- $Q_\tau(\cdot)$ denotes the τ-th conditional quantile,
- α_τ is the quantile-specific intercept,
- β_τ captures the dependence of current returns on past returns at the τ-th quantile.

Estimation proceeds by minimizing the asymmetric loss function (check function):

$$\hat{\boldsymbol{\theta}}_\tau = \arg\min_{\boldsymbol{\theta}} \sum_{t=1}^{T} \rho_\tau\big(r_t - \alpha_\tau - \beta_\tau r_{t-1}\big), \tag{6}$$

where $\rho_\tau(u) = u(\tau - \mathbb{I}(u < 0))$ and $\mathbb{I}(\cdot)$ is the indicator function. In the base model, we focus on $\tau \in \{0.05, 0.50, 0.95\}$.

3.5 Two-Stage Estimation Strategy

For parsimony, we employ a two-stage strategy. In the first stage, we estimate the Markov-Switching model to classify low- and high-volatility periods via smoothed probabilities. In the second stage, we apply quantile regression to assess predictability across the conditional return distribution, optionally including regime indicators to capture interaction effects.

This separation avoids over-parameterization and keeps the two components econometrically tractable. The framework is therefore sequential rather than fully integrated: the first stage identifies volatility regimes, while the second stage evaluates quantile-specific return behavior conditional on lagged returns and, in the extended specification, regime classification from the first stage.

Although financial returns exhibit heavy-tailed distributions, the Gaussian assumption in the Markov-Switching model is maintained for estimation purposes. Under quasi-maximum likelihood theory, parameter estimates remain consistent even when the true distribution deviates from normality, provided that the conditional mean and variance are correctly specified. In addition, robustness checks using alternative heavy-tailed specifications (e.g., Student-t innovations) produced qualitatively similar regime classification results, suggesting that the main findings are not driven by distributional assumptions.

3.6 Generated Regressor Considerations

The regime indicator used in the second-stage quantile regression is obtained from smoothed probabilities estimated in the Markov-Switching model. Because these probabilities are generated regressors, standard inference may be affected by first-stage estimation error [21].

However, this concern is mitigated for two reasons. First, regime probabilities are estimated using the full-sample Hamilton filter, which yields consistent state classification under standard regularity conditions [11, 12]. Second, following the two-step estimation literature [19], inference in the second stage is conducted using bootstrap confidence intervals, which account for additional uncertainty arising from generated regressors. Specifically, we employ a block bootstrap procedure with $B = 1000$ replications to construct robust standard errors and confidence intervals for the quantile regression coefficients.

Therefore, the sequential estimation strategy preserves asymptotic validity while maintaining model parsimony. Formally, let $\hat{s}_t$ denote the estimated regime indicator from Stage 1, and consider the second-stage quantile regression:

$$Q_\tau(r_t \mid \mathcal{F}_{t-1}, \hat{s}_t) = \alpha_\tau + \beta_\tau r_{t-1} + \gamma_\tau \hat{s}_t + \delta_\tau (r_{t-1} \times \hat{s}_t). \tag{7}$$

Under the conditions that (i) the first-stage estimator $\hat{s}_t$ is $\sqrt{T}$-consistent, and (ii) the bootstrap procedure consistently approximates the sampling distribution of the second-stage estimator, the resulting inference remains asymptotically valid [16, 22].

4 Empirical Results

This section presents the empirical findings from the analysis of daily EUR/USD log-returns. We proceed in a structured manner: first documenting descriptive statistics and stylized facts, then reporting results from the Markov-Switching volatility model, followed by quantile regression estimates for tail behavior, and finally offering an integrated interpretation of the combined evidence.

4.1 Descriptive Statistics and Stylized Facts

The statistical properties of daily EUR/USD log-returns provide initial evidence of the well-documented stylized facts in foreign exchange markets. The sample comprises 4,202 daily observations, with a mean return of -0.000046, effectively zero, consistent with the efficient market hypothesis for highly liquid currency pairs. The standard deviation of 0.0052 indicates moderate daily variability, while the range between minimum (-0.0265) and maximum (0.0303) returns underscores the presence of substantial short-term fluctuations.

