GARCH-FX: A Modular Framework for Stochastic and Regime-Aware GARCH Forecasting

Tony Paul, Nitin

10 July 2025

GARCH-FX: A Modular Framework for Stochastic and Regime-Aware GARCH Forecasting*

Nitin Tony Paul
nitintonypaul@gmail.com

July 2025

Abstract

Traditional GARCH models, while robust, are deterministic and their long-horizon forecasts converge to a static mean, failing to capture the dynamic nature of real markets. Conversely, classical stochastic volatility models often introduce significant implementation and calibration complexity. This paper introduces **GARCH-FX** (GARCH Forecasting eXtension), a novel and accessible framework that augments the classic GARCH model to generate realistic, stochastic volatility paths without this prohibitive complexity.

GARCH-FX is built upon the core strength of GARCH—its ability to estimate long-run variance—but replaces the deterministic multi-step forecast with a stochastic simulation engine. It injects controlled randomness through a Gamma-distributed process, ensuring the forecast path is non-smooth and jagged. Furthermore, it incorporates a modular regime-switching multiplier, providing a flexible interface to inject external views or systematic signals into the forecast's mean level.

The result is a powerful and intuitive framework for generating dynamic long-term volatility scenarios. By separating the drivers of mean-level shifts from local stochastic behavior, GARCH-FX aims to provide a practical tool for applications requiring realistic market simulations, such as stress-testing, risk analysis, and synthetic data generation.

Keywords: Stochastic Volatility Forecasting, GARCH Extensions, Regime-Switching Volatility, Gamma-Distributed, Volatility, Volatility Forecast Uncertainty, Nonlinear GARCH Models, Stochastic Vol Forecast, Financial Time Series, Heteroskedasticity Dynamics, Gamma Noise in Volatility

*Version 2: Benchmarks improved, Gamma distribution clarified, and tests corrected.

Contents

1 Introduction

1.1 The Challenge of Volatility Forecasting

Volatility forecasting is foundational in modern finance, serving as a critical input for risk management, derivative pricing, and algorithmic trading strategies. The challenge for practitioners lies in selecting a model that balances accuracy with tractability. On one hand, the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) family of models offers simplicity and robustness. However, their deterministic forecast structure struggles over long horizons, where forecasts tend to "flatline" to a constant mean, rendering them inaccurate for capturing the persistent, shifting nature of market risk. On the other hand, more complex stochastic volatility models, such as the Heston model or Volterra processes, offer greater realism but introduce significant implementation and calibration complexity, making them less accessible for rapid prototyping and analysis.

1.2 The GARCH Family of Models

The GARCH framework, introduced by Bollerslev (1986), has become an industry benchmark for its ability to model volatility clustering. The canonical GARCH(1,1) model defines the conditional variance σ_t^2 as a function of its own recent history and past shocks:

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \tag{1}$$

Here, the variance at time t is a weighted sum of a long-run constant term (ω), the previous period's squared error or "news" (ϵ_{t-1}^2, weighted by α), and the previous period's forecasted variance (σ_{t-1}^2, weighted by β). Over the years, numerous extensions have been developed to capture more nuanced market behaviors. For instance, the Exponential GARCH (EGARCH) model accounts for the leverage effect, while the Fractionally Integrated GARCH (FIGARCH) model addresses long-memory properties in volatility. Despite these advances, the majority of these models retain a deterministic forecasting mechanism.

1.3 The Long-Horizon Problem

The primary limitation of this deterministic structure emerges in long-horizon forecasting. When projecting volatility multiple steps into the future, a GARCH model recursively relies on its own previous forecasts. This process inherently smooths out variance, causing the forecast to converge monotonically toward its unconditional long-run mean, $V_L = \omega/(1 - \alpha - \beta)$. This "flatlining" effect stands in stark contrast to the behavior of realized volatility in financial markets, which is characterized by jagged paths, persistent volatility clusters, and abrupt structural breaks rather than a smooth reversion to a single, static average.

1.4 Contribution: The GARCH-FX Framework

To address this gap, we propose **GARCH-FX**[1] (GARCH Forecasting eXtension)—a forecasting framework that maintains the well-understood structure of GARCH but enhances it with three key innovations designed to inject realistic, non-deterministic behavior:

- **Stochastic Evolution:** We replace the deterministic recursion with a **Gamma-distributed process**. This allows each future variance step to be drawn from a distribution, generating the realistic, noisy bursts and jagged paths absent in standard GARCH forecasts.

[1] The source code for the GARCH-FX framework is publicly available at: `https://github.com/nitintonypaul/GARCH-FX`

- **Volatility-of-Volatility Control:** The model introduces a tunable parameter, **Theta** (θ), which serves as the scale of the Gamma distribution. By adjusting θ, the user can control the dispersion of the forecast. While not the volatility-of-volatility itself, θ is directly **proportional to it**, allowing for intuitive control over the forecast's smoothness or jaggedness.
- **Regime Modulation:** We introduce a multiplier, **Delta (Δ)**, which is applied directly to the long-run variance component (ω). Rather than prescribing a specific regime-switching model, Δ serves as a modular framework extension point. This design empowers the user to integrate their own systematic logic for modulating the forecast's mean level, making the model agnostic to the source of the regime signal. The mechanism driving Δ can range from sophisticated models, such as Hidden Markov Models or functions of market entropy, to a simple, pre-defined sequence of regime states. This component can also be deactivated, allowing GARCH-FX to operate with purely stochastic evolution, thus providing maximum flexibility for different research and testing scenarios.

These modifications transform the GARCH forecast from a static, decaying path into a dynamic simulation engine. GARCH-FX can thus generate volatility paths that are more responsive to external conditions and better reflect the true, uncertain nature of financial markets, making it a powerful tool for stress-testing and scenario analysis.

