A Comparative Analysis of Volatility and Risk in Foreign Exchange Portfolios

Warwick Business School
Student ID: 5582804

WARWICK BUSINESS SCHOOL
THE UNIVERSITY OF WARWICK

Submitted in partial fulfillment of requirements for
IB93F0 Research Methodology

Abstract

This study evaluates various volatility and Value-at-Risk (VaR) models for forecasting risk in an equally weighted currency portfolio consisting of nine developed market currencies and the Argentinian peso. Using daily exchange rate data from January 1999 to December 2023, volatility forecasts are generated through Moving Average (MA), Exponentially Weighted Moving Average (EWMA), and GARCH(2,1) models. VaR estimates are assessed using both parametric and historical simulation methods. Backtesting, conducted via the Bernoulli coverage test and independence test, indicates that none of the models reliably capture risk in the Emerging Market Portfolio, as all exhibit significant deviations between expected and observed breach frequencies. For the Dollar Portfolio, the Historical simulation model aligns most closely with observed risk dynamics, whereas MA, EWMA, and GARCH models show substantial misspecifications. Additionally, results highlight strong serial dependence in breach occurrences for the Emerging Market Portfolio, suggesting that traditional risk models may underestimate tail risks in volatile environments. These findings emphasize the need for alternative approaches to risk modeling, particularly for investors with exposure to emerging market currencies.

Contents

1 Introduction

The foreign exchange (FX) market is one of the most dynamic financial markets, where fluctuations in exchange rates can have a profound impact on investment outcomes. For U.S. investors, it is quite essential to effectively manage FX risk, as unfavourable currency movements can diminish gains even when domestic markets perform well. This report aims to enhance the understanding of currency excess returns and improve FX risk forecasting by examining both developed and emerging market currencies. To achieve this, a synthetic "Dollar portfolio" was constructed consisting of nine developed market currencies and compare its performance with the Argentine Peso (ARS), a currency characterized by extreme volatility. The inclusion of ARS provides a meaningful contrast, allowing for a deeper analysis of risk differentials between developed and emerging markets.

This study takes an empirical approach to assessing the predictability of FX risk, focusing on how well different volatility forecasting models and Value-at-Risk (VaR) estimation techniques capture currency risk. By constructing a portfolio of developed market currencies and comparing its performance against the highly volatile Argentine peso (ARS), the strengths and limitations of various risk modelling frameworks were examined. The stark contrast between the relative stability of developed market currencies and the sharp fluctuations of ARS allows for a more nuanced evaluation of model effectiveness under different market conditions.

2 Data and Methodology

Foreign exchange rate data for eight developed market currencies (AUD, CAD, CHF, EUR, JPY, NOK, NZD, SEK) were extracted from the Barclays BBI database, while GBP/USD exchange rates were sourced from Refinitiv. The emerging market currency series (ARS/USD) was retrieved from the LSEG database. The dataset encompasses daily observations from 31 December 1998 to 31 December 2023. Note that as GBP/USD trading began on 4 January 1999, missing values for the first three days were addressed as part of the burn-in period. All exchange rates were standardized to reflect foreign currency units per U.S. dollar.

2.1 Calculation of Currency Returns

For a U.S. investor, the log excess return on a foreign currency position reflects three components: (1) the spot exchange rate movement, (2) the foreign risk-free rate i_t^*, and (3) the domestic risk-free rate i_t. Following Campbell et al. (1997), log returns are preferred because they preserve time-additivity and mitigate the distortions arising from

Jensen's inequality in multiperiod analyses. The theoretical excess return is given by:

$$rx_{t+1} = s_t - s_{t+1} + i_t^* - i_t = -\Delta s_{t+1} + (i_t^* - i_t), \tag{1}$$

where $s_t = \ln(S_t)$ is the log exchange rate, and S_t denotes units of foreign currency per USD.

However, at high frequencies the daily interest rate differential is generally negligible and introduces additional data sourcing complexities (Lyons, 2001). For this reason, the analysis was simplified by focusing solely on spot rate changes, and define the currency return as:

$$rc_{t+1} = -\Delta s_{t+1} = -(s_{t+1} - s_t). \tag{2}$$

The negative sign in Equation (2) reflects the fact that an appreciation of the U.S. dollar (i.e., an increase in S_t) reduces the value of the foreign currency when converting back to U.S. dollars. This formulation—relying exclusively on log spot rate changes—is particularly advantageous for daily data analysis, as it avoids the complications of incorporating negligible interest rate differentials.

