STRATOS
Research

Volatility Clustering and Persistence: A Comparative Study of GARCH Models and Long-Memory Estimation

This paper goes through one of the most important concepts when analyzing volatility: does volatility caused by a shock fade rapidly or does it show long-memory behavior?

Filippo De Angelis

March 13, 2026

1. Introduction

Volatility plays a central role in financial risk management, derivative pricing, and portfolio allocation. One of the most documented empirical regularities in financial markets is volatility clustering: large price changes tend to be followed by large changes, and small changes by small movements. Understanding whether this clustering reflects short-term persistence or a stronger form of long-range dependence is crucial for accurate volatility modeling.

Traditional econometric models such as GARCH assume that volatility is conditionally heteroskedastic but ultimately mean-reverting, implying a short-memory structure in which shocks dissipate. In contrast, alternative approaches inspired by the fractal market theory [4] suggest that volatility may exhibit long-memory behavior, characterized by a much slower decay of dependence.

This paper reviews the theoretical foundations of ARCH and GARCH models and contrasts them with the concept of long-memory volatility. The comparison is then explored empirically through the analysis of volatility dynamics surrounding the Covid-19 crisis, which provides a natural experiment to assess how quickly volatility reverts after a major exogenous shock.

2. The Emergence of Conditional Heteroskedasticity

Prior to the 1980s, standard econometric models assumed homoskedasticity that means that the variance of the error term was constant over time. Under this framework, financial time series were modeled with time-invariant volatility, implying that uncertainty was stable. However, empirical evidence increasingly suggested that financial returns exhibit periods of increased volatility, contradicting the assumption of constant variance.

A major breakthrough occurred with Engle (1982) [3], who introduced the Autoregressive Conditional Heteroskedasticity (ARCH) model. Engle begins by specifying a conventional linear regression model for the conditional mean-first moment-where the dependent variable is regressed with lagged exogenous and endogenous variables, but the innovation is in the specification of the second moment, the variance of the error term could evolve dynamically over time as a function of past squared information. Formally, the ARCH(1) model specifies:

$$E(y_t \mid \Omega_{t-1}) = x_t'\beta \tag{1}$$

$$\varepsilon_t \mid \Omega_{t-1} \sim N(0, h_t) \tag{2}$$

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 \tag{3}$$

where h_t represents the conditional variance at time t, and Ω_{t-1} denotes the information set available at time $t-1$.

The central innovation of Engle's framework was the idea that volatility is conditionally predictable. Instead of treating variance as constant, the ARCH model allows large shocks to increase future volatility, indeed, if a shock occurred at t-1 and therefore ε_{t-1}^2 is large, the variance will be large itself; conversely, if markets are calm, ε_{t-1}^2 is small implying a lower variance. This encapsulates the idea of the already mentioned volatility clustering, that is, periods of high volatility are followed by periods of high volatility and periods of low volatility tend to be followed by periods of low volatility.

Engle also derived the conditions under which the process remains covariance stationary. For an ARCH(1) model, the unconditional variance exists only if $\alpha_1 < 1$, and is given by:

$$\sigma^2 = \frac{\alpha_0}{1-\alpha_1} \tag{4}$$

As α_1 approaches the value of 1, volatility becomes increasingly persistent, although the dependence structure remains finite due to the limited lag structure of the model.

3. Interpretation and estimation of the parameters

Given the aforementioned variance equation:

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 \tag{5}$$

α_0 represents the baseline level of volatility, intuitively it represents the level of volatility in the market if no past shocks have occurred.

While α_1 represents the sensitivity of current volatility to past squared information. If it is large, current volatility would increase more given past shocks. In a more mathematical way, the partial derivative of the conditional variance with respect to past squared innovations:

$$\frac{\partial h_t}{\partial \varepsilon_{t-1}^2} = \alpha_1 \tag{6}$$

Thus, α_1 represents the marginal impact of past squared shocks on current conditional variance.

In the ARCH model, variance is modeled as a function of squared past shocks; although volatility is not directly observable from the market, what is directly observable are returns and residuals. This is to say that even if the ARCH resembles a standard regression model, in reality variance is unobserved and it depends on past innovations, which have a normal distribution with mean 0 and time-varying variance. This probabilistic characteristic of the model requires a method of estimating the parameters that fully reflects it. The Maximum Likelihood Estimation (MLE) method is the natural choice, since another estimation method such as OLS would ignore the density, that is, the probabilistic feature of the model. Instead Maximum Likelihood uses it in its favor by choosing the parameter that makes the data most probable, in simple words, it chooses the volatility path that makes the observed shocks plausible. This is done by maximizing the log-likelihood function

The introduction of ARCH fundamentally altered the econometric treatment of financial time series. It shifted the focus from modeling only the conditional mean to explicitly modeling the conditional second moment. This marked the beginning of modern volatility modeling and laid the foundation for subsequent extensions such as the GARCH framework.