2 Methodology: The GARCH-FX Framework

2.1 The GARCH(1,1) Model and its Deterministic Forecast

The foundation of our model is the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model, first proposed by Bollerslev (1986) as an extension of Engle's (1982) ARCH model. The GARCH(1,1) process has become a cornerstone of financial econometrics for its parsimonious yet powerful ability to capture volatility clustering. It models the conditional variance, σ_t^2, as follows:

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \tag{2}$$

where $\omega > 0$, $\alpha, \beta \geq 0$, and the stationarity condition $\alpha + \beta < 1$ is imposed. The model's parameters represent the long-run average variance (ω), the reaction to past shocks (α), and the persistence of volatility (β).

While effective for in-sample fitting, the model's structure reveals its primary limitation when used for out-of-sample forecasting. As future shocks ϵ_t are unknown, the multi-step forecast for the conditional variance at time $t + k$ becomes a deterministic recursion based on the previous forecast:

$$\hat{\sigma}_{t+k}^2 = \omega + (\alpha + \beta)\hat{\sigma}_{t+k-1}^2 \tag{3}$$

As the forecast horizon k increases, this equation causes the projected variance $\hat{\sigma}_{t+k}^2$ to monotonically converge to the unconditional long-run variance, $V_L = \frac{\omega}{1-\alpha-\beta}$. This convergence results in the characteristic **flatlining** of the forecast, producing a smooth path that fails to represent the jagged, stochastic behavior observed in real financial markets.

2.2 The GARCH-FX Stochastic Forecast Equation

The GARCH-FX framework is designed to solve this flatlining problem by building upon GARCH's core strength: its exceptional ability to capture the long-run variance of an asset. It is crucial to note that GARCH-FX is a multi-step forecasting engine that activates specifically when realized shocks are no longer available. For the 1-step-ahead forecast, the standard GARCH equation is superior as it incorporates the most recent squared error (ϵ_{t-1}^2). For all subsequent steps ($k > 1$), GARCH-FX takes over, replacing the deterministic recursion with a stochastic process.

We posit that while the deterministic forecast *path* is flawed, the long-run variance (V_L) implied by the GARCH parameters (ω, α, β) is a robust and meaningful baseline. GARCH-FX therefore uses this baseline not as a fixed destination, but as a dynamic center of gravity. It preserves the core GARCH parameters but introduces controlled randomness and regime awareness into the forecast evolution. Instead of projecting a single, decaying path, GARCH-FX generates a distribution of possible volatility paths that "dance" around this GARCH-derived baseline, thus honoring the underlying structure while reflecting real-world uncertainty.

The GARCH-FX forecast for variance at time $t + k$ is generated as follows:

$$\sigma^2_{t+k} = \omega \cdot \Delta_{t+k} + (\alpha + \beta) \cdot \tilde{\sigma}^2_{t+k-1} \tag{4}$$

$$\text{where } \tilde{\sigma}^2_{t+k-1} \sim \Gamma\left(\frac{\sigma^2_{t+k-1}}{\theta} + 1,\ \theta\right) \tag{5}$$

The key components of this formulation are:

- **Stochastic Innovation:** The deterministic term $\hat{\sigma}^2_{t+k-1}$ is replaced by $\tilde{\sigma}^2_{t+k-1}$, a random variable drawn from a Gamma distribution, $\Gamma(k, \theta)$. The shape parameter k is critically defined as $k = (\sigma^2_{t+k-1}/\theta) + 1$. This specific formulation ensures that the **mode** of the Gamma distribution is precisely the previous period's variance, σ^2_{t+k-1}. Consequently, the most likely value for the stochastic component is anchored to the last known state of the system, ensuring that the randomness is not arbitrary but is centered around the model's recent history. This injection of structured stochasticity allows the forecast path to be non-smooth and produce the variance bursts characteristic of real markets, while the use of the Gamma distribution inherently guarantees non-negative variance.

- **Regime-Awareness (Δ_t):** The constant term ω is multiplied by a time-varying regime modulator, Δ_{t+k}. This acts as an injection point for user-defined logic, allowing the long-run variance level to shift dynamically based on external signals or models.

- **Volatility-of-Volatility Control (θ):** The parameter θ is the scale of the Gamma distribution. It directly governs the dispersion of the stochastic component $\tilde{\sigma}^2_{t+k-1}$ around its mean. By adjusting this single parameter, the user can control the jaggedness of the forecast path, effectively tuning a proxy for the volatility-of-volatility.

This framework thus transforms the GARCH forecast from a static calculation into a dynamic simulation engine, capable of generating more realistic long-term volatility scenarios.

3 Formulation of the Stochastic Component

The selection of the Gamma distribution for the stochastic component, $\tilde{\sigma}^2_{t+k-1}$, is not arbitrary. It is a deliberate design choice central to the model's ability to realistically mimic the behavior of financial volatility, based on three key properties.

- **Non-Negativity:** The most fundamental requirement of a variance forecast is that it must be strictly non-negative. The Gamma distribution is defined on the support $[0, \infty)$, inherently enforcing this critical constraint without the need for truncation or other ad-hoc adjustments.

- **Right Skewness and Stylized Facts:** It is a well-documented stylized fact that financial volatility exhibits asymmetric behavior. It tends to spike upward rapidly during periods of market stress and then drift downward more slowly during periods of calm. This behavior is best captured by a right-skewed distribution. The Gamma distribution's inherent right skewness allows it to naturally model these large, infrequent upward bursts, providing a more realistic representation than a symmetric distribution like the Normal.

- **Analytical Tractability of the Mode:** A core principle of the GARCH-FX forecast is that each stochastic step should be anchored to the previous period's variance. The most intuitive way to achieve this is to set the previous variance, σ^2_{t+k-1}, as the *mode* (the most likely outcome) of the distribution for the next step, $\tilde{\sigma}^2_{t+k-1}$. The Gamma distribution makes this implementation exceptionally elegant.