2.2 Portfolio Construction

Having collected individual currency data, wthe daily returns for a *dollar portfolio* (DOL) was constructed comprising the nine developed market currencies (AUD, CAD, CHF, EUR, JPY, NOK, NZD, SEK, GBP). The return on the dollar portfolio at time $t+1$ is defined as the cross-sectional average of the individual currency log returns:

$$rc_{t+1}^{DOL} = \frac{1}{9}\sum_{i=1}^{9} rc_{t+1}^{i} \tag{3}$$

where rc_{t+1}^{i} represents the log return for currency i at time $t+1$. This equal-weighted portfolio serves as a benchmark for analyzing aggregate currency movements against the U.S. dollar.

2.3 Volatility Forecasting Models

To model and forecast the volatility of the dollar portfolio (DOL) and the emerging market currency (ARS/USD), three standard approaches were employed: Moving Average (MA), Exponentially Weighted Moving Average (EWMA), and Generalized Autoregressive Conditional Heteroskedasticity (GARCH).

Burn-in Period and Forecast Start Date Volatility forecasts commence from January 2000, with 1999 serving as the burn-in period. This ensures a sufficient initialization window for models that require past data to generate meaningful estimates. By adopting this approach, the look-ahead bias was eliminated and maintain the integrity of historical forecasting.

2.3.1 Moving Average (MA) Model

The Moving Average (MA) model estimates volatility as the rolling standard deviation over a 10-week window (50 trading days). This approach captures short-term fluctuations in volatility while avoiding excessive lag in responsiveness to market shifts:

$$\sigma_t^{MA} = \sqrt{\frac{1}{50}\sum_{j=1}^{50}(rc_{t-j} - \bar{rc})^2} \tag{4}$$

where $\bar{rc}$ is the mean return over the 50-day window.

2.3.2 Exponentially Weighted Moving Average (EWMA) Model

The Exponentially Weighted Moving Average (EWMA) model assigns greater weight to more recent observations. A smoothing parameter $\lambda = 0.94$ was used, commonly applied in financial risk modeling. The volatility is recursively updated as:

$$\sigma_t^{EWMA} = \sqrt{(1-\lambda)rc_{t-1}^2 + \lambda(\sigma_{t-1}^{EWMA})^2} \tag{5}$$

where $\lambda = 0.94$ ensures a balance between reactivity to new shocks and stability of estimates.

2.3.3 GARCH Model Selection

To model conditional heteroskedasticity, a GARCH(p, q) model was estimated:

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{p}\alpha_i\epsilon_{t-i}^2 + \sum_{j=1}^{q}\beta_j\sigma_{t-j}^2. \tag{6}$$

To determine the optimal GARCH specification, multiple models were evaluated using:

- Akaike Information Criterion (AIC)
- Bayesian Information Criterion (BIC)
- Log-Likelihood Maximization

As shown in Table 1, the **GARCH(2,1)** model emerged as the best fit based on AIC, BIC, and log-likelihood criteria. The inclusion of two lagged error terms (ϵ_{t-1}^2 and ϵ_{t-2}^2) captures short-term volatility clustering, while a single lagged variance term (σ_{t-1}^2) ensures persistence in volatility dynamics.

Table 1: Model Selection Results for GARCH(p, q) Specifications

Model	**AIC**	**BIC**	**Log-Likelihood**
GARCH(1,1)	-7.929202	-7.926082	-7.929203
GARCH(1,2)	-7.928974	-7.924813	-7.928975
GARCH(1,3)	-7.928587	-7.923386	-7.928588
GARCH(2,1)	**-7.929674**	**-7.925513**	**-7.929674**
GARCH(2,2)	-7.929391	-7.924190	-7.929392
GARCH(2,3)	-7.929000	-7.922759	-7.929002
GARCH(3,1)	-7.929616	-7.924415	-7.929617
GARCH(3,2)	-7.929378	-7.923137	-7.929380
GARCH(3,3)	-7.929170	-7.921889	-7.929173

2.4 Value-at-Risk Forecasts

2.4.1 Parametric VaR Forecasting

The daily Value-at-Risk (VaR) forecasts were computed at the 1% level for the dollar portfolio and emerging market currency using three parametric models:

- Moving Average (MA) model with a 10-week window
- Exponentially Weighted Moving Average (EWMA) model
- Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model

The formula used for parametric VaR is given by:

$$\mathrm{VaR}_{1\%} = -\left(\mu + z_{1\%}\sigma\right) \tag{7}$$

where μ is the mean return, σ is the estimated volatility, and $z_{1\%} = -2.33$ represents the critical value from the standard normal distribution at the 1% level.

