Studying the Volatility Structure of a Time Series: A Complete Guide

Andrea Buffoli, Davide Rolfi

July 2025

Understanding Volatility: Why ARCH and GARCH Matter

When analyzing a time series, particularly in finance, it's not enough to understand just the expected value (mean). Equally important is the variability of the series over time, that is, its **volatility**. In financial data, volatility often clusters: periods of large movements are followed by large movements, and quiet periods by further quiet. This behavior motivates models with **conditional heteroskedasticity**.

ARCH and **GARCH** models allow the variance to change over time as a function of past information. They differ from classical models with constant variance (homoskedasticity), and provide a more realistic framework for modeling risk and uncertainty.

ARCH(q): Autoregressive Conditional Heteroskedasticity

The ARCH model was introduced by Engle (1982) to model time-varying volatility in a univariate time series.

Model Setup

Let $\{X_t\}$ be a weakly stationary time series. Assume the process follows:

$$X_t = \mu + \varepsilon_t \tag{1}$$

$$\varepsilon_t = \sigma_t z_t \quad z_t \sim \text{i.i.d. } \mathcal{N}(0, 1). \tag{2}$$

The innovation, or error term, ε_t has zero mean and time-varying conditional variance σ_t^2. The ARCH(q) model specifies:

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2, \tag{3}$$

with $\alpha_0 > 0$ and $\alpha_i \geq 0$ to ensure $\sigma_t^2 > 0$.

Conditional and Unconditional Moments in ARCH Models

To better understand how ARCH models capture volatility dynamics, we need to distinguish between two types of moments: those that are *conditional* on past information, and those that are *unconditional*, or marginal.

Conditional mean The conditional mean of X_t given past information $\mathcal{F}_{t-1}$ is:

$$\mathbb{E}[X_t \mid \mathcal{F}_{t-1}] = \mu + \mathbb{E}[\varepsilon_t \mid \mathcal{F}_{t-1}] = \mu$$

because $\mathbb{E}[z_t \mid \mathcal{F}_{t-1}] = 0$ and $\varepsilon_t = \sigma_t z_t$ is mean-zero conditionally.

Interpretation: The ARCH model assumes that the expected return (or the expected value of the series) is constant over time. All the dynamics and temporal dependence are captured by the second moment, i.e., the conditional variance. Therefore, even though the mean is constant, the risk or uncertainty around that mean can vary over time depending on the size of past shocks.

Conditional variance The conditional variance of X_t given past information is:

$$\mathbb{V}\text{ar}[X_t \mid \mathcal{F}_{t-1}] = \mathbb{V}\text{ar}[\varepsilon_t \mid \mathcal{F}_{t-1}] = \sigma_t^2 = \alpha_0 + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2$$

Interpretation: The current volatility depends on past squared shocks. If large shocks occurred in the past, the conditional variance increases, even if the past returns themselves are mean-zero. This leads to a key empirical feature observed in financial time series: *volatility clustering*. Volatility is serially correlated, even if returns themselves are not. Large returns tend to be followed by large returns (in magnitude), regardless of their sign.

Unconditional mean To compute the unconditional (i.e., long-run average) mean of X_t, we take the expectation over all information:

$$\mathbb{E}[X_t] = \mathbb{E}[\mu + \varepsilon_t] = \mu + \mathbb{E}[\varepsilon_t] = \mu$$

Interpretation: Again, since ε_t is centered, the long-term average of the process is simply the drift μ. The ARCH process does not change the expected level of the series, it modulates its uncertainty.

Unconditional variance Deriving the unconditional variance is more subtle. For general ARCH(q), there is no simple closed-form expression. However, in the special case of ARCH(1):

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2$$

with $\varepsilon_t = \sigma_t z_t$, and $z_t \sim \mathcal{N}(0, 1)$.

Taking expectations on both sides, and assuming covariance stationarity (i.e., constant unconditional variance $\mathbb{V}[X_t]$ exists), we can write:

$$\mathbb{E}[\varepsilon_t^2] = \mathbb{E}[\sigma_t^2] = \alpha_0 + \alpha_1 \mathbb{E}[\varepsilon_{t-1}^2] \Rightarrow \sigma^2 = \frac{\alpha_0}{1-\alpha_1}, \quad \text{if } \alpha_1 < 1$$

Interpretation: The unconditional variance exists and is finite only if $\alpha_1 < 1$. As α_1 approaches 1, the volatility persistence increases and the unconditional variance becomes very large. This explains why ARCH models are suitable to model periods of extreme volatility, when past shocks have long-lasting effects.