  The mode of a Gamma distribution with shape k and scale θ is given by $(k-1)\theta$, for $k > 1$. By defining our shape parameter as:

$$k = \frac{\sigma^2_{t+k-1}}{\theta} + 1$$

  We ensure that the mode of the distribution for $\tilde{\sigma}^2_{t+k-1}$ is precisely σ^2_{t+k-1}:

$$\begin{aligned} \text{Mode} &= (k-1) \cdot \theta \\ &= \left(\left(\frac{\sigma^2_{t+k-1}}{\theta} + 1 \right) - 1 \right) \cdot \theta \\ &= \left(\frac{\sigma^2_{t+k-1}}{\theta} \right) \cdot \theta \\ &= \sigma^2_{t+k-1} \end{aligned}$$

  This property is not as easily or directly achievable with other common right-skewed distributions like the log-normal or Weibull, making the Gamma distribution uniquely suited to the autoregressive, path-dependent structure of our model.

4 Parameter Interpretation and Model Components

4.1 Delta (Δ): Regime-Sensitive Mean Modulation

The parameter Δ_t serves as a dynamic, time-varying scalar multiplier applied to the GARCH constant term, ω. This modification directly alters the long-run variance component of the forecast equation (3), effectively shifting the mean level or "center of gravity" around which the stochastic volatility process evolves. Its properties are distinct and intentional:

- **Regime Awareness:** Δ_t is designed to reflect the prevailing market regime. A value of $\Delta_t > 1$ elevates the mean volatility level, corresponding to a high-volatility or crisis period, while a value of $\Delta_t < 1$ suppresses it, reflecting a calm or quiet period. A value of $\Delta_t = 1$ represents the baseline GARCH regime.

- **Mean Level Shift:** Its primary function is to adjust the forecast's mean. It does not directly influence the local jaggedness or the "volatility of volatility," which is governed by a separate parameter, θ.

A key design principle of GARCH-FX is the modularity of the Δ_t component. The framework does not enforce a specific regime-detection model; instead, it provides an interface for the user to programmatically inject their own logic. This mechanism can be driven by a wide array of models, such as Hidden Markov Models (HMM), functions of market entropy, or Time-Varying Transition Probability (TVTP) models. Furthermore, this feature is entirely optional. Within our implementation, the regime modulation can be deactivated by setting $\Delta_t = 1$ for all t, allowing GARCH-FX to operate with purely stochastic evolution around the standard GARCH baseline.

4.1.1 Demonstration with a 3-State Markov Chain

For the purpose of demonstrating its capability within this paper, we implement a 3-state Markov chain to drive the evolution of Δ_t. This provides a structured yet dynamic approach to regime-switching, modeling transitions between low, normal, and high volatility states. The transition probability matrix is defined as:

$$\mathbf{P} = \begin{pmatrix} 0.970 & 0.029 & 0.001 \\ 0.015 & 0.950 & 0.035 \\ 0.000 & 0.040 & 0.960 \end{pmatrix}$$

Each of the three states is assigned a specific multiplier. The vector of multipliers, corresponding to "Low Volatility," "Normal," and "High Volatility" states, is:

$$\mathbf{\Delta}_S = [0.5, 1.0, 1.5]$$

The effect of this mechanism is starkly illustrated in Figure 1. Figure 1a shows the GARCH-FX forecast when the regime modulation is deactivated ($\Delta_t = 1$ for all steps). The forecast is stochastic but gravitates around a single, constant mean. In contrast, Figure 1b shows the forecast with the 3-state Markov chain enabled. The path clearly exhibits dynamic mean-reversion, shifting its central tendency as the underlying state transitions between different volatility regimes.

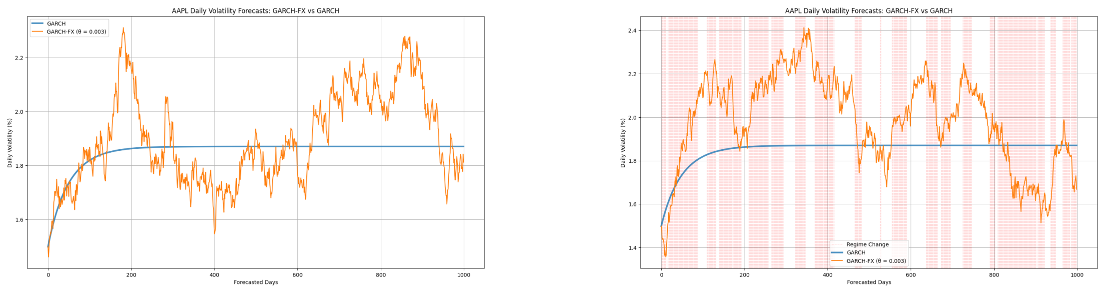

(a) Forecast with regime modulation deactivated. (b) Forecast with 3-state Markov chain activated.

Figure 1: Comparison of a GARCH-FX forecast without (left) and with (right) regime modulation. The right panel demonstrates dynamic mean-level shifting. Vertical red lines indicate transitions between distinct volatility regimes.

4.1.2 Showcasing the Potential of Delta

To showcase the full potential of the Δ_t parameter, Figure 2 presents a GARCH-FX forecast where the regime multipliers were manually adjusted ex-post to align with observed realized volatility. While this exercise does not represent a predictive model, it serves as a powerful proof-of-concept. It demonstrates that if a sufficiently accurate external signal for regime changes were available, the Δ_t mechanism provides the necessary tool to translate that signal into a highly accurate and responsive volatility forecast.

The generation of such a state-of-the-art signal could be approached through numerous avenues, targeting the complex and often non-linear drivers of market regimes. These include:

- **Advanced Econometric Models:** Such as Time-Varying Transition Probability (TVTP) models that allow the likelihood of regime changes to adapt based on market conditions.
- **Information-Theoretic Measures:** Using metrics like Shannon entropy calculated on market returns or order book data to quantify system-wide uncertainty and detect tipping points.
- **Machine Learning and AI:** Employing a range of techniques from classical classifiers to deep learning. This could involve Natural Language Processing (NLP) models to extract sentiment

from financial news, regulatory filings, or social media, or using sophisticated time-series models like LSTMs and Transformers to learn predictive patterns from a vast set of market, macroeconomic, and alternative data.

This highlights the unique power of separating the mean-level driver (Δ_t) from the stochastic engine (θ), allowing for targeted model enhancement. It effectively creates a plug-and-play architecture where GARCH-FX provides the core volatility dynamics, while external, potentially alpha-generating signals can be injected to guide its path.

