Volatility Clustering and GARCH

Week 3 — Financial Management: Volatility, Risk, and AI

Yi-Hao Lai

2026 Spring

Today's Roadmap

1

The Story: Why This Matters

The VolTech Challenge

Last week: Alex built EWMA volatility dashboards.

This week's question: EWMA measures volatility *now*. Can we *predict* what it will be **tomorrow**?

Yes — because volatility has a memory.

Mandelbrot (1963)

"Large changes tend to be followed by large changes — of either sign — and small changes by small changes."

Today's Learning Objectives

By the end of this session, you will:

1. Detect volatility clustering using ACF of squared returns
2. Understand the ARCH model foundation
3. Master the GARCH(1,1) equation and its three parameters
4. Explain why GARCH improves over EWMA (mean reversion)
5. Estimate GARCH(1,1) in Python using the `arch` library

2

Detecting Volatility Clustering

Returns vs. Squared Returns

Raw returns r_t:

- Approximately uncorrelated
- Direction is unpredictable

Squared returns r_t^2:

- Strongly autocorrelated
- **Magnitude** is predictable

Key Insight

Returns are like the **direction** of the wind — hard to predict. Squared returns are like the **strength** — when a storm starts, it doesn't stop suddenly.

The Ljung–Box Test

Ljung–Box Statistic

$$Q(m) = T(T+2)\sum_{k=1}^{m}\frac{\hat{\rho}(k)^2}{T-k} \sim \chi^2(m) \quad \text{under } H_0$$

- H_0: No autocorrelation in squared returns
- Small p-value (< 0.05) $\Rightarrow$ reject H_0 $\Rightarrow$ volatility clustering confirmed
- In practice: always rejects for financial data

3

From ARCH to GARCH

Return Decomposition

The Building Block

$$r_t = \sigma_t z_t, \qquad z_t \sim \text{i.i.d.}(0, 1)$$

- σ_t = **conditional volatility** (time-varying "volume knob")
- z_t = standardized shock (random "direction")
- Goal: model how σ_t evolves over time

The ARCH Model (Engle, 1982)

ARCH(q)

$$\sigma_t^2 = \omega + \alpha_1 r_{t-1}^2 + \alpha_2 r_{t-2}^2 + \cdots + \alpha_q r_{t-q}^2$$

- Today's variance depends on recent past squared returns
- If yesterday's $|r_{t-1}|$ was large $\Rightarrow$ σ_t^2 is large $\Rightarrow$ **clustering**
- Problem: Needs many lags q for realistic decay $\Rightarrow$ too many parameters

GARCH(1,1): The Workhorse (Bollerslev, 1986)

GARCH(1,1)

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

Parameter	Role	Typical Value
ω	Floor / base variance	Small positive
α	Reaction to new shocks	0.05–0.15
β	Memory / persistence	0.80–0.95
$\alpha + \beta$	Total persistence	0.95–0.99

Why Only 3 Parameters?

GARCH(1,1) = ARCH(∞)

By recursive substitution:

$$\sigma_t^2 = \frac{\omega}{1-\beta} + \alpha \sum_{i=1}^{\infty} \beta^{i-1} \, r_{t-i}^2$$

- The β term encodes the **entire history** through exponential decay
- One GARCH lag replaces infinitely many ARCH lags
- Three numbers capture volatility dynamics for most financial assets

4

Mean Reversion and Half-Life

GARCH vs. EWMA: The Missing Piece

EWMA ($\omega = 0$):

$$\sigma_t^2 = (1-\lambda)\, r_{t-1}^2 + \lambda\, \sigma_{t-1}^2$$

- No mean reversion
- Volatility drifts after shocks

GARCH ($\omega > 0$):

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

- Mean-reverts to $\frac{\omega}{1-\alpha-\beta}$
- ω acts as a **spring**

Key Difference

After a shock, GARCH pulls volatility back to its long-run level. EWMA has no anchor — it drifts indefinitely.

Long-Run Variance and Half-Life

Long-Run (Unconditional) Variance

$$\bar{\sigma}^2 = \frac{\omega}{1 - \alpha - \beta}$$

Volatility Half-Life

$$h = \frac{\ln(0.5)}{\ln(\alpha + \beta)}$$

5

Maximum Likelihood Estimation

How Do We Find ω, α, β?