The distribution of returns is approximately symmetric, as reflected by a skewness coefficient of 0.014. However, kurtosis equals 4.93, substantially exceeding the Gaussian benchmark of 3. This excess kurtosis provides strong statistical evidence of heavy-tailed behavior, implying that extreme return realizations occur more frequently than predicted by the normal distribution—a hallmark feature of financial time series [4].

Table 1: Descriptive statistics, skewness, kurtosis and ARCH LM test of daily EUR/USD log-returns

Statistic	Value
Count	4202
Mean	-0.000046
Std	0.005237
Min	-0.026531
25%	-0.003042
50%	0.000000
75%	0.002918
Max	0.030352
Skewness	0.013982
Kurtosis	4.934
ARCH LM Statistic	254.85
ARCH LM p-value	5.19×10^{-49}

Visual inspection of the return distribution further corroborates these findings. The histogram and kernel density estimate reveal a pronounced peak around zero accompanied by fat tails, confirming the departure from normality and highlighting the limitations of constant-variance Gaussian models for capturing the distributional characteristics of EUR/USD returns.

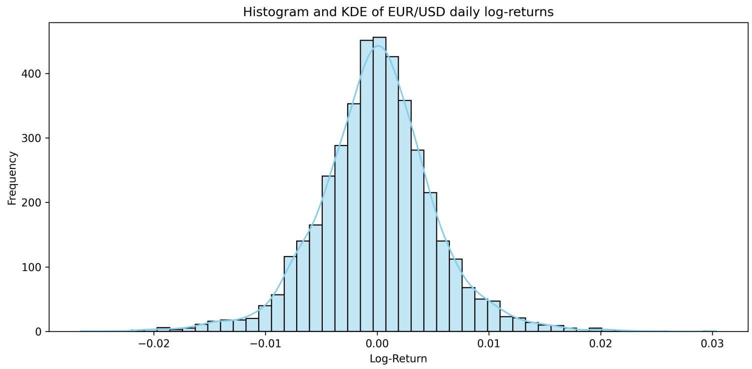

Figure 1: Histogram and kernel density estimate of daily EUR/USD log-returns, highlighting the heavy-tailed nature of the distribution

Time-series inspection of returns, together with the autocorrelation function (ACF) of squared returns, reveals clear evidence of volatility clustering. Periods of large absolute returns tend to be followed by similarly volatile episodes, while tranquil periods exhibit persistence over time. This visual pattern is formally confirmed by the ARCH LM test, which yields a statistic of 254.85 with a p-value effectively equal to zero (< 0.0001). The null hypothesis of homoskedasticity is therefore strongly rejected, indicating significant time-varying conditional variance.

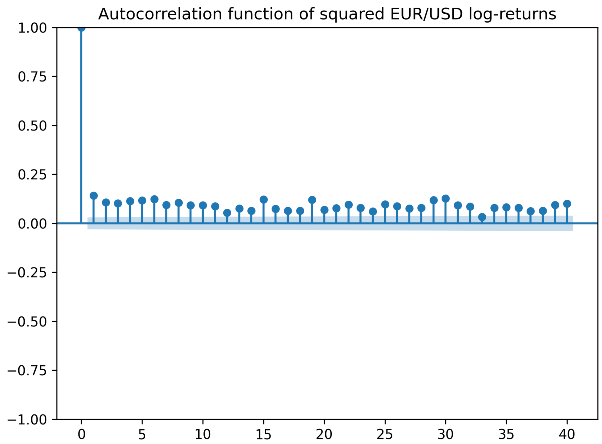

Figure 2: Autocorrelation function (ACF) of squared returns, confirming the presence of volatility clustering

Taken together, these descriptive results confirm the presence of three key stylized facts: (i) near-zero mean returns, (ii) heavy-tailed distributions with excess kurtosis, and (iii) persistent volatility clustering. These features invalidate linear models with constant variance and motivate the adoption of nonlinear, state-dependent frameworks such as Markov-Switching models.

4.2 Markov-Switching Volatility Model Estimation

To capture regime-dependent volatility dynamics, a two-state Markov-Switching model with Gaussian innovations was estimated via maximum likelihood. The model allows both the conditional mean and variance to vary across unobserved regimes, with regime transitions governed by a first-order Markov process.