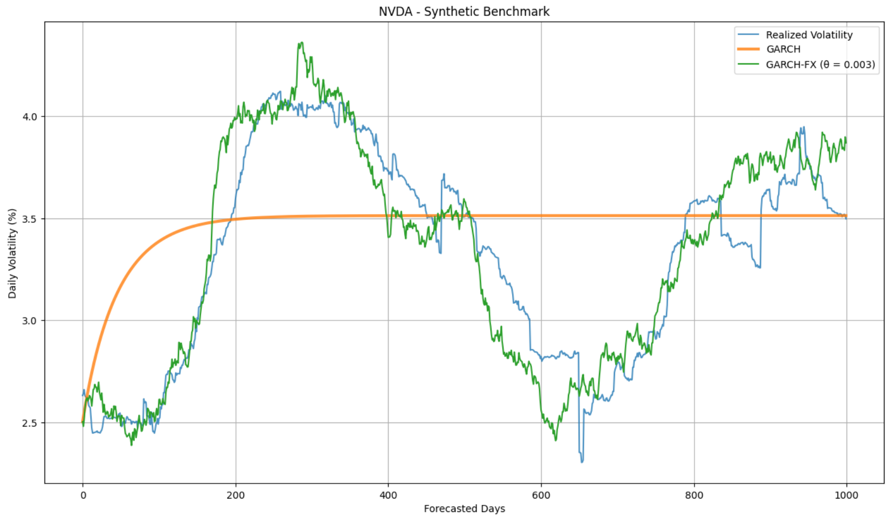

Figure 2: GARCH-FX forecast with manually adjusted Δ_t values (green line) to demonstrate its potential in tracking realized volatility (blue line). This illustrates the high degree of forecast control achievable with an accurate regime signal.

4.2 Theta (θ): Volatility-of-Volatility Control

While the Δ parameter governs the mean level of the forecast, the Theta (θ) parameter controls its local texture and erraticness. As the **scale** of the Gamma distribution in our model's forecast equation (4), θ directly influences the dispersion of the stochastic variance component. It can be intuitively understood as a "spikiness index" or a control for the forecast's jaggedness.

It is crucial to note that θ is not the volatility-of-volatility itself, but rather a parameter that is directly **proportional to it**. A larger θ results in a wider Gamma distribution, allowing for more aggressive, jump-prone volatility behavior, whereas a smaller θ constrains the forecast path to be smoother and stay closer to its GARCH-derived baseline.

4.2.1 Heuristic Application and Sensitivity

At present, a direct historical calibration technique for θ has not been developed; it is treated as a heuristic parameter. Our empirical tests show that its value has a predictable and intuitive effect on the forecast path:

- Values in the range of `1e-5` to `1e-4` produce a forecast that "hugs" the deterministic GARCH line, introducing minimal stochasticity.
- Values between `1e-3` and `5e-3` yield a balanced forecast with realistic, jagged behavior.
- Values approaching `1e-2` or higher can cause the forecast to "overshoot," generating a highly volatile path that may be suitable for stress-testing but less so for baseline forecasting.

This heuristic nature offers flexibility. For instance, if a user disables the regime-shifting Δ mechanism, a higher θ value can be selected. This is a riskier approach, but it allows the purely stochastic process to generate large variance spikes that might, by chance, capture periods of high realized volatility. Figure 3 demonstrates this sensitivity, showing how different fixed values of θ impact the character of the forecast.

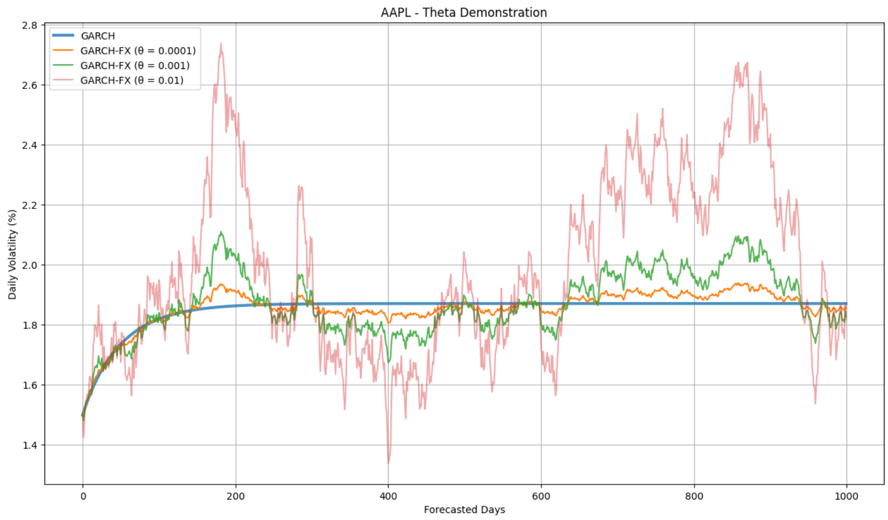

Figure 3: Sensitivity analysis of the GARCH-FX forecast to different fixed values of θ. Higher values clearly lead to a more jagged and dispersed volatility path.

4.2.2 The Potential of a Time-Varying Theta (θ_t)

While a fixed θ is effective, the true potential of this parameter lies in making it time-varying (θ_t). This would allow the model to adapt the "volatility of volatility" based on changing market dynamics, such as periods of high uncertainty followed by quiet consolidation.

To demonstrate the direct and immediate impact of such a parameter, we conduct a controlled experiment shown in Figure 4. Instead of attempting to align the forecast with a noisy realized volatility series, we programmatically alter the value of θ_t at pre-defined steps within a single forecast path for JPMorgan Chase (JPM). The parameter begins at a low value of `1e-4`, then steps up to `1e-3`, `1e-2`, and an aggressive `1e-1`, before reverting to `1e-4` near the end of the horizon. The effect on the forecast's texture is explicit and unambiguous: the path transitions from smooth, to moderately jagged, to extremely volatile, and back to calm, directly corresponding to the specified value of θ_t.