Note: For ARCH(q) with $q > 1$, analytical expressions become complex, and moments are usually computed numerically or via simulation.

Stationarity Conditions for ARCH(q)

The ARCH(q) model is said to be **weakly stationary** (i.e., with constant unconditional mean and variance, and time-invariant autocovariances) if the following conditions hold:

- The error terms $\{z_t\}$ are i.i.d. with $\mathbb{E}[z_t] = 0$, $\mathbb{V}[z_t] = 1$.
- The parameters satisfy: $\alpha_0 > 0$ and $\alpha_i \geq 0$ for all $i = 1, \ldots, q$.
- The total impact of past shocks is bounded:

$$\sum_{i=1}^{q} \alpha_i < 1$$

Why is this condition important? If $\sum \alpha_i < 1$, then the process has a finite unconditional variance:

$$\mathbb{V}[X_t] = \frac{\alpha_0}{1 - \sum_{i=1}^{q} \alpha_i}$$

This guarantees that the volatility does not explode and that the model captures time-varying but bounded risk.

If the condition is violated (i.e., $\sum \alpha_i \geq 1$), then the variance becomes infinite or undefined, and the process is nonstationary. In that case, the volatility tends to grow over time, which is usually undesirable in modeling stable financial returns.

Key Properties and Economic Interpretation

ARCH models embody several features that make them appealing for modeling financial time series:

- **Volatility clustering:** Due to dependence on past squared residuals, the model naturally generates periods of high and low volatility.
- **Uncorrelated but dependent:** Even though the returns X_t may be uncorrelated (white noise), their squared values ε_t^2 are autocorrelated, creating a dependence structure in higher-order moments.

- **Leptokurtosis:** ARCH models can produce heavy tails in the marginal distribution of X_t (excess kurtosis), matching empirical data better than constant-variance models.
- **Stationarity conditions:** The model is weakly stationary only if $\sum_{i=1}^{q} \alpha_i < 1$. Violating this condition may result in infinite variance or explosive volatility.

ARCH-type models separate the modeling of the expected value (mean) and uncertainty (variance), which is particularly suitable for financial applications where the focus is on risk prediction rather than return forecasting.

GARCH(p, q): Generalized Autoregressive Conditional Heteroskedasticity

The ARCH model captures short-term dependence in volatility, but often fails to reproduce the strong and persistent volatility clustering observed in financial time series. To address this, **Bollerslev (1986)** proposed the GARCH model, which includes lagged conditional variances in the specification of σ_t^2. This allows past volatility to directly influence current volatility, introducing a form of *memory* into the volatility process.

Model Setup

Let $\{X_t\}$ be a weakly stationary time series such that:

$$X_t = \mu + \varepsilon_t \tag{4}$$

$$\varepsilon_t = \sigma_t z_t, \quad z_t \sim \text{i.i.d. } \mathcal{N}(0, 1) \tag{5}$$

Then, the conditional variance σ_t^2 evolves as:

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

with $\alpha_0 > 0$, $\alpha_i \geq 0$, $\beta_j \geq 0$. This is the GARCH(p, q) specification.
The simplest case is GARCH(1,1), where:

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2$$

Conditional and Unconditional Moments

Conditional mean

$$\mathbb{E}[X_t \mid \mathcal{F}_{t-1}] = \mu$$

The conditional expectation remains constant and equal to μ, since the shock z_t has zero conditional mean.

Conditional variance

$$\mathbb{V}\mathrm{ar}[X_t \mid \mathcal{F}_{t-1}] = \sigma_t^2$$

This evolves over time and depends on both past squared shocks and past variances, capturing both the short-run impact of volatility and its persistence.

Unconditional mean

$$\mathbb{E}[X_t] = \mu$$

The long-run mean of the process remains constant, since $\mathbb{E}[\varepsilon_t] = 0$.