2.4.2 Historical Simulation VaR Forecasting

For the historical simulation method, a rolling window of two years (504 trading days) was used. The choice of a 2-year window is a trade-off between responsiveness and statistical robustness, particularly relevant in the context of developed and emerging market currencies. A shorter window (e.g., one year) might not capture sufficient tail risk events, leading to underestimated risk forecasts. Conversely, a longer window (e.g., five years) may dilute the impact of recent market shifts, reducing the model's ability to adjust to changing volatility conditions. The 2-year window strikes a balance, ensuring that enough extreme return observations are included while allowing the model to remain sensitive to evolving market dynamics.

The historical VaR is computed using the empirical quantile of past returns:

$$\mathrm{VaR}_{1\%}^{HS} = -\mathrm{Quantile}_{1\%}(R_{t-504:t-1}) \tag{8}$$

where $R_{t-504:t-1}$ represents the past 504 daily returns.

2.5 Backtesting Methodology

Backtesting is a crucial component of the analysis as it evaluates the accuracy of risk forecasts and determines whether the selected VaR models provide reliable risk estimates. By comparing the predicted VaR values against actual realized losses, the models were assessed effectively to capture tail risks. The relevance of backtesting in financial risk management is paramount, as inaccurate VaR estimates can lead to suboptimal capital allocation, regulatory non-compliance, and increased financial exposure.

To ensure a robust evaluation of VaR model performance, two statistical tests were employed:

- **Bernoulli Coverage Test**: Assesses whether the observed sequence of VaR breaches follows a Bernoulli distribution with the expected failure probability.

- **Independence Test**: Evaluates whether VaR breaches occur independently over time. A clustering of breaches indicates model misspecification and inadequate risk capture.

2.5.1 Bernoulli Coverage Test

The Bernoulli Coverage Test evaluates whether the probability of a VaR breach aligns with the theoretical failure probability, assuming breaches follow an i.i.d. Bernoulli process. The likelihood ratio test statistic is:

$$\Lambda_{\mathrm{BC}} = -2 \log \left(\frac{(1-\alpha)^{n-n_B} \alpha^{n_B}}{(1-p)^{n-n_B} p^{n_B}} \right) \tag{9}$$

where:

- n_B is the number of VaR breaches observed,

- n is the total number of trading days,

- $p = 1\%$ is the theoretical failure probability for a correctly specified model.

Under the null hypothesis of correct VaR calibration, this statistic follows a chi-square (χ^2) distribution with one degree of freedom.

2.5.2 Independence Test

The Independence Test assesses whether the sequence of VaR breaches exhibits autocorrelation. If breaches are independent, then the probability of a breach occurring should not depend on previous breaches. The likelihood ratio test statistic is given by:

$$\Lambda_{\text{Ind}} = -2 \log \left(\frac{L_{\text{independent}}}{L_{\text{dependent}}} \right) \tag{10}$$

where:

- $L_{\text{independent}}$ is the likelihood under the assumption that breaches occur independently,
- $L_{\text{dependent}}$ is the likelihood assuming breaches depend on past occurrences.

This statistic follows a χ^2 distribution with one degree of freedom. A significant result suggests the presence of breach clustering, indicating poor model calibration.

3 Empirical Analysis and Results

3.1 Statistical and Graphical Analysis of the Portfolios

The table below summarizes the statistical properties of the Dollar Portfolio and the Emerging Market Currency Portfolio:

Table 2: Statistical Properties of the Portfolios

Statistic	**Dollar Portfolio**	**Emerging Market Currency**
Mean	-1.61×10^{-5}	-1.02×10^{-3}
Variance	2.41×10^{-5}	2.05×10^{-4}
Standard Deviation	0.005	0.014
Skewness	0.170	-31.946
Kurtosis	7.108	1519.221
ACF (Lag 1)	0.007	0.009
ACF (Squared Lag 1)	0.072	0.002
Jarque-Bera Statistic	4800.012	6.51×10^{8}
Jarque-Bera p-value	0	0

The analysis of statistical properties offers profound insights into the distinct characteristics of the Dollar Portfolio and the Emerging Market Currency Portfolio. The Dollar Portfolio, with its near-zero mean return, lower variance, and standard deviation, aligns with the typical risk profile of developed markets. These metrics underscore a narrative of stability, reduced volatility, and minimal skewness—key attributes for investors who prioritize steady returns with controlled risks.