This controlled demonstration serves as a powerful illustration of the mechanism. In a practical application, the signal to drive a dynamic θ_t could come from several advanced sources:

- **Nested Volatility Models:** Fitting a secondary GARCH-type model to the volatility of the primary volatility series.
- **High-Frequency Data:** Using metrics derived from intraday data, such as realized quarticity, to measure the variance of variance.
- **Market Microstructure and Liquidity:** Tying θ_t to measures like bid-ask spreads or market depth, which often correlate with market uncertainty and fragility.

This highlights that θ provides a powerful and intuitive lever for controlling the forecast's texture, with significant potential for data-driven enhancement in future work.

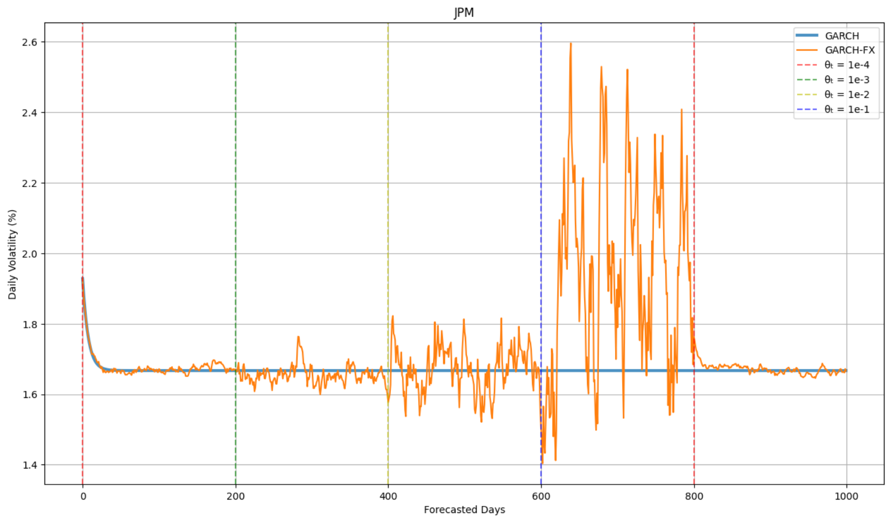

Figure 4: A GARCH-FX forecast demonstrating the effect of a time-varying θ_t. The parameter is programmatically changed at different forecast horizons (e.g., from `1e-4` to `1e-1`), directly controlling the path's jaggedness in a single simulation.

5 Experimental Setup

To empirically validate the performance of the GARCH-FX framework, we design a rigorous backtesting experiment comparing it against a deterministic GARCH forecast, a classical stochastic volatility model (Heston), and the observed realized volatility.

5.1 Data and Backtesting Protocol

The experiment is conducted on the daily closing prices of three large-cap US equities: NVIDIA (NVDA), Apple (AAPL), and Coca Cola (KO). For each stock, we utilize a dynamic data window of 2000 consecutive trading days, sourced from Yahoo Finance.

The backtesting protocol employs a single, fixed train-test split for each asset:

- **In-Sample (Training) Period:** The first 1000 days of data are used to calibrate the baseline GARCH(1,1) model.
- **Out-of-Sample (Forecasting) Period:** The subsequent 1000 days are reserved for forecasting, where the performance of all models is evaluated against the ground truth.

5.2 Ground Truth: Realized Volatility

The ground truth for our comparison is the realized daily volatility. For each day t in the forecasting period, this value is calculated as the standard deviation of daily logarithmic returns over a trailing 180-day lookback window.

This window was chosen to provide a stable, yet responsive, measure of historical volatility against which the long-horizon forecasts can be evaluated. No annualization or other scaling is applied; all model forecasts and the realized volatility are kept in their native daily units to ensure a direct and consistent comparison.

5.3 Model Implementation

5.3.1 GARCH(1,1)

A standard GARCH(1,1) model is fitted to the 1000-day in-sample period to obtain the parameters ω, α, β. Its out-of-sample forecast is generated using the deterministic multi-step forecast equation (2), serving as our baseline.

5.3.2 GARCH-FX Configurations

The GARCH-FX model inherits the ω, α, β parameters directly from the fitted GARCH(1,1) model. To fully evaluate the contribution of each component of the framework, we test two distinct configurations for the 1000-day out-of-sample forecast:

1. **GARCH-FX (Regime-Aware):** This is the full implementation of the model.
   - *Regime Modulation (Δ_t):* The 3-state Markov chain mechanism, as defined in a previous section, is activated to allow for dynamic shifts in the mean volatility level.
   - *Stochasticity Control (θ):* The θ parameter is held constant, with a value selected from the range of `1e-5` to `1e-1` to ensure a realistic level of stochasticity.
2. **GARCH-FX (Stochastic Only):** This configuration is designed to isolate the performance of the core stochastic engine.
   - *Regime Modulation (Δ_t):* The regime-switching mechanism is deactivated by fixing the multiplier $\Delta_t = 1$ for all forecast steps.
   - *Stochasticity Control (θ):* The θ parameter is again held constant within the same range of `1e-5` to `1e-1`.

This dual-configuration approach allows us to disentangle the performance contribution of the stochastic Gamma process from that of the explicit regime-switching overlay.

5.3.3 Heston Model

The Heston model is included as a benchmark representing the classic stochastic volatility family. The model's variance process is simulated using the **Quadratic-Exponential (QE) discretization scheme**, a highly accurate method proposed by Andersen (2008).

This scheme is chosen over simpler methods for its superior properties: it guarantees non-negative variance and provides high accuracy by dynamically switching between two approximation methods based on the local characteristics of the variance process. At each step, it calculates the exact conditional moments of the variance process and, based on the squared coefficient of variation, draws the next variance value from either a quadratic or an exponential approximation of the true non-central chi-squared distribution.