Unconditional variance For the GARCH(1,1) model, the unconditional variance exists only if $\alpha_1 + \beta_1 < 1$. In that case:

$$\mathbb{V}[X_t] = \frac{\alpha_0}{1 - \alpha_1 - \beta_1}$$

Interpretation: the condition ensures that past volatility and past shocks do not indefinitely amplify the variance. The closer $\alpha_1 + \beta_1$ is to 1, the more persistent the volatility, and the longer it takes for shocks to dissipate.

Stationarity Conditions for GARCH(1,1)

The GARCH(1,1) process is weakly stationary if:

- The innovations z_t are i.i.d. with $\mathbb{E}[z_t^2] = 1$
- Parameters satisfy $\alpha_0 > 0$, $\alpha_1 \geq 0$, $\beta_1 \geq 0$
- The persistence condition holds: $\alpha_1 + \beta_1 < 1$

Why it matters: When the persistence condition fails (i.e., sum ≥ 1), the process may become nonstationary and exhibit explosive behavior in volatility. This is undesirable in most applications and invalidates standard inference techniques.

Economic Interpretation of Parameters

- α_0: Long-run baseline level of volatility.
- α_1: Sensitivity to recent news (short-term shocks).
- β_1: Persistence or memory in volatility.
- $\alpha_1 + \beta_1$: Total persistence. If close to 1, volatility decays slowly.

The GARCH model is thus capable of reproducing key empirical facts in financial data: volatility clustering, leptokurtosis, and the persistence of risk over time. Unlike ARCH, it achieves this with parsimony: GARCH(1,1) often suffices where ARCH would require many lags.

Simulation-Based Estimation and Diagnostics of GARCH(1,1)

To illustrate the behavior of GARCH models in practice, we simulate a return series (see 1) of 10,000 observations using a GARCH(1,1) process with the following parameters:

$$\alpha_0 = 0.1, \quad \alpha_1 = 0.4, \quad \beta_1 = 0.5$$

These values satisfy the stationarity condition $\alpha_1 + \beta_1 < 1$, ensuring that the unconditional variance is finite and constant over time. The simulated series evolves according to:

$$\sigma_t^2 = \alpha_0 + \alpha_1 r_{t-1}^2 + \beta_1 \sigma_{t-1}^2, \quad r_t = \sigma_t z_t, \quad z_t \sim \mathcal{N}(0,1)$$

Note: In standard GARCH notation, the conditional variance depends on past squared residuals $\epsilon_{t-1}^2 = (r_{t-1} - \mu)^2$. *Here we use* r_{t-1}^2 *because the simulated process assumes zero mean, i.e.,* $\mu = 0$, *so that* $\epsilon_t = r_t$.

The goal of this simulation is to replicate the main empirical features captured by GARCH models: persistence in volatility and the clustering of large shocks. Even though the underlying innovations z_t are standard normal and independent, the resulting return series displays stretches of high and low volatility.

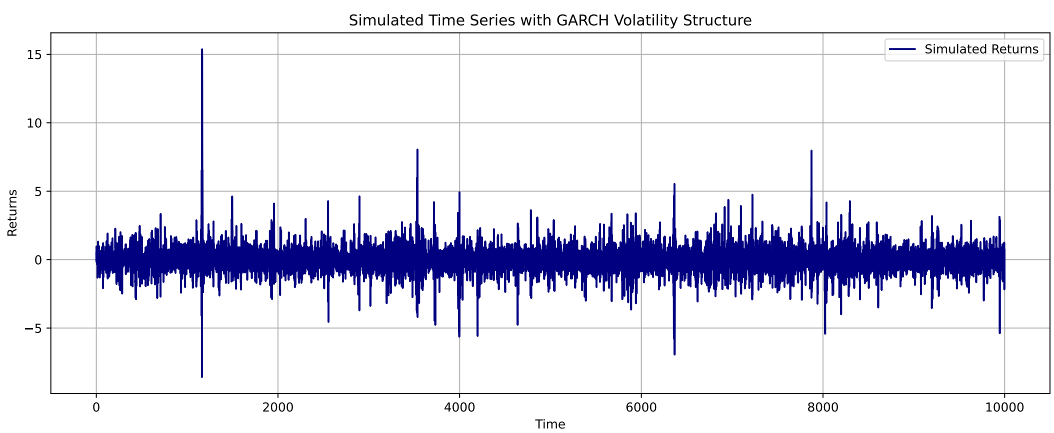

Figure 1: Simulated return series from a GARCH(1,1) process. We show the simulated path of r_t over 10,000 periods. The plot illustrates volatility clustering, where large returns (in absolute value) tend to be followed by other large returns. This dynamic behavior is not due to correlation in the returns themselves, but rather in their second moments, and is well captured by the GARCH structure.