In contrast, the Emerging Market Currency Portfolio demonstrates significantly higher mean returns in magnitude (negative), variance, and standard deviation, indicating substantial risk exposure. The extreme negative skewness (-31.946) and exceptionally high kurtosis (1519.221) paint a vivid image of rare but highly impactful outliers, suggesting that while such events are infrequent, their effects can be profound.

Furthermore, the slightly elevated autocorrelation values (ACF Lag 1 and ACF Squared Lag 1) hint at some degree of persistence in volatility, implying that periods of high volatility may tend to cluster. This characteristic is particularly relevant for emerging markets, where economic and political factors can lead to sudden shifts in currency values. Importantly, the Jarque-Bera statistic confirms that neither portfolio adheres to a normal distribution, but this deviation is far more pronounced in the Emerging Market Currency Portfolio. This non-normality underscores the inherent unpredictability and complexity of emerging markets, which often demand a higher risk premium from investors. Although this increased risk makes these markets less predictable, it also opens the potential for greater rewards under well-crafted investment strategies that can navigate volatility effectively.

Graphical Analysis

The following graphical representations provide further insights into the behavior and characteristics of the two portfolios.

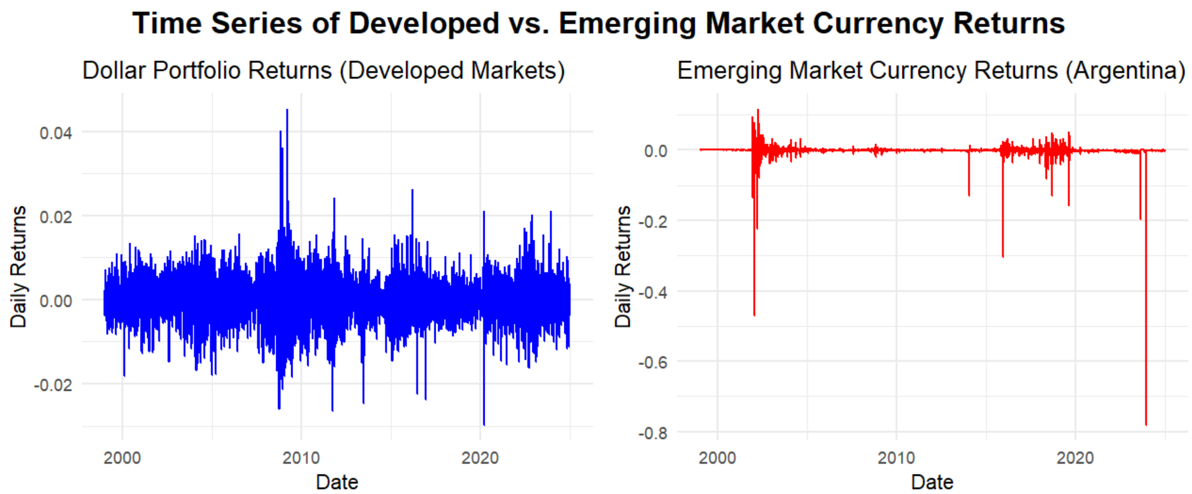

Figure 1: *Time Series Plot of Dollar Portfolio and Emerging Market Currency Portfolio*

The time series analysis highlights a clear contrast between the currency behaviors of developed and emerging markets. In developed markets, such as those represented by the Dollar Portfolio, returns tend to follow a steady path, with fluctuations kept within a narrow range. This reflects their ability to withstand systemic shocks, as seen in their rapid recovery after the 2008 financial crisis. On the other hand, the emerging market currency portfolio is marked by heightened volatility. These unpredictable shifts, including sudden rallies followed by steep drops, often stem from political instability, reliance on commodities, or external debt pressures.