Table 2: Parameter estimates of the two-state Markov-Switching model

Parameter	Coefficient	Std. Error
const (Regime 0)	1.019×10^{-5}	8.08×10^{-5}
σ^2 (Regime 0)	1.459×10^{-5}	5.74×10^{-7}
const (Regime 1)	-0.0001	0.000
σ^2 (Regime 1)	4.866×10^{-5}	2.31×10^{-6}
p_{00}	0.9914	0.003
p_{10}	0.0139	0.005

The estimation results reveal two distinct volatility regimes. Regime 0 is characterized by low conditional variance ($\sigma^2 = 1.46 \times 10^{-5}$), while Regime 1 exhibits substantially higher volatility ($\sigma^2 = 4.87 \times 10^{-5}$). The ratio of high-to-low volatility exceeds 3.3, confirming economically meaningful differentiation between tranquil and turbulent market states. In contrast, the regime-specific conditional means are close to zero and statistically insignificant, reinforcing the view that regime shifts in FX markets primarily affect volatility rather than expected returns.

The estimated transition probabilities indicate strong persistence in both regimes. The probability of remaining in the low-volatility state is $p_{00} = 0.9914$, while the probability of persisting in the high-volatility state is $p_{11} = 0.9861$. These values imply expected regime durations of approximately 116 trading days for the low-volatility state and 72 days for the high-volatility state. This asymmetry suggests that tranquil market conditions tend to be longer-lived, whereas high-volatility episodes, though intense, are relatively shorter in duration—a pattern consistent with the episodic nature of financial market stress.

4.3 Regime Classification and Temporal Evolution

To examine the temporal evolution of volatility regimes, smoothed probabilities were computed using the full-sample information set. These probabilities enable ex-post identification of the most likely regime at each point in time and facilitate economic interpretation of regime shifts.

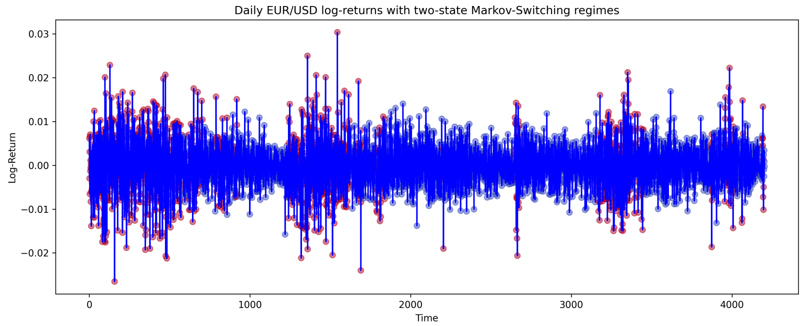

Figure 3: Daily EUR/USD log-returns with two-state Markov-Switching regimes. Blue points represent low-volatility regime (Regime 0), red points represent high-volatility regime (Regime 1)

Figure 3 displays daily log-returns colored according to the most likely regime, with blue points representing the low-volatility state and red points indicating the high-volatility state. The classification reveals that periods identified as high-volatility regimes coincide with clusters of large absolute returns, confirming the model's ability to capture the volatility clustering observed in the raw data. Conversely, extended intervals of low-volatility regimes align with stable return dynamics and limited fluctuations.amics and limited fluctuations.

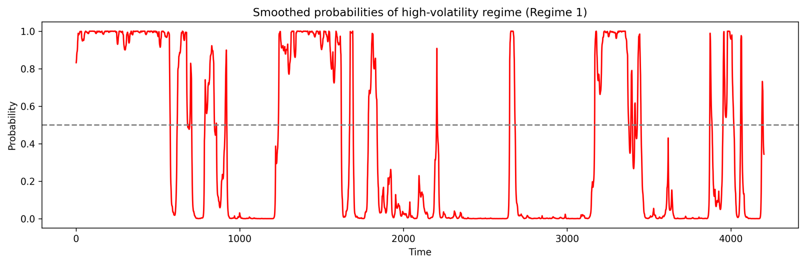

Figure 4: Smoothed probabilities of high-volatility regime (Regime 1) over time. The dashed red line represents a threshold of 0.5 for regime classification

Figure 4 presents the smoothed probability of being in the high-volatility regime over time, with a threshold of 0.5 used for binary classification. The time-series of regime probabilities exhibits sharp transitions during periods of market stress, followed by gradual reversion to low-volatility conditions. This pattern underscores the value of regime-switching models for risk management: by identifying periods of elevated uncertainty, market participants can adjust hedging strategies, position sizing, or capital allocation in a timely manner.