To further examine whether the simulated series indeed exhibits conditional heteroskedasticity, we analyze the autocorrelation structure of the squared returns. Figure 2 shows the ACF and PACF of r_t^2. Both plots reveal significant autocorrelation at low lags, providing visual evidence of volatility clustering. This is consistent with the presence of ARCH effects, where the variance depends on past squared returns, even though the returns themselves are uncorrelated.

To formally test for ARCH-type behavior, we apply the Lagrange Multiplier test proposed by Engle (1982). The null hypothesis is that the series is homoskedastic, that is, the conditional variance is constant over time. As shown in Table 1, the test statistics strongly reject the null at conventional significance levels, confirming the presence of conditional heteroskedasticity in the simulated series.

Table 1: Engle ARCH Test Results (10 lags)

Statistic	Value	p-value
LM Statistic	3973.2741	0.0000
F Statistic	658.9847	0.0000

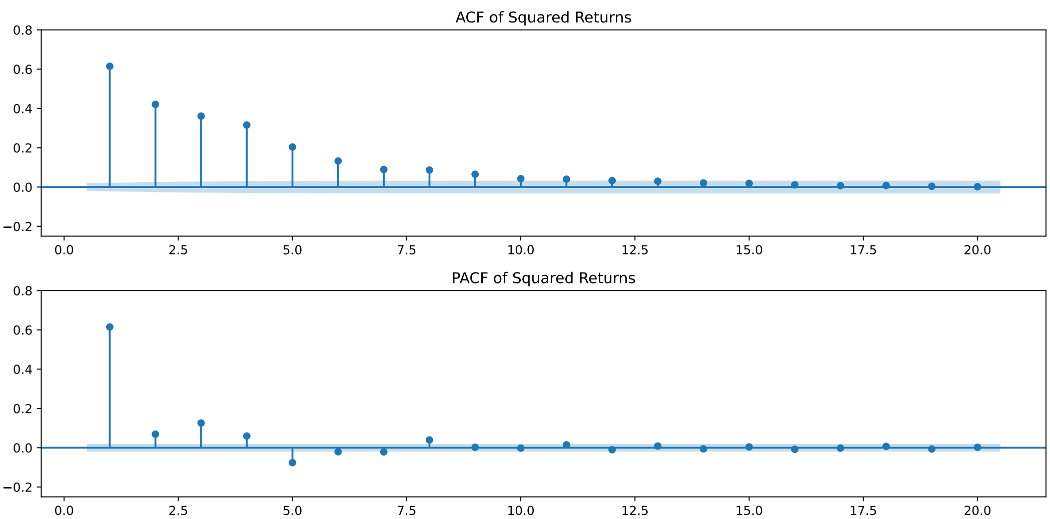

Figure 2: ACF and PACF of squared returns from the simulated GARCH(1,1) series. Significant autocorrelations at short lags suggest the presence of time-varying volatility, as expected under conditional heteroskedasticity.

Having established the presence of conditional heteroskedasticity, we proceed to fit an ARCH(q) model to the simulated data. We estimate a sequence of ARCH(q) models for $q = 1, \ldots, 8$ and select the specification with the lowest Bayesian Information Criterion (BIC). The optimal lag length corresponds to the model with $q = 5$, indicating that the conditional variance of returns is best explained by the past four squared residuals.

Table 2 reports the estimated parameters for the ARCH(5) model, along with standard errors, t-statistics, p-values, and confidence intervals. All coefficients are statistically significant at the 5% level, and the magnitude of the α_i parameters decreases with lag order, reflecting a decaying influence of past shocks on current volatility. The model also confirms a positive and significant α_0, capturing the baseline level of variance.