To create a structured and theoretically grounded comparison, we do not perform a separate, unconstrained calibration of the Heston model. Instead, we initialize its key parameters by mapping them directly from the properties of the fitted GARCH(1,1) model. The mapping is as follows:

- **Long-Run Variance (θ):** The Heston long-run variance is set equal to the unconditional variance ($V_L = \omega/(1 - \alpha - \beta)$) derived from the GARCH model.
- **Mean Reversion Speed (κ):** The Heston speed of mean reversion is derived from the GARCH persistence parameter ($\alpha + \beta$). To bridge the discrete-time nature of GARCH with the continuous-time Heston process, we use the standard conversion for daily data: $\kappa = -\ln(\alpha + \beta) \cdot 252$.

- **Volatility-of-Volatility (σ_v):** This parameter is treated as a fixed, asset-specific hyperparameter, chosen to provide a realistic level of variance fluctuation for each stock.

This theoretically-motivated parameter mapping ensures that both the GARCH-FX and Heston models begin their forecasts from an identical, GARCH-derived understanding of long-run variance and persistence. The experiment therefore becomes a direct comparison of the distinct structural assumptions and evolutionary dynamics of each forecasting engine.

5.4 Evaluation Methodology

To quantify forecast accuracy, we use two standard error metrics: Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE).

$$\text{RMSE} = \sqrt{\frac{1}{N}\sum_{i=1}^{N}(\hat{\sigma}_i - \sigma_i)^2} \tag{6}$$

$$\text{MAE} = \frac{1}{N}\sum_{i=1}^{N}|\hat{\sigma}_i - \sigma_i| \tag{7}$$

where $\hat{\sigma}_i$ is the forecasted volatility and σ_i is the realized volatility.

Since both GARCH-FX and the Heston model are stochastic, a single forecast path is insufficient for robust evaluation. Therefore, for each asset, we generate **10 independent simulation paths** using different random seeds. The final RMSE and MAE reported for these models are the **average** of the metrics calculated across all 10 simulations. The results are compiled into a summary table, and representative forecast plots are generated to provide qualitative visual analysis.

6 Experimental Results and Analysis

This section presents the results of our backtesting protocol across the three selected assets. For each asset, we provide a qualitative visual analysis from a single, representative simulation run, followed by a quantitative summary of the average performance over 10 independent runs.

6.1 Asset 1: NVIDIA Corp (NVDA)

Figure 5 displays a representative 1000-day forecast for NVDA, offering a clear visual comparison of the models' behaviors.

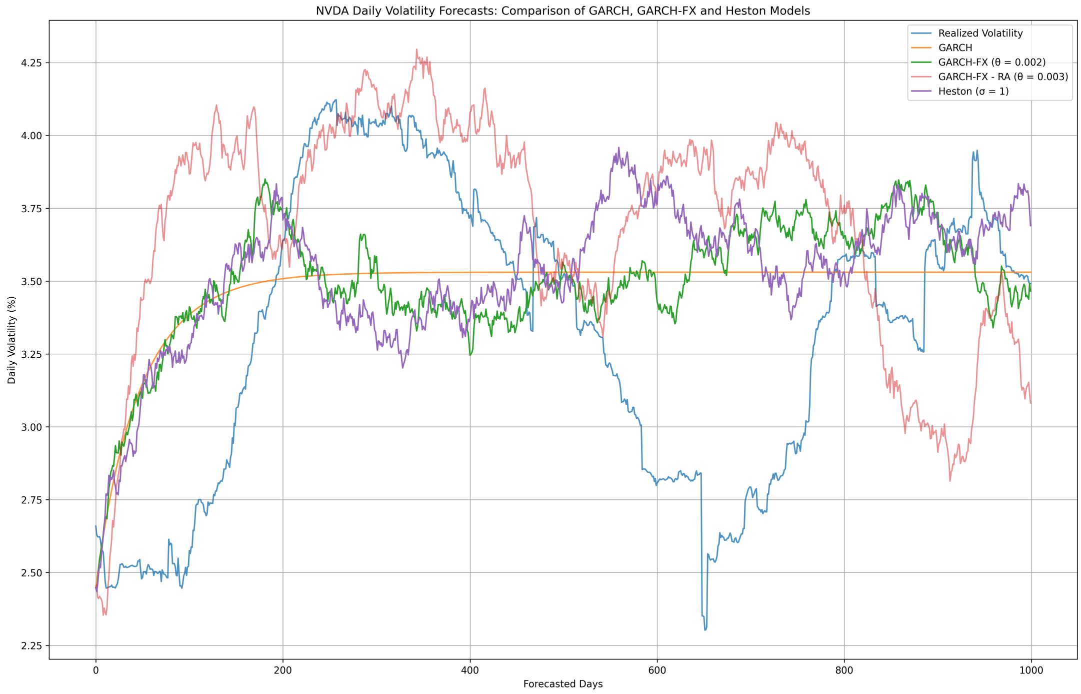

Figure 5: A representative forecast run for NVDA. Realized Volatility (blue) is compared against GARCH-FX (Stochastic Only) (green), GARCH-FX (Regime Aware) (red), Heston (purple), and the deterministic GARCH forecast (orange).

The quantitative results, averaged over 10 simulations, are presented in Table 1.

Table 1: Average Performance Metrics for NVDA Forecasts (%)

Model	**RMSE**	**MAE**
GARCH(1,1)	0.4797	0.3897
Heston	0.4997	0.4155
GARCH-FX	0.4828	0.3972
GARCH-FX (RA)	0.6515	0.5350

The results for NVIDIA (NVDA), a stock renowned for its high growth and trending volatility, provide strong validation of the **GARCH-FX (Stochastic Only)** framework. This configuration demonstrates exceptional performance with an RMSE of 0.4828 and MAE of 0.3972—only marginally higher than the deterministic GARCH(1,1) baseline. The model successfully introduces realistic, dynamic forecast paths absent in the standard GARCH framework while maintaining quantitative accuracy, representing a highly efficient trade-off.

The GARCH-FX (Stochastic Only model also decisively outperforms the Heston model, which recorded a less accurate RMSE of 0.4997 and MAE of 0.4155. Conversely, the **GARCH-FX (Regime-Aware)** configuration proved significantly less effective, confirming that for NVDA's complex volatility dynamics, the simple 3-state Markov chain assumption was inappropriate and detrimental to model performance.inappropriate and detrimental assumption.