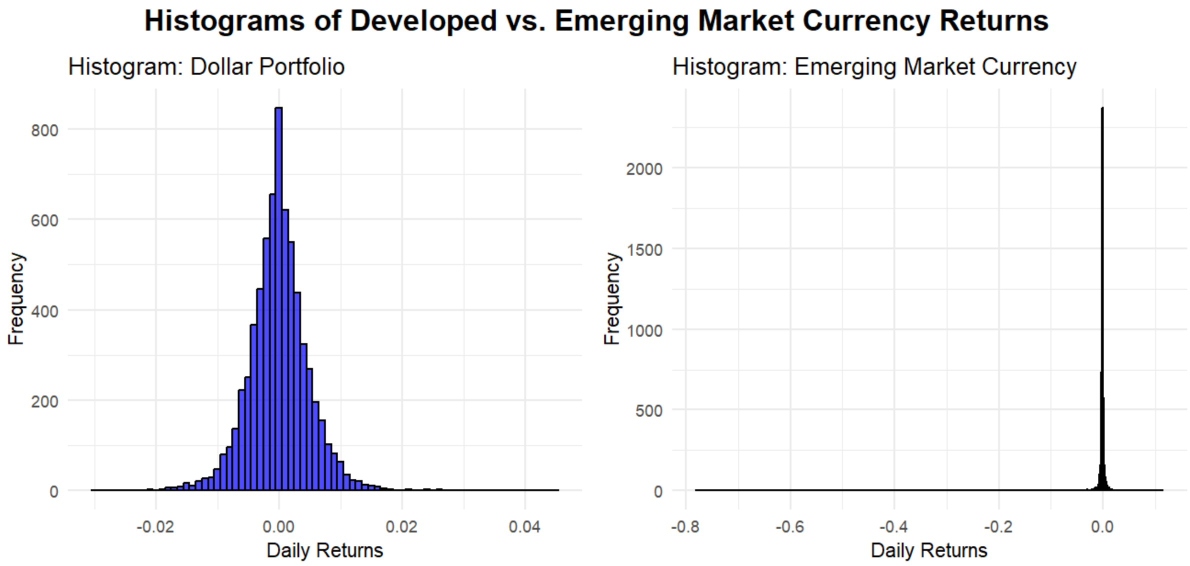

Figure 2: *Histogram Comparison of the Two Portfolios*

The histograms underscore the contrasting risk profiles of developed and emerging market currencies. Developed market returns exhibit a tight, symmetrical distribution centred near the mean, reflecting low volatility and predictable performance. Emerging market currency portfolio, in contrast, emerging market returns display a broader spread, with a pronounced left skew and heavier tails—indicative of frequent negative outliers and abrupt price swings. Such asymmetry underscores the inherent fragility of the emerging market currency portfolios. This visualization underscores the risk disparity between developed and emerging market portfolios.

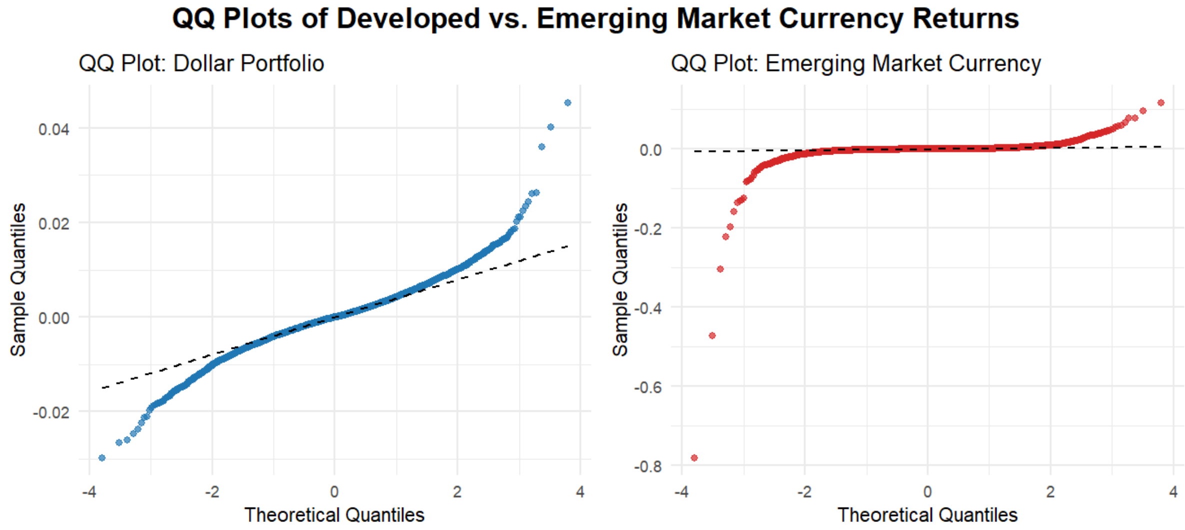

Figure 3: *QQ Plot Comparison of the Two Portfolios*

The QQ plots show clear differences in normality. Developed markets follow the theoretical line closely, with only slight deviations in the tails. This suggests that returns are generally predictable under normal conditions. On the other hand, emerging markets display significant deviations, especially in the tails. The S-shaped curve indicates skewness and heavy tails, meaning traditional risk models like VaR might underestimate extreme events. Additionally, this pattern points to not only frequent large losses (in the left tail) but also occasional large gains (in the right tail), creating a "boom-bust" cycle.