After estimating the model, we assess its adequacy by analyzing the standardized residuals $\hat{z}_t$:

$$\hat{z}_t = \frac{r_t - \hat{\mu}}{\hat{\sigma}_t}$$

Table 2: Estimated parameters of the ARCH(5) model

Parameter	Coefficient	Std. Error	t-Statistic	p-value	95% Conf. Interval
α_0	0.2183	0.00900	24.260	5.171×10^{-130}	[0.201, 0.236]
α_1	0.3926	0.01921	20.439	7.544×10^{-93}	[0.355, 0.430]
α_2	0.1967	0.01581	12.446	1.477×10^{-35}	[0.166, 0.228]
α_3	0.0778	0.01191	6.533	6.445×10^{-11}	[0.0545, 0.101]
α_4	0.0696	0.01215	5.726	1.028×10^{-8}	[0.0458, 0.0934]
α_5	0.0338	0.00927	3.642	2.704×10^{-4}	[0.0156, 0.0519]

If the model is well specified, these standardized residuals should behave like white noise, i.e., they should exhibit no autocorrelation, and their squared values should not display conditional dependence.

Figure 3 shows the ACF and PACF of the standardized residuals $\hat{z}_t$. The absence of significant autocorrelation suggests that the ARCH(5) model has effectively captured the linear dependence in the data. Figure 4 presents the ACF and PACF of $\hat{z}_t^2$. The lack of remaining autocorrelation in the squared standardized residuals confirms that no further ARCH effects are present, validating the adequacy of the chosen model specification.

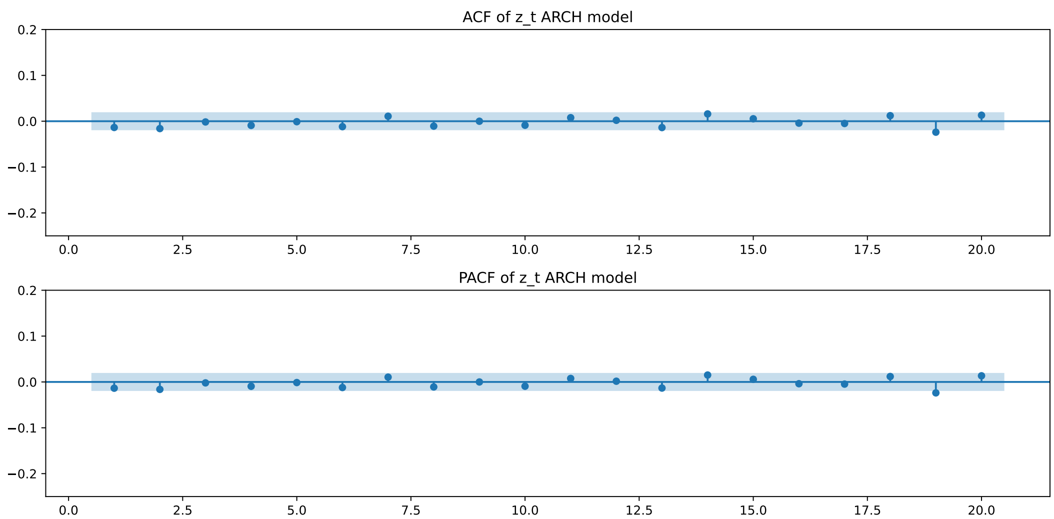

Figure 3: ACF and PACF of standardized residuals $\hat{z}_t$ from the ARCH(5) model. The residuals behave as white noise, indicating that the model has properly captured the conditional mean and volatility dynamics.

We now turn to fitting a GARCH(p,q) model, which extends the ARCH specification by incorporating past conditional variances into the volatility equation. We estimate all combinations with $1 \leq p, q \leq 5$, and select the optimal specification using the Bayesian Information Criterion. The best-performing model is GARCH(1,1), which coincides with the true data-generating process used in the simulation.

Table 3 reports the estimated parameters of the selected model. All coefficients are statistically significant at conventional levels. The estimated value of α_0 is 0.1046, indicating