6.2 Asset 2: Apple Inc (AAPL)

The forecast for Apple (AAPL) in Figure 6 provides a clear visual summary.

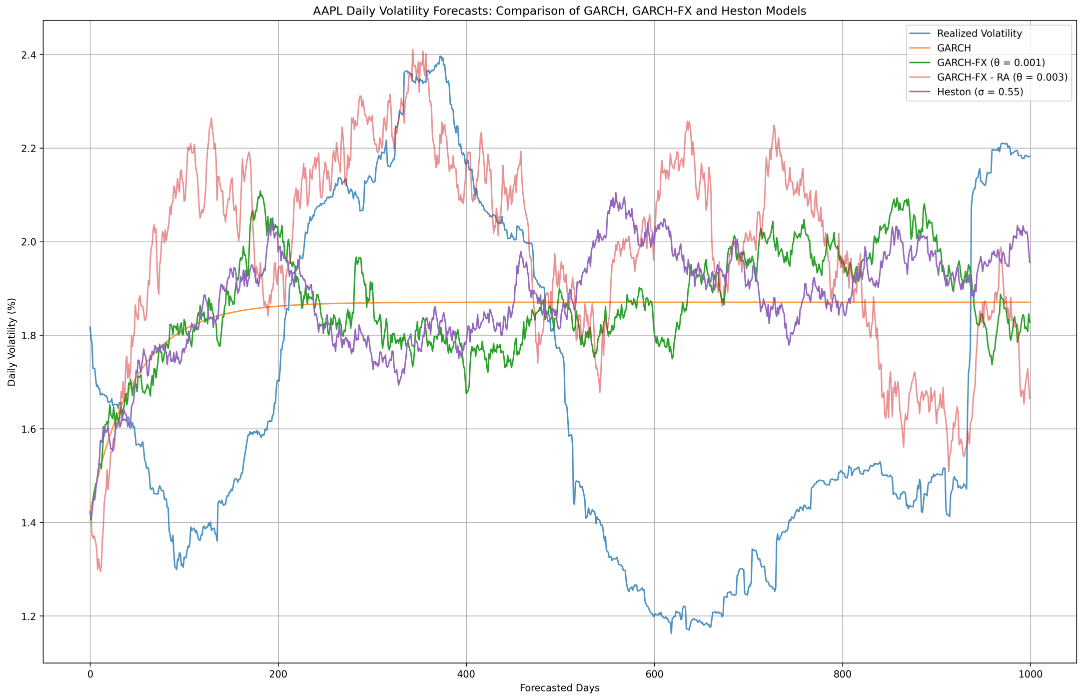

Figure 6: A representative forecast run for AAPL. Realized Volatility (blue) is compared against GARCH-FX (Stochastic Only) (green), GARCH-FX (Regime Aware) (red), Heston (purple), and the deterministic GARCH forecast (orange).

The quantitative analysis in Table 2 provides the definitive performance metrics for AAPL, which serves as a representative case for a mature, large-cap stock

Table 2: Average Performance Metrics for AAPL Forecasts (%)

Model	**RMSE**	**MAE**
GARCH(1,1)	0.4035	0.3646
Heston	0.4163	0.3718
GARCH-FX	0.4153	0.3703
GARCH-FX (RA)	0.5095	0.4340

The performance analysis for Apple (AAPL), representing a mature, large-cap equity, reveals a clear and competitive hierarchy among the tested models. As is typical, the deterministic **GARCH(1,1) model** provides the lowest error scores (RMSE 0.4035, MAE 0.3646), benefiting from its inherent smoothing.

The crucial insights, however, come from the contest between the dynamic forecasting models, where the two leading stochastic approaches are extremely closely matched. **The GARCH-FX (Stochastic Only)** configuration emerges as the top-performing stochastic model, demonstrating a slight but consistent advantage over the Heston model. It achieves a marginally lower RMSE (0.4153 vs. 0.4163) and a similarly lower MAE (0.3703 vs. 0.3718).

In contrast, the **GARCH-FX (Regime-Aware)** model performed poorly, confirming that for a stable asset like AAPL, the simple 3-state Markov chain was an inappropriate structural assumption that hindered performance. This outcome validates the core GARCH-FX stochastic engine as a highly efficient framework, capable of performing on par with established benchmarks like Heston on a major large-cap stock.

6.3 Asset 3: Coca-Cola Co (KO)

Finally, the results for KO are shown in Figure 7 and Table 3.