4. The Generalized ARCH Model, GARCH

Although the ARCH framework represented a substantial innovation in volatility modeling, empirical applications quickly revealed an important limitation. In order to adequately capture the high persistence typically observed in financial volatility, ARCH models often require a large number of lagged squared residuals. Such high-order specifications reduce parsimony, efficiency and may lead to estimation difficulties.

To address this issue, Bollerslev (1986) [1] introduced the Generalized Autoregressive Conditional Heteroskedasticity (GARCH) model, which extends the ARCH structure by allowing past conditional variances to enter the variance equation. This extension is analogous to the generalization from autoregressive (AR) to autoregressive moving average (ARMA) [2] models in the conditional mean.

Formally, a GARCH(p,q) process is defined as:

$$\varepsilon_t \mid \Omega_{t-1} \sim N(0, h_t) \tag{7}$$

$$h_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j h_{t-j} \tag{8}$$

where $\alpha_0 > 0$, $\alpha_i \geq 0$, and $\beta_j \geq 0$. When $p = 0$, the model reduces to the ARCH(q) specification. The inclusion of lagged conditional variances allows volatility to depend not only on past shocks but also on its own past realizations.

An important theoretical result derived by Bollerslev (1986) is that the GARCH(p,q) process is covariance stationary if and only if:

$$\sum_{i=1}^{q} \alpha_i + \sum_{j=1}^{p} \beta_j < 1 \tag{9}$$

Under this condition, the unconditional variance exists and is given by:

$$\sigma^2 = \frac{\alpha_0}{1 - \sum_{i=1}^{q} \alpha_i - \sum_{j=1}^{p} \beta_j} \tag{10}$$

This result generalizes the ARCH stationarity condition and highlights the joint role of ARCH and GARCH effects in determining volatility persistence.

The GARCH(1,1) Specification

In empirical finance, the most widely used specification is the GARCH(1,1) model:

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 h_{t-1} \tag{11}$$

Despite its simplicity, this specification captures volatility clustering remarkably well. The degree of persistence is summarized by the quantity $\alpha_1 + \beta_1$. When this sum is close to one, shocks to volatility decay slowly, generating highly persistent volatility dynamics. However, as long as $\alpha_1 + \beta_1 < 1$, the process remains mean-reverting and exhibits finite unconditional variance.

There is only one source of randomness given by the driving white noise ε_{t-1}^2, while the volatility process conditioned to current information is deterministic.

From an economic perspective, the GARCH specification may be interpreted as an adaptive learning mechanism: large past shocks increase current volatility, and elevated volatility feeds into future variance through the lagged h_{t-1} term. Importantly, volatility shocks decay at a rate determined by $\alpha_1 + \beta_1$, implying a short-memory structure in which the dependence eventually vanishes.

5. Long Memory and Fractal Volatility

Although the GARCH framework captures volatility clustering effectively, it assumes that shocks to volatility eventually fade at a relatively fast rate. In particular, volatility shocks decay progressively over time, and the process ultimately returns to its long-run average level, which is the mean reverting property. This property is often referred to as short-memory persistence.

However, several studies suggest that financial volatility may display a stronger form of dependence known as long memory. In a long-memory process, the impact of a shock declines much more slowly. Periods of high volatility can influence the market for an extended time, sometimes over weeks or months.

The essential difference lies in the speed at which the shocks disappear. In short-memory models such as GARCH, volatility reverts to its mean at a relatively quick pace given by $\alpha_1 + \beta_1$. In contrast, long-memory models imply that volatility clusters are more persistent and that dependence between distant observations remains significant.

From a real life financial perspective, this distinction is truly important . If volatility exhibits long memory, risk forecasts may remain elevated for longer periods, and shocks in the market may have more durable effects. Consequently, determining whether volatility follows a short-memory or long-memory structure becomes an empirical question.

6. Empirical analysis

This section presents the empirical analysis aimed at examining the behavior of volatility following a major exogenous shock. The Covid-19 crisis is examined because it was an exogenous and global shock that generated one of the highest spikes of volatility in financial history, making it an ideal event to understand whether volatility shocks dissipate rapidly, as predicted by short-memory models, or whether they exhibit prolonged persistence. VIX index data were obtained from the Federal Reserve Economic Data (FRED) database (https://fred.stlouisfed.org) for the period from 2018/01/01 to 2023/01/01. All calculations and graphical representations were performed using MATLAB. The data were divided into three subcategories:

After that, the first indicator of volatility was computed as the sample mean for each of the

Table 1: Sample Division Around the Covid-19 Crisis

Period	Start Date	End Date
Pre-Covid	01/01/2018	31/12/2019
Covid Shock	01/02/2020	30/06/2020
Post-Covid	01/07/2020	31/12/2022

three periods in order to assess whether volatility changed over time.