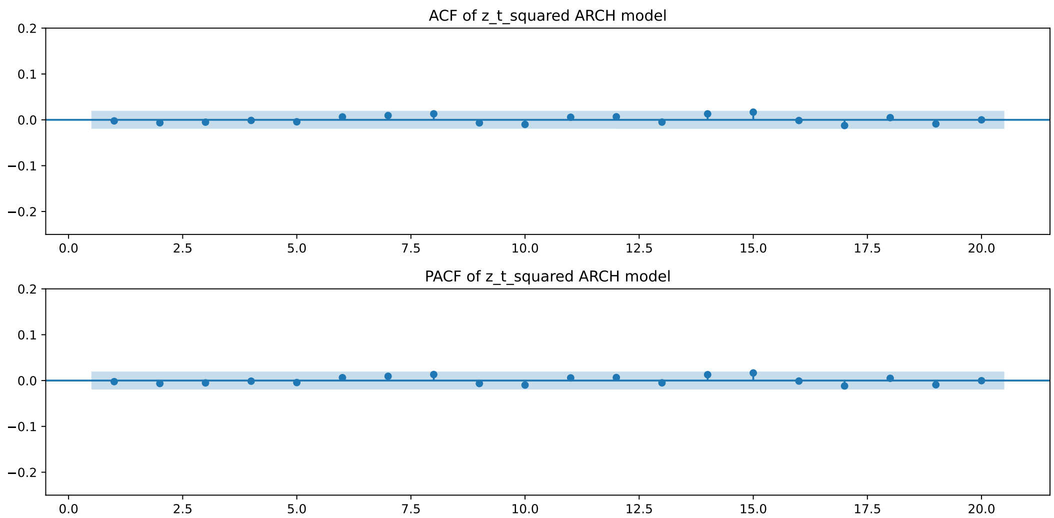

Figure 4: ACF and PACF of squared standardized residuals $\hat{z}_t^2$ from the ARCH(5) model. The disappearance of autocorrelations confirms that no ARCH effects remain in the residuals, supporting the validity of the model.

a positive long-run baseline level of volatility. The parameter $\alpha_1 = 0.3841$ captures the sensitivity to recent shocks, while $\beta_1 = 0.5055$ reflects the persistence of past volatility. The sum $\alpha_1 + \beta_1 = 0.8896$ is less than 1, confirming that the process is stationary, yet exhibits strong volatility persistence.

Table 3: Estimated parameters of the GARCH(1,1) model

Parameter	Coefficient	Std. Error	t-Statistic	p-value	95% Conf. Interval
α_0	0.1046	0.00707	14.800	1.460×10^{-49}	[0.0908, 0.118]
α_1	0.3841	0.01720	22.333	1.777×10^{-110}	[0.350, 0.418]
β_1	0.5055	0.01789	28.261	1.034×10^{-175}	[0.470, 0.541]

After estimating the model, we assess its adequacy by analyzing the standardized residuals:

$$\hat{z}_t = \frac{r_t - \hat{\mu}}{\hat{\sigma}_t}$$

In a well-specified GARCH model, these residuals should behave as white noise, uncorrelated and with constant variance over time. Moreover, the squared standardized residuals should exhibit no remaining ARCH effects.

Figure 5 displays the ACF and PACF of the standardized residuals from the GARCH(1,1) model. There is no evidence of significant autocorrelation, suggesting that the mean and volatility dynamics are adequately captured. Figure 6 presents the ACF and PACF of the squared standardized residuals. The absence of significant correlations here further confirms that the GARCH(1,1) specification successfully absorbs the conditional heteroskedasticity in the data. This result, together with its parsimony, makes the GARCH(1,1) model particularly attractive: it nests an ARCH(∞) process and captures persistence in volatility with fewer parameters than a higher-order ARCH model. In our case, the GARCH(1,1) model achieves a substantially lower Bayesian Information Criterion (BIC = 2239.22) compared to the ARCH(5) model (BIC = 23705.4), providing strong evidence in favor of the GARCH specification despite its simpler structure.