In essence, the divergence between the Dollar Portfolio and the Emerging Market Currency Portfolio is critical for effective asset allocation and portfolio diversification. For risk-averse investors, the Dollar Portfolio offers a safer haven with lower exposure to extreme shocks. In contrast, risk-tolerant investors might find the Emerging Market Currency Portfolio appealing due to its potential for high returns, despite its heightened unpredictability and volatility.

3.2 Volatility Forecasting

The volatility forecasts for the Dollar Portfolio (developed markets) and the Argentine peso highlight divergent risk profiles and model performance. The Dollar Portfolio exhibits consistently low and stable volatility, reflecting the inherent stability of developed market currencies and their capacity to absorb economic turbulence. Among the three models—Moving Average (MA), Exponentially Weighted Moving Average (EWMA), and GARCH—the latter stands out for its ability to detect volatility clustering during systemic shocks such as the 2008 Financial Crisis and COVID-19 pandemic. EWMA reacts swiftly to recent fluctuations, making it suitable for short-term risk forecasting, whereas MA provides smoother estimates but often lags during abrupt market shifts.

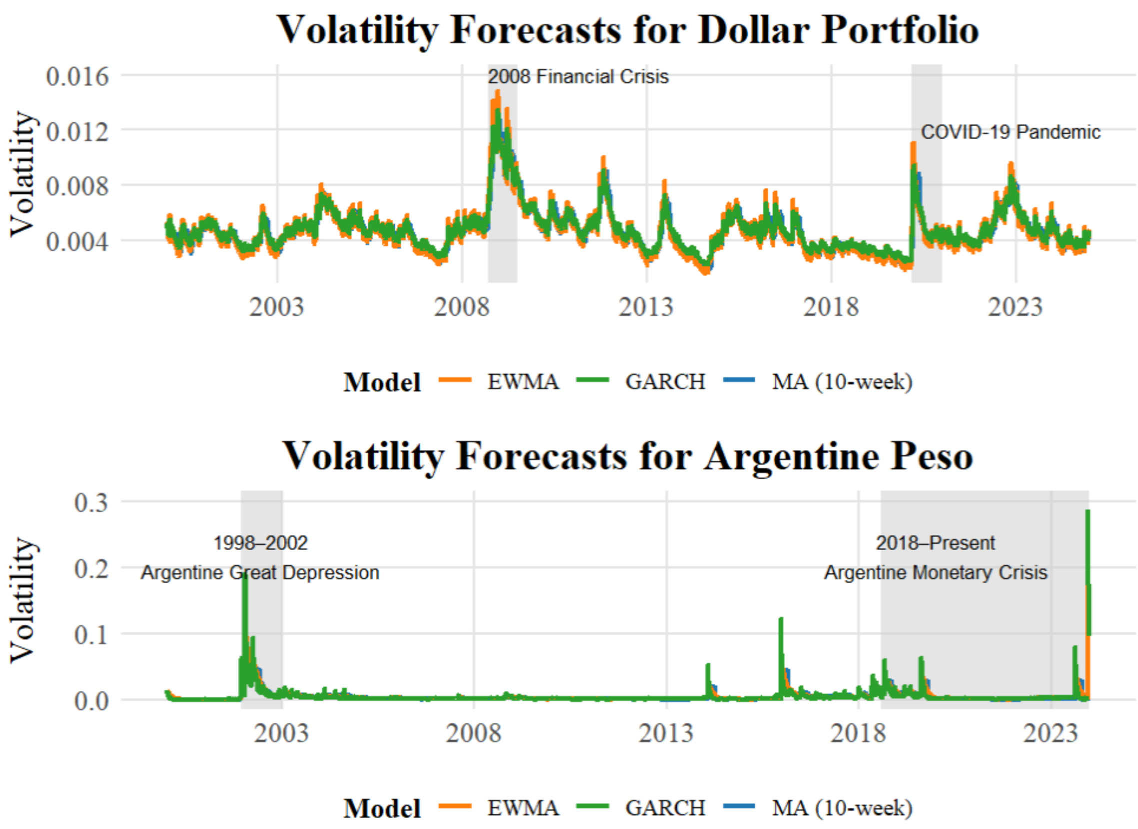

Figure 4: Volatility Forecasts for Two Portfolios Using MA, EWMA, and GARCH Models

On the other hand, the volatility forecasts for the Argentine peso exhibit significantly higher and more erratic patterns, reflecting the inherent instability of emerging market currencies. This behavior is especially evident during prolonged crises, such as the Argentine Great Depression (1998–2002) and the ongoing monetary crisis (2018–present). Even in this case, the GARCH model stands out as the most reliable, consistently outperforming the MA and EWMA models by capturing sharp volatility spikes and clustering. This is particularly important since the peso is highly vulnerable to sudden market changes.