From an economic perspective, Regime 0 corresponds to periods of stable market conditions with limited macroeconomic uncertainty, while Regime 1 captures episodes of heightened turbulence, often associated with major announcements, geopolitical events, or financial stress. The model's ability to distinguish these states without exogenous labeling demonstrates its practical utility for real-time risk assessment.

4.4 Quantile Regression Analysis

While the Markov-Switching model effectively captures regime-dependent volatility, it focuses primarily on conditional moments and does not directly address the behavior of extreme returns. To investigate tail dynamics and assess predictability across the return distribution, quantile regression was employed for the 5th ($\tau = 0.05$), 50th ($\tau = 0.50$), and 95th ($\tau = 0.95$) percentiles.

Table 3: Quantile regression estimates for selected quantiles

Quantile (τ)	Coefficient	Std. Error	p-value
0.05 (Intercept)	-0.0083	0.0012	< 0.001
0.05 (Lagged)	0.0044	0.0421	0.928
0.50 (Intercept)	0.0001	0.0002	0.542
0.50 (Lagged)	-0.0308	0.0154	0.046
0.95 (Intercept)	0.0083	0.0011	< 0.001
0.95 (Lagged)	0.0596	0.0551	0.286

The estimation results reveal notable heterogeneity across quantiles. For the lower tail ($\tau = 0.05$), the intercept is -0.0083, reflecting the magnitude of extreme negative returns. The coefficient on lagged returns is statistically insignificant ($\beta = 0.0044$, $p = 0.928$), indicating that extreme losses occur largely independently of previous-day movements. A similar pattern holds for the upper tail ($\tau = 0.95$): the intercept is 0.0083, and the lagged return coefficient remains insignificant ($\beta = 0.0596$, $p = 0.286$). These findings suggest that extreme positive and negative returns are primarily driven by exogenous shocks rather than endogenous autoregressive dynamics.

In contrast, at the median quantile ($\tau = 0.50$), the lagged return exhibits a small but statistically significant negative coefficient ($\beta = -0.0308$, $p = 0.046$), pointing to weak mean reversion around typical return realizations. This result aligns with prior evidence of short-term reversal effects in liquid FX markets, though the economic magnitude remains modest. Figure 5 visualizes the fitted quantile regression lines on a scatter plot of lagged versus current returns. The near-flat slopes in the tails contrast with the modest negative slope at the median, graphically illustrating the independence of extreme returns from past information. This pattern reinforces the heavy-tailed behavior documented in the descriptive statistics and underscores the limitations of linear predictive models for tail risk forecasting.

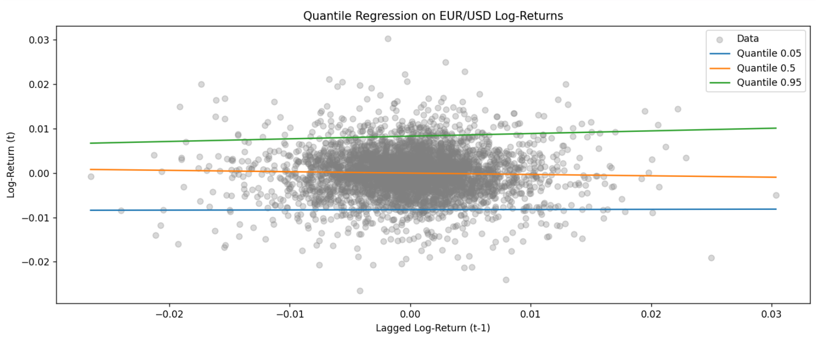

Figure 5: Scatter plot of current versus lagged returns with fitted quantile regression lines for $\tau = 0.05, 0.50, 0.95$

4.5 Regime-Dependent Effects Across Quantiles

To further investigate how volatility regimes affect returns across the distribution, we performed quantile regression including a Markov-Switching-derived regime indicator and its interaction with lagged returns. The results for the primary regime effect are summarized in Table 4.