Table 2: Average VIX Levels Across Subperiods, the calculation were computed in Matlab.

Subperiod	Mean VIX Level
Pre-Covid (Jan 2018 – Dec 2019)	16.01
Covid Crisis (Feb 2020 – Jun 2020)	36.69
Post-Covid (Jul 2020 – Dec 2022)	23.28

By looking at the table, it is immediately understandable that the process shows persistence and slow mean reversion, in fact, the post-crisis average volatility remains approximately 45 percent higher than the pre-crisis level. Furthermore, the autocorrelation function (ACF) was also computed to better assess the persistence of volatility dynamics. The ACF measures the degree of dependence between the current and past values of a time series. The autocorrelation results are shown in Figure 1.

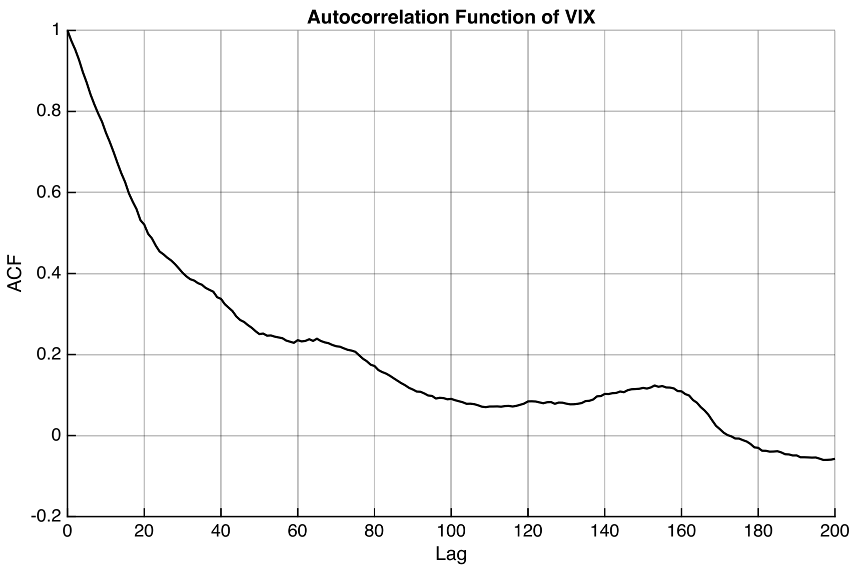

Figure 1: Autocorrelation Function of the VIX. The calculations were computed using MATLAB.

The autocorrelation exhibits very slow decay. The first-lag autocorrelation is close to one, indicating strong dependence between consecutive observations. Moreover, positive autocorrelation persists for a large number of lags and approaches zero only after approximately 175 trading days.

7. Conclusion

This paper examined the behavior of financial volatility following the Covid-19 crisis, with the objective of assessing whether volatility shocks have dissipated rapidly, as implied by short-memory models, or exhibit more prolonged persistence.

The theoretical discussion highlighted the distinction between short-memory frameworks, such as GARCH, where shocks decay at a geometric rate, and long-memory approaches, where dependence may persist over extended horizons. The empirical analysis, based on daily VIX data from 2018 to 2023, provided descriptive evidence on volatility dynamics during and after the Covid shock.

The results show that the average volatility level in the post-crisis period remained substantially higher than in the pre-Covid benchmark. Moreover, the autocorrelation function exhibits a slow decay, with significant positive dependence persisting for a large number of trading days. These findings indicate a high degree of persistence and a slow mean reversion in volatility following the crisis.

Even though the analysis does not constitute a formal test of long-memory dynamics, it provides recent evidence that suggests that major exogenous shocks can generate sustained volatility regimes. This has important implications for volatility modeling, forecasting, and risk management, as models must be able to capture prolonged dependence in periods of financial stress.

References

[1] Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986.

[2] George E. P. Box, Gwilym M. Jenkins, Gregory C. Reinsel, and Greta M. Ljung. *Time Series Analysis: Forecasting and Control*. Wiley, Hoboken, NJ, 5 edition, 2015.

[3] Robert F. Engle. Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007, 1982.

[4] Benoit B. Mandelbrot. *Fractals and Scaling in Finance: Discontinuity, Concentration, Risk*. Springer, New York, 1997.