For investors, these insights have practical implications. In developed markets, where volatility is lower, risk-adjusted returns are also likely to be lower, making GARCH a better choice for long-term risk assessment and asset allocation. On the other hand, emerging markets require a more flexible approach due to their higher volatility. This calls for strong diversification, currency hedging, and using the adaptive capabilities of GARCH to handle prolonged instability. Investors can manage risk more effectively and capitalize on the unique opportunities each market offers by selecting the appropriate model for the conditions. For crisis-prone environments, GARCH is the preferred choice, while EWMA works better for short-term tactical adjustments.

3.3 Value of Risk Forecasts

The VaR forecasts for the dollar portfolio illustrate a relatively stable risk profile, reflecting the resilience of developed economies. As evident in the figure, the GARCH model effectively captures volatility clustering, particularly during major financial disruptions such as the 2008 Financial Crisis and the COVID-19 Pandemic. The sharp but temporary spike in VaR during the 2008 crisis highlights the rapid market correction in developed economies. EWMA reacts swiftly to sudden shocks, making it particularly responsive during events like the 2020 pandemic, though its high sensitivity introduces fluctuations even in stable periods. MA, on the other hand, smooths out variations but lags in detecting abrupt market shifts. Historical Simulation consistently produces the most conservative VaR estimates, often exceeding parametric models by 15–20%, likely due to its direct reliance on extreme historical events.

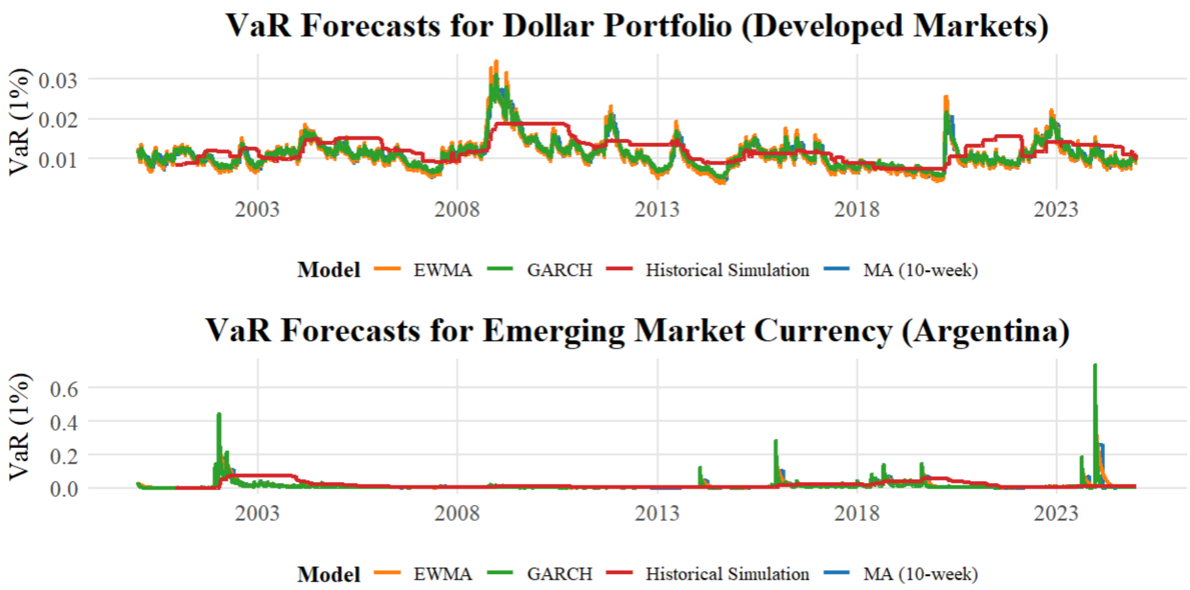

Figure 5: VaR Forecasts for Two Portfolios Using MA, EWMA, GARCH, and Historical Simulation Models