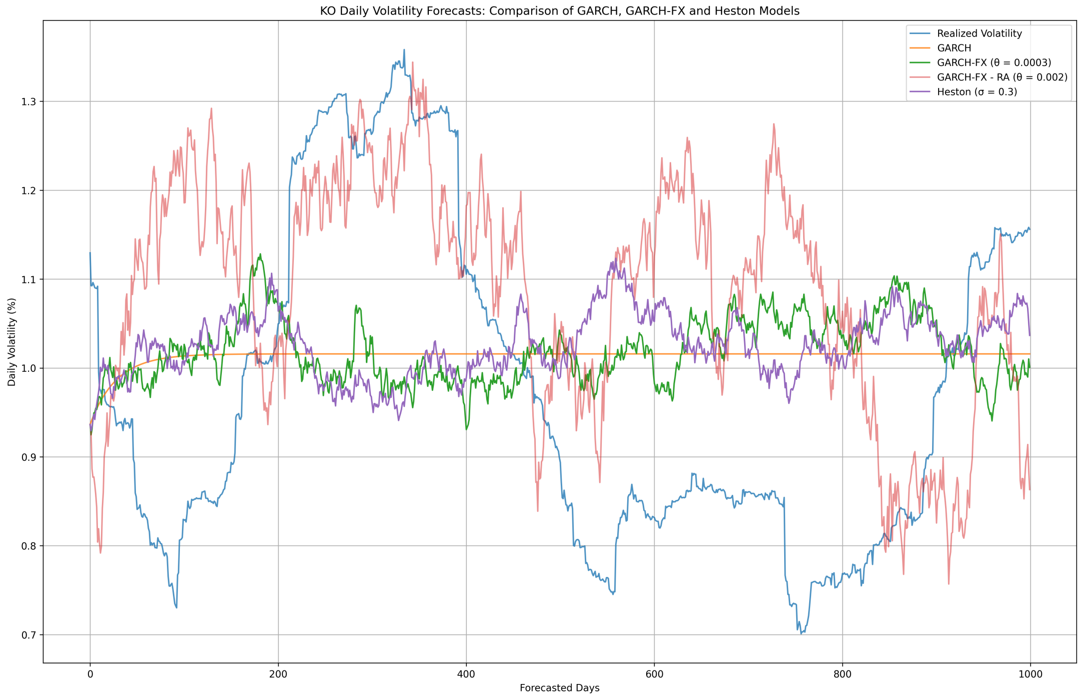

Figure 7: A representative forecast run for KO. Realized Volatility (blue) is compared against GARCH-FX (Stochastic Only) (green), GARCH-FX (Regime Aware) (red), Heston (purple), and the deterministic GARCH forecast (orange).

Table 3: Average Performance Metrics for KO Forecasts (%)

Model	RMSE	MAE
GARCH(1,1)	0.1870	0.1650
Heston	0.1932	0.1698
GARCH-FX	0.1902	0.1687
GARCH-FX (RA)	0.2618	0.2188

The results for Coca-Cola (KO), a classic low-volatility stock, provide a clear performance hierarchy. As expected, the deterministic **GARCH(1,1) model** yields the lowest error scores (RMSE 0.1870, MAE 0.1650) due to its smoothing properties, establishing the baseline for comparison.

The critical insight, however, comes from the competition between the stochastic models, where the **GARCH-FX (Stochastic Only)** framework demonstrates a distinct advantage. It outperforms the Heston model on both key metrics, achieving a lower RMSE (0.1902 vs. 0.1932) and a lower MAE (0.1687 vs. 0.1698). This indicates that for this stable asset, the GARCH-FX core engine was more effective at both avoiding large errors and minimizing the average forecast error compared to the classical Heston framework.

As with other assets, the **GARCH-FX (Regime-Aware)** configuration performed poorly, reinforcing that for a stable stock like KO, the addition of a simple regime-switching mechanism was an unnecessary and detrimental complication. Ultimately, the results for this low-volatility asset demonstrate that the GARCH-FX stochastic engine is a robust framework, capable of producing results that are highly competitive with an established benchmark like the Heston model under our experimental design.

7 Conclusion

In this paper, we introduced GARCH-FX, a novel framework designed not to replace the traditional GARCH model but to extend its forecasting capabilities into the long-horizon, stochastic domain. It addresses a fundamental limitation of deterministic GARCH forecasts: their tendency to "flatline" to a constant mean, which fails to capture the chaotic and contextual nature of real-world volatility. GARCH-FX is built upon the core strength of GARCH—its exceptional ability to model the long-run variance of an asset—and uses this as a baseline for a more realistic, dynamic forecasting engine.

Our experimental results highlight the central contribution of this framework. Across a diverse range of assets, the **GARCH-FX (Stochastic Only)** configuration produced results that were highly competitive with the established Heston model. The key takeaway is not one of superiority, but of parity with simplicity: GARCH-FX demonstrates that it is possible to achieve a comparable level of forecasting performance to classical stochastic volatility models while retaining the implementation simplicity and intuitive structure of the GARCH family. It is crucial, however, to contextualize this comparison as a **structural** test; both models were seeded from the same GARCH-derived parameters to ensure a fair comparison of their underlying engines.

Conversely, the performance of the **GARCH-FX (Regime-Aware)** model serves as a critical insight. Our implementation with a simple 3-state Markov chain often failed to improve accuracy, demonstrating that the regime modulator is a double-edged sword. While a poorly specified signal can degrade performance, this also highlights its immense potential; as shown in our earlier qualitative demonstrations, an accurate external signal fed into the Δ parameter could yield exceptionally precise forecasts. This finding powerfully reinforces the idea that GARCH-FX is a foundational framework, offering a fertile ground for further research and experimentation. It is an open invitation to researchers and practitioners to build upon its modular architecture.

Potential avenues for enhancement are numerous and include:

- **Advanced Regime Drivers for Δ_t:** The user-defined nature of the Delta parameter invites the integration of sophisticated external signals, such as functions of market entropy, Time-Varying Transition Probability (TVTP) models, or machine learning classifiers.
- **Dynamic Calibration of θ_t:** Developing a data-driven method for a time-varying Theta would allow the model to adapt its "volatility of volatility" to changing market conditions, potentially through nested models or high-frequency data.
- **Alternative Stochastic Processes:** While the Gamma distribution is effective, the model's structure allows for experimentation with other non-negative distributions, such as the Weibull or Generalized Gamma.
- **Calibration via Reverse Parameter Mapping:** The theoretical mapping from GARCH to Heston is bidirectional. This opens a novel avenue for initializing GARCH-FX: one could perform a full Heston calibration (potentially using option-implied data) and map the continuous-time parameters back to derive the GARCH persistence and constant term, allowing GARCH-FX to be driven by a market-implied, continuous-time perspective.

Ultimately, GARCH-FX is not a fixed model but a customizable volatility engine. By retaining GARCH's robust structure while injecting stochastic flexibility, it provides a powerful tool for more realistic forecasting, stress-testing, and risk analysis under true market uncertainty.

References

[1] Andersen, L. (2008). *Simple and Efficient Simulation of the Heston Stochastic Volatility Model*. Journal of Computational Finance, 11(3), 1–42.

[2] Bollerslev, T. (1986). *Generalized Autoregressive Conditional Heteroskedasticity*. Journal of Econometrics, 31(3), 307–327.

[3] Engle, R. F. (1982). *Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation*. Econometrica, 50(4), 987–1007.

[4] Hamilton, J. D. (1989). *A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle*. Econometrica, 57(2), 357–384.

[5] Heston, S. L. (1993). *A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options*. The Review of Financial Studies, 6(2), 327–343.