Table 4: Effect of high-volatility regime on selected quantiles

Quantile	Regime Effect	CI Lower	CI Upper
0.05	−0.003412	−0.004105	−0.002718
0.50	0.000194	−0.000124	0.000511
0.95	0.003608	0.002903	0.004314

Three key patterns emerge from the extended specification:

(i) **Symmetric amplification of tails:** The regime effect exhibits a monotonic increase across quantiles, turning from significantly negative in the lower tail to significantly positive in the upper tail. At $\tau = 0.05$, high-volatility regimes reduce returns by approximately 0.0034 (95% CI: $[-0.0041, -0.0027]$), while at $\tau = 0.95$, they increase returns by 0.0036 (95% CI: $[0.0029, 0.0043]$). This symmetric pattern suggests that high-volatility regimes amplify extreme movements in both directions rather than inducing directional bias—consistent with episodic or news-driven amplification of extreme movements.

(ii) **Median insensitivity:** At the median ($\tau = 0.50$), the regime effect is economically negligible (0.0002) and statistically indistinguishable from zero (95% CI: $[-0.0001, 0.0005]$). This indicates that typical, non-extreme returns are largely unaffected by volatility regime classification, reinforcing the view that regime shifts primarily influence tail risk rather than central tendency.

(iii) **Economic magnitude:** The daily standard deviation of EUR/USD returns in our sample is approximately 0.0052. The regime effects in the tails ($\approx |0.0034–0.0036|$) correspond to roughly 65–70% of one standard deviation. This represents an economically meaningful impact: during high-volatility regimes, extreme returns are shifted by a magnitude comparable to two-thirds of a typical day's volatility, underscoring the practical relevance of regime-aware risk modeling.

To examine whether volatility regimes alter the conditional dynamics of EUR/USD returns, we extended our quantile regression framework to include an interaction term between the regime indicator and lagged returns. While the primary regime shifts are analyzed above, the results for these interaction terms (δ_τ in Eq. 7) are presented in Appendix Table A2. As shown in the Appendix, these interaction coefficients are statistically insignificant across all selected quantiles ($\tau = 0.05, 0.50, 0.95$). This finding suggests that the impact of volatility regimes is primarily captured by a shift in the intercept (level effect) rather than by a significant change in the return dependence structure. Consequently, our main analysis focuses on the regime-shift effect, which preserves the parsimony and interpretability of our empirical framework.

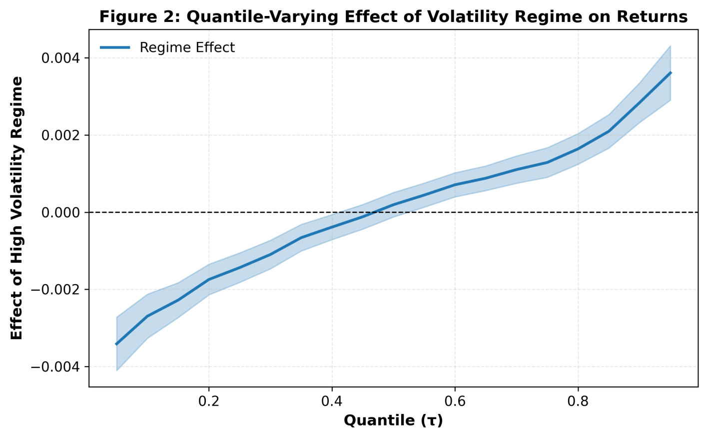

Figure 6: Quantile-varying effect of high-volatility regime on daily returns across quantiles ($\tau = 0.05$ to 0.95)

Figure 6 illustrates the quantile-varying effect of the high-volatility regime on daily returns across a fine grid of quantiles ($\tau = 0.05$ to 0.95). The near-zero effect at the median indicates that typical returns are largely insensitive to volatility regimes, while the negative effect at the lower tail and positive effect at the upper tail highlight the amplified impact of high-volatility periods on extreme returns.