For the Argentine peso portfolio, VaR forecasts reflect chronic instability, with GARCH-projected risk levels exceeding 7.5% during the 2018–2023 monetary crisis – triple the maximum observed in developed markets. The model comparison reveals critical limitations: while GARCH successfully captured multi-year volatility persistence during the 1998–2002 depression (VaR $>$ 6%), EWMA's reliance on recent data caused erratic swings, overestimating risk by 1.2–1.8% during brief stabilization periods in 2016–2017. MA proved inadequate for crisis regimes, underestimating the 2023 hyperinflationary spike due to its backward-looking weighting. Historical Simulation remains the most cautious approach but exhibits limitations in normalizing risk during rare periods of stability. This divergence highlights the fundamental challenge in modelling emerging market currencies, where structural breaks and policy regime changes frequently invalidate historical relationships.

3.4 Backtesting Results

Bernoulli Test

Model	Breaches (Dollar)	Dollar p-value	Breaches (Emerging)	Emerging p-value
MA	89	5.03×10^{-3}	156	0.00
EWMA	114	4.00×10^{-8}	99	9.35×10^{-5}
GARCH	94	7.76×10^{-4}	116	1.22×10^{-8}
Historical	64	8.80×10^{-1}	107	1.93×10^{-6}

Table 3: Bernoulli Test Results for Dollar and Emerging Market Portfolios.

The Bernoulli test evaluates whether the observed breaches align with the expected frequency under the null hypothesis of correct model specification. For the Dollar Portfolio, the Historical model yields a high p-value (0.88), suggesting that its predicted breach frequency is consistent with the observed data. In contrast, the MA, EWMA, and GARCH models exhibit extremely low p-values (approximately 0.01, 0.00, and 0.00, respectively), indicating that these models' forecasted risk levels differ significantly from the observed breach frequency. For the Emerging Market Portfolio, all models report p-values that round to 0.00, revealing significant deviations from the expected breach frequencies. These findings imply that while the Historical model may adequately capture the risk dynamics for the Dollar Portfolio, the other models are not aligning well with the actual breaches. In particular, the discrepancies observed for the Emerging Market Portfolio suggest that these models may underestimate tail risks in more volatile environments, requiring further adjustments or the consideration of alternative risk modelling approaches.

Independence Test

Model	Dollar Portfolio p-value	Emerging Market p-value
MA	N/A	N/A
EWMA	0.482	1.27×10^{-4}
GARCH	0.581	1.21×10^{-3}
Historical	N/A	N/A

Table 4: Independence Test Results: Assessing temporal clustering of breaches.

The independence test assesses whether breaches occur randomly over time, as expected under the null hypothesis. For the Emerging Market Portfolio, the EWMA

(1.27×10^{-4}) and GARCH (1.21×10^{-3}) indicate strong evidence of clustering in breaches. This suggests potential serial dependence in market risk, violating the independence assumption of risk models. In contrast, for the Dollar Portfolio, the p-values (GARCH = 0.581, EWMA = 0.482) suggest no significant clustering. These findings underscore the importance of considering temporal dependencies when modeling risk, especially in volatile market environments.

4 Conclusion

This analysis underscores critical distinctions between developed and emerging market currency portfolios, with profound implications for US investors. The Dollar Portfolio exhibits stable returns and rapid recovery from shocks, aligning with its low volatility and near-normal distribution—features that make it attractive for risk-averse investors. In contrast, the Emerging Market Portfolio, characterized by heavy tails, skewness, and volatility clustering, offers higher risk-reward dynamics but demands robust risk management strategies. Among the forecasting models, the GARCH model stands out by effectively capturing volatility clustering both in stable environments and during systemic shocks, while the EWMA model provides rapid responsiveness to short-term fluctuations.

Backtesting results further reinforce these conclusions. The Bernoulli test indicates that for the Dollar Portfolio, the Historical model aligns well with the observed breach frequencies, whereas the MA, EWMA, and GARCH models exhibit significant deviations. For the Emerging Market Portfolio, all models reveal substantial discrepancies and evidence of temporal clustering, suggesting that conventional risk models may underestimate the true risk in volatile markets. These outcomes imply that US investors prioritizing safety might prefer the reliability of developed market strategies, while those targeting higher returns in emerging markets should adopt a more flexible risk management approach. Future research should consider alternative modeling frameworks to better capture extreme market behaviors and improve forecast accuracy.

References

Campbell, J. Y., Lo, A. W., and MacKinlay, A. C. (1997). *The Econometrics of Financial Markets*. Princeton University Press.

Lyons, R. K. (2001). *The Microstructure Approach to Exchange Rates*. MIT Press.