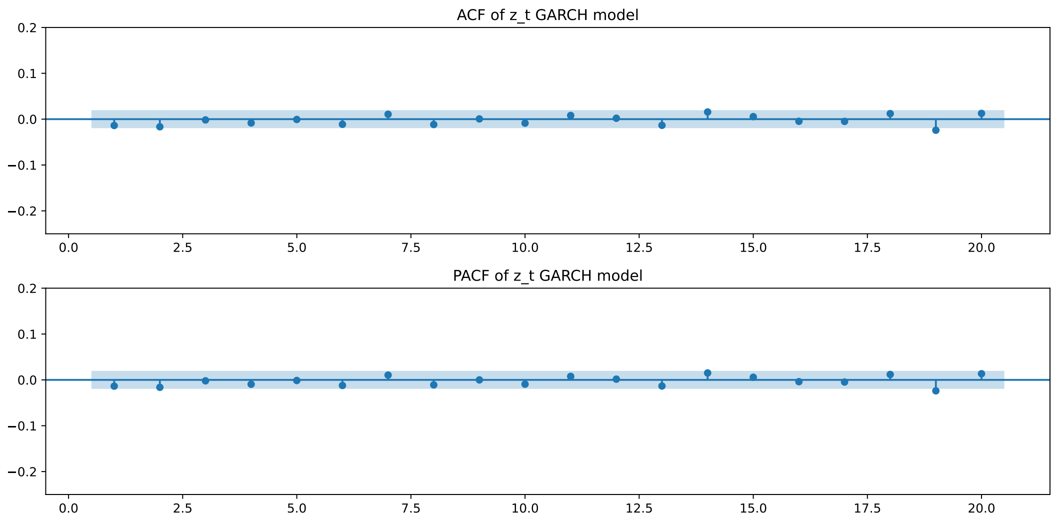

Figure 5: ACF and PACF of standardized residuals $\hat{z}_t$ from the GARCH(1,1) model. The residuals appear uncorrelated, supporting the adequacy of the fitted model in capturing the time dependence in volatility.

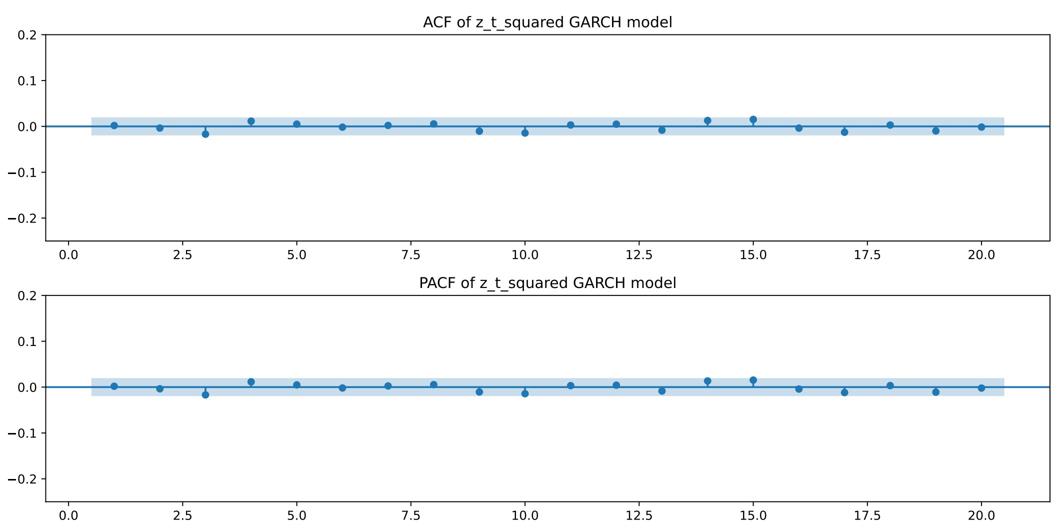

Figure 6: ACF and PACF of squared standardized residuals $\hat{z}_t^2$ from the GARCH(1,1) model. The lack of autocorrelation in the squared residuals confirms that no ARCH effects remain, validating the model's specification.

This analysis has illustrated a systematic procedure for modeling volatility in time series data using ARCH and GARCH frameworks. Starting from the simulation of a GARCH(1,1) process, we diagnosed the presence of conditional heteroskedasticity by analyzing the autocorrelation structure of squared returns and formally testing for ARCH effects using the Lagrange Multiplier test. We then proceeded to estimate an ARCH(q) model, selected using the Bayesian Information Criterion (BIC), and evaluated its adequacy through residual diagnostics. Subsequently, we estimated a GARCH(p,q) model, again guided by BIC, and verified that the model adequately captured both the conditional mean and volatility dynamics.

These steps, diagnosis, testing, specification search, estimation, and diagnostic checking, form a robust framework for identifying and validating volatility models. Following such a structured workflow ensures that the chosen model not only fits the data well but also respects key assumptions such as the absence of autocorrelation in standardized residuals.

References

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

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

- Nelson, D. B. (1991). *Conditional Heteroskedasticity in Asset Returns: A New Approach. Econometrica*, 59(2), 347–370.

- Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). *On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. Journal of Finance*, 48(5), 1779–1801.