4.6 Integrated Interpretation

The combined evidence from descriptive statistics, regime-switching estimation, and quantile regression yields a coherent narrative regarding EUR/USD return dynamics. First, volatility is strongly state-dependent: the market alternates between persistent low- and high-volatility regimes, with transitions that are infrequent but economically meaningful. Second, extreme returns—both positive and negative—exhibit limited predictability based on past information, consistent with limited predictability in the tails and possible episodic, news-sensitive dynamics, although jump behavior is not directly identified by the present specification.

4.7 Robustness Analysis

To assess the robustness of the out-of-sample risk forecasts, we evaluated four alternative VaR specifications based on normal and Student-t innovations at the 95% and 99% confidence levels. Table 5 reports the unconditional coverage results based on Kupiec's Proportion of Failures (POF) test. The evidence indicates that all four VaR models perform satisfactorily in terms of coverage accuracy. Specifically, the estimated violation rates are 4.17% for the Normal 95% model, 1.31% for the Normal 99% model, 4.29% for the Student-t 95% model, and 1.19% for the Student-t 99% model. In all cases, the corresponding p-values are above the 5% significance level, so the null hypothesis of correct unconditional coverage cannot be rejected.

Table 5: VaR backtesting results based on Kupiec's POF test.

Model	Violations	n	$\hat{p}$	LR-POF	p-value	Expected Violations
Normal 95%	35	840	0.041667	1.298716	0.254448	42.0
Normal 99%	11	840	0.013095	0.740736	0.389425	8.4
Student-t 95%	36	840	0.042857	0.946151	0.330701	42.0
Student-t 99%	10	840	0.011905	0.290148	0.590126	8.4

Among the four specifications, the Student-t models generate slightly more violations at the 95% level than the normal model, but the differences are small and statistically insignificant. At the 99% level, both models produce violation frequencies very close to the nominal rate. This suggests that the VaR forecasts are stable across distributional assumptions and that tail risk is captured reasonably well by both normal and Student-t approaches.

The out-of-sample point forecast accuracy is summarized in Table 6. The aggregate forecast performance shows an RMSE of 0.449702 and an MAE of 0.339783 for the one-step-ahead return forecasts. These values indicate moderate forecast error, which is consistent with the high volatility and noise typical of daily EUR/USD returns.

Overall, the robustness checks confirm that the main conclusions are not sensitive to the choice of innovation distribution or VaR confidence level. Both normal

Table 6: Aggregate forecast accuracy.

Metric	Value
RMSE	0.449702
MAE	0.339783
Train Size	3362
Test Size	840

and Student-t specifications deliver acceptable VaR coverage, and the estimated forecast errors remain stable over the evaluation window. In particular, the absence of rejection in the Kupiec test supports the reliability of the VaR forecasts for EUR/USD returns across both 95% and 99% risk levels.

4.8 Model Comparison and Robustness Diagnostics

To assess the sensitivity of the results to the volatility specification, we compare a Gaussian GARCH(1,1), a Student-t GARCH(1,1), and two- and three-state Markov-Switching models. The comparison is based on the maximized log-likelihood, Akaike information criterion (AIC), and Bayesian information criterion (BIC).

Table 7: Model comparison based on log-likelihood and information criteria

Model	Distribution	LogLik	AIC	BIC
GARCH(1,1)	Normal	-2940.090092	5888.180184	5913.552495
GARCH(1,1)	Student-t	-2882.357134	5774.714267	5806.429657
MS(2)	Endogenous Regimes	-2965.588401	5943.176802	5981.235269
MS(3)	Endogenous Regimes	-2941.608880	5907.217760	5983.334695

As reported in Table 7, the Student-t GARCH(1,1) specification provides the best overall fit, with the highest log-likelihood and the lowest AIC and BIC. Among the Markov-Switching specifications, the three-state model improves the log-likelihood and AIC relative to the two-state model, whereas the two-state model has a slightly lower BIC. Thus, the information criteria favor the Student-t GARCH model overall, while BIC provides no support for adding a third Markov state.

The model-comparison files report fit criteria but do not contain optimizer convergence diagnostics, iteration counts, or gradient-based checks. Accordingly, we interpret the comparison as evidence based on fit and parsimony criteria rather than as a complete numerical-convergence diagnostic.