Log-Likelihood (Gaussian)

$$\mathcal{L} = -\frac{1}{2}\sum_{t=1}^{T}\left[\ln(2\pi) + \ln(\sigma_t^2) + \frac{r_t^2}{\sigma_t^2}\right]$$

MLE algorithm:

1. Initialize σ_1^2 (sample variance)
2. For candidate (ω, α, β): compute $\sigma_2^2, \ldots, \sigma_T^2$ recursively
3. Evaluate log-likelihood
4. Optimizer searches for the maximum

In practice: Use Student-t distribution ($\nu \approx 5$–7) to accommodate residual fat

6

Python Live Demo

Step 1: Visualize Clustering

```python
import numpy as np, pandas as pd
import yfinance as yf
import matplotlib.pyplot as plt

sp500 = yf.download("^GSPC",
    start="2015-01-01", end="2024-12-31")
prices = sp500["Close"].squeeze()
returns = 100 * np.log(
    prices / prices.shift(1)).dropna()

fig, axes = plt.subplots(2, 1, figsize=(14, 8))
axes[0].plot(returns, linewidth=0.5)
axes[0].set_title("Daily Log Returns (%)")
axes[1].plot(returns ** 2, color="red",
    linewidth=0.5)
axes[1].set_title("Squared Returns")
```

Step 2: Ljung–Box Test

```
from statsmodels.stats.diagnostic import (
    acorr_ljungbox)
from statsmodels.graphics.tsaplots import plot_acf

# ACF: raw returns (flat) vs. squared (persistent)
fig, axes = plt.subplots(1, 2, figsize=(14, 4))
plot_acf(returns.values, lags=30, ax=axes[0],
    title="ACF of Returns")
plot_acf(returns.values ** 2, lags=30, ax=axes[1],
    title="ACF of Squared Returns")

lb = acorr_ljungbox(returns ** 2, lags=10,
    return_df=True)
print(lb)  # Very small p-values!
```

Step 3: Fit GARCH(1,1)

```
from arch import arch_model

model = arch_model(returns, vol="Garch",
    p=1, q=1, dist="t")
result = model.fit(disp="off")
print(result.summary())

omega = result.params["omega"]
alpha = result.params["alpha[1]"]
beta  = result.params["beta[1]"]

print(f"alpha + beta = {alpha + beta:.4f}")
print(f"Long-run vol = {np.sqrt(omega /
    (1-alpha-beta) * 252):.2f}% ann.")
print(f"Half-life = {np.log(0.5) /
    np.log(alpha+beta):.1f} days")
```

Step 4: Conditional Volatility Plot

```python
cond_vol = result.conditional_volatility

fig, ax = plt.subplots(figsize=(14, 6))
ax.plot(returns.index, returns,
    color="gray", linewidth=0.4, alpha=0.6)
ax.plot(cond_vol.index, cond_vol,
    color="red", linewidth=1.5,
    label="GARCH cond. vol")
ax.plot(cond_vol.index, -cond_vol,
    color="red", linewidth=1.5)
ax.fill_between(cond_vol.index,
    cond_vol, -cond_vol, alpha=0.1,
    color="red")
ax.legend()
```

7

Application: COVID Crash

The Onset Problem

Test: Estimate GARCH pre-COVID (before Jan 2020), then forecast into the crash.

Result: GARCH was slow to react:

- Day 1: Market drops 3%. GARCH barely moves ($\beta \times \sigma_{t-1}^2$ dominates)
- Day 2–3: Still below realized volatility
- Day 5+: Finally catches up

⚠ Why?

GARCH is **reactive**, not proactive. It updates only *after* observing a large return. The VIX spikes *before* the worst returns.

Model Diagnostics

After GARCH fitting, check standardized residuals $\hat{z}_t = r_t/\hat{\sigma}_t$:

1. **Ljung–Box on $\hat{z}_t^2$**: p-values should be large $\Rightarrow$ clustering removed
2. **ACF of $\hat{z}_t^2$**: Should be flat (vs. raw squared returns)
3. **Skewness**: Still negative (≈ -0.3)

The Remaining Problem

Negative skewness $\Rightarrow$ GARCH cannot capture the **leverage effect**. It uses r_{t-1}^2, which erases the sign.

A -3% day and a $+3\%$ day are treated identically.

8

Key Takeaways

Six Things to Remember

1. **Volatility clusters**: Detectable via ACF of r_t^2 and the Ljung–Box test
2. **ARCH**: Today's variance depends on past squared returns. Needs many lags.
3. **GARCH(1,1)**: Three parameters (ω, α, β) capture most volatility dynamics
4. **Mean reversion**: $\bar{\sigma}^2 = \omega/(1 - \alpha - \beta)$ is what EWMA lacks $(\omega = 0)$
5. **MLE**: Parameters chosen to maximize data probability. Student-t improves tail fit.
6. **Limitation**: GARCH squares r_{t-1}, erasing the sign $\Rightarrow$ cannot capture the leverage effect

Mission 3: Building the GARCH Engine

Deliverables

1. Download Nikkei 225 and S&P 500 (2010–2024)
2. ACF plots and Ljung–Box test for each index
3. Fit GARCH(1,1) with Student-t for each index
4. Report: $\hat{\omega}$, $\hat{\alpha}$, $\hat{\beta}$, long-run vol, half-life
5. Diagnostic plots: ACF of squared standardized residuals
6. Compare GARCH vs. EWMA conditional volatility

Next Week Preview

Week 4: Asymmetric Volatility — GJR-GARCH

Alex discovered that a -3% day increases tomorrow's volatility **far more** than a $+3\%$ day.

GARCH is blind to this asymmetry.

Topics: Leverage effect, GJR-GARCH, EGARCH, News Impact Curve, likelihood ratio test