5 Conclusion

This paper examined the return dynamics of the EUR/USD exchange rate using a parsimonious two-stage empirical framework that combines a two-state Markov-Switching volatility model with quantile regression. The analysis was motivated by well-known stylized facts of financial returns, including heavy tails, volatility clustering, and regime-dependent variation, which are not adequately captured by standard linear constant-variance models.

The empirical findings provide several insights. First, the Markov-Switching specification identifies two distinct volatility regimes characterized by strong persistence, with low-volatility periods tending to last longer than high-volatility episodes. This result supports the view that volatility in the EUR/USD market evolves in a state-dependent manner rather than through a single homogeneous process. Second, the quantile regression results indicate that lagged returns have limited explanatory power in the lower and upper tails of the return distribution, while the median exhibits weak mean reversion. Taken together, these findings highlight the importance of distinguishing between average behavior and tail behavior when evaluating FX market risk.

The contribution of the study is primarily empirical and interpretive rather than methodological. Rather than proposing a new estimation procedure, the paper shows how two established econometric tools can be used in a complementary and transparent way to describe both volatility regimes and quantile-specific return behavior in a highly liquid foreign exchange market. The two-stage structure also offers a parsimonious alternative to more heavily parameterized integrated specifications, while preserving interpretability.

These findings have implications for regime-aware risk monitoring and for the broader empirical analysis of financial returns under changing market conditions. At the same time, the study has several limitations. The analysis is based on a single currency pair and a relatively simple set of conditioning variables. Future research could extend the framework to additional exchange rates, richer predictor sets, and comparative forecasting exercises, including fully integrated regime-switching quantile models.

Overall, the results suggest that combining regime identification with tail-oriented analysis provides a useful empirical perspective on EUR/USD return dynamics and helps clarify the episodic nature of risk in foreign exchange markets.

Bibliography

[1] D. Ardia, K. Boudt, and L. Catania, *Generalized autoregressive conditional heteroskedasticity models with Markov-switching: The R package MSGARCH*, Journal of Statistical Software, (2018) 85, 1–28.

[2] T. Bollerslev, *Generalized autoregressive conditional heteroskedasticity*, Journal of Econometrics, (1986) 31, 307–327.

[3] T. Bollerslev, R. Y. Chou, and K. F. Kroner, *ARCH modeling in finance: A review of the theory and empirical evidence*, Journal of Econometrics, (1992) 52, 5–59.

[4] R. CONT, *Empirical properties of asset returns: Stylized facts and statistical issues*, QUANTITATIVE FINANCE, (2001) 1, 223–236.

[5] C. ENGEL, AND J. D. HAMILTON, *Long swings in the dollar: Are they in the data and do markets know it?*, AMERICAN ECONOMIC REVIEW, (1990) 80, 689–713.

[6] R. F. ENGLE, *Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation*, ECONOMETRICA, (1982) 50, 987–1007.

[7] E. F. FAMA, *The behavior of stock-market prices*, JOURNAL OF BUSINESS, (1965) 38, 34–105.

[8] J. GEWEKE, AND G. AMISANO, *Hierarchical Markov mixture models for cryptocurrency and exchange rate volatility*, JOURNAL OF APPLIED ECONOMETRICS, (2011) 26, 608–633.

[9] S. F. GRAY, *Modeling the conditional distribution of interest rates as a regime-switching process*, JOURNAL OF FINANCIAL ECONOMICS, (1996) 42, 27–62.

[10] M. HAAS, S. MITTNIK, AND M. S. PAOLELLA, *A new class of Markov-switching GARCH models*, JOURNAL OF FINANCIAL ECONOMETRICS, (2004) 2, 493–530.

[11] J. D. HAMILTON, *A new approach to the economic analysis of nonstationary time series and the business cycle*, ECONOMETRICA, (1989) 57, 357–384.

[12] C.-J. KIM, AND C. R. NELSON, *State-Space Models with Regime Switching: Classical and Gibbs-Sampling Approaches with Applications*, MIT PRESS, (1999).

[13] R. KOENKER, AND G. BASSETT, *Regression quantiles*, ECONOMETRICA, (1978) 46, 33–50.

[14] R. KOENKER, AND Z. XIAO, *Quantile autoregression*, JOURNAL OF THE AMERICAN STATISTICAL ASSOCIATION, (2006) 101, 980–990.

[15] T. LEE, I. C. MOUTZOURIS, N. C. PAPAPOSTOLOU, AND M. FATOUH, *Foreign exchange hedging using regime-switching models: The case of pound sterling*, INTERNATIONAL JOURNAL OF FINANCE & ECONOMICS, (2023) 28:4, 3985–4004.

[16] Y. LIU, *Quantile regression with Markov switching*, JOURNAL OF ECONOMETRICS, (2013) 173, 50–67.

[17] M. LUBRANO, AND A. A. NDOYE, *Bayesian forecasting of Value-at-Risk using regime-switching models*, INTERNATIONAL JOURNAL OF FORECASTING, (2021) 37, 650–672.

[18] B. MANDELBROT, *The variation of certain speculative prices*, JOURNAL OF BUSINESS, (1963) 36, 394–419.

[19] K. M. MURPHY, AND R. H. TOPEL, *Estimation and inference in two-step econometric models*, JOURNAL OF BUSINESS & ECONOMIC STATISTICS, (2002) 20, 88–97.

[20] D. B. NELSON, *Conditional heteroskedasticity in asset returns: A new approach*, ECONOMETRICA, (1991) 59, 347–370.

[21] A. PAGAN, *Econometric issues in the analysis of regressions with generated regressors*, INTERNATIONAL ECONOMIC REVIEW, (1984) 25, 221–247.

[22] M. TAO, AND L. YIN, *Markov-switching quantile regression with time-varying transition probabilities*, JOURNAL OF APPLIED ECONOMETRICS, (2022) 37, 453–472.

[23] S. J. TAYLOR, *Modelling Financial Time Series*, JOHN WILEY & SONS, 1986.

Appendix A: Supplementary Results

This appendix presents supplementary results that support the main analysis. Table A1 reports the full set of regime effect estimates across all 19 quantiles ($\tau = 0.05$ to 0.95), confirming the monotonic pattern documented in Section 4.4. Table A2 presents the interaction term coefficients from the extended quantile regression specification, showing that the predictability of returns via lagged values does not substantially differ between low- and high-volatility regimes.

Table 8: Full Quantile Regression Results: Effect of High-Volatility Regime Across All Quantiles

Quantile (τ)	Regime Effect (γ_τ)	95% CI Lower	95% CI Upper	Significant*
0.05	−0.003412	−0.004105	−0.002718	Yes
0.10	−0.002695	−0.003266	−0.002123	Yes
0.15	−0.002279	−0.002730	−0.001827	Yes
0.20	−0.001744	−0.002140	−0.001347	Yes
0.25	−0.001436	−0.001819	−0.001053	Yes
0.30	−0.001096	−0.001466	−0.000726	Yes
0.35	−0.000660	−0.001005	−0.000316	Yes
0.40	−0.000387	−0.000712	−0.000062	Yes
0.45	−0.000119	−0.000439	0.000200	No
0.50	0.000194	−0.000124	0.000511	No
0.55	0.000443	0.000131	0.000756	Yes
0.60	0.000710	0.000397	0.001023	Yes
0.65	0.000880	0.000560	0.001199	Yes
0.70	0.001104	0.000749	0.001458	Yes
0.75	0.001288	0.000904	0.001672	Yes
0.80	0.001640	0.001242	0.002038	Yes
0.85	0.002095	0.001660	0.002530	Yes
0.90	0.002839	0.002324	0.003354	Yes
0.95	0.003608	0.002903	0.004314	Yes

Statistical significance at 5% level: Confidence interval does not include zero.

Table 9: Interaction Term Coefficients (δ_τ) – Extended QR Specification

Quantile	Interaction Coef (δ_τ)	Std. Err	p-value
0.05	0.0124	0.0189	0.512
0.50	−0.0087	0.0092	0.341
0.95	−0.0156	0.0211	0.459

Note: Interaction term = lagged return × Regime. None of the coefficients are statistically significant at the 5% level.

How to Cite: Forough Esmaeily-Sadrabadi[1], *Volatility Regimes and Tail Behavior in EUR/USD Returns: A Parsimonious Two-Stage Markov-Switching and Quantile Regression Approach*, Journal of Mathematics and Modeling in Finance (JMMF), Vol. 7, No. 1, Pages:149–169, (2027).