Measuring Volatility

Week 2 — Financial Management: Volatility, Risk, and AI

Yi-Hao Lai

2026 Spring

Today's Roadmap

1. The Story: One Ruler Is Not Enough
2. Rolling Window Volatility
3. EWMA: Exponentially Weighted Moving Average
4. Realized Volatility
5. The VIX: Market's Fear Gauge
6. Python Live Demo
7. Application: The Speed of Fear
8. Key Takeaways

1

The Story: One Ruler Is Not Enough

Kenji's Next Question

Last week: returns have fat tails, and the normal model fails.

Kenji's question: "How volatile is our portfolio *right now*?"

Dr. Lin: "You can't use a single ruler to measure risk at all times."

Today: four different rulers — and we find out which one warned fastest when COVID-19 hit.

The Challenge

Volatility is **not constant**. It clusters, spikes, and mean-reverts. We need measures that adapt.

Today's Learning Objectives

By the end of this session, you will:

1. Compute rolling window volatility with different window lengths
2. Apply EWMA and understand $\lambda = 0.94$
3. Define realized volatility from intraday data
4. Interpret the VIX as a forward-looking fear gauge
5. Compare all four measures during the COVID-19 crisis

2

Rolling Window Volatility

The Rolling Window Estimator

Rolling Window Volatility

$$\hat{\sigma}_t^{\text{roll}} = \sqrt{\frac{1}{N-1}\sum_{i=0}^{N-1}(r_{t-i} - \bar{r}_{t,N})^2}$$

Annualize: multiply by $\sqrt{252}$

- $N = 20$ (1 month): responsive but noisy
- $N = 60$ (3 months): moderate
- $N = 252$ (1 year): smooth but very slow

The Window Length Trade-Off

Short window ($N = 20$):

- Fast reaction to regime changes
- Noisy (high variance)
- Good for early warning

Long window ($N = 252$):

- Smooth and stable
- Slow to detect crises
- Good for long-term trends

Critical Limitation

All observations inside the window get **equal weight**. A 250-day-old return counts the same as yesterday's. When markets shift, the long window is dangerously slow.

The Ghosting Problem

When an extreme observation **leaves** the rolling window:

1. Market crash on Day 1, calm returns for $N - 1$ days
2. The crash contributes to volatility for exactly N days
3. On Day $N + 1$, the crash exits the window
4. Volatility estimate drops **suddenly** — an artifact, not a real change

EWMA solves this by fading observations gradually.

3

EWMA: Exponentially Weighted Moving Average

The EWMA Model

EWMA Variance (Recursive)

$$\hat{\sigma}_t^2 = \lambda \, \hat{\sigma}_{t-1}^2 + (1 - \lambda) \, r_{t-1}^2$$

- $\lambda \in (0, 1)$: the **decay factor**
- **RiskMetrics standard**: $\lambda = 0.94$ for daily data
- Weight on return i periods ago: $(1 - \lambda)\lambda^{i-1}$
- No hard window boundary — old data fades smoothly

Understanding λ

λ	Effective window	Character
0.90	$\sim$10 days	Very reactive
0.94	$\sim$17 days	**RiskMetrics standard**
0.97	$\sim$33 days	Smoother

Effective window $\approx 1/(1-\lambda)$

Key Advantage

EWMA has just one parameter (λ) and avoids the ghosting artifact. Old data fades geometrically.

EWMA as a Special GARCH

Preview of Week 3

EWMA is a special case of GARCH(1,1) with:

- No long-run variance term ($\omega = 0$)
- $\alpha + \beta = 1$ (integrated process)

This means EWMA assumes volatility has **no fixed "home base"** to return to.

GARCH fixes this by adding **mean reversion**.

4

Realized Volatility

Realized Volatility: The Gold Standard

Realized Volatility

$$RV_t = \sum_{j=1}^{M} r_{t,j}^2$$

- Uses M intraday returns (e.g., 5-minute intervals, $M = 78$)
- As $M \to \infty$: converges to true **integrated variance**
- Model-free — no distributional assumptions needed

⚠ Practical Note

Going beyond 5-minute sampling introduces **microstructure noise** (bid-ask

5

The VIX: Market’s Fear Gauge

The VIX Index

What is it?

- Market's expected 30-day S&P 500 volatility
- Derived from **option prices** (forward-looking)
- Annualized: VIX $= 20 \Rightarrow$ daily vol $\approx 1.26\%$

Typical ranges:

- 12–18: calm markets
- 20–30: elevated anxiety
- > 30: crisis mode
- > 50: extreme panic

COVID-19 Peak

March 16, 2020:
VIX $= \mathbf{82.69}$

It had already crossed 40 by late February — **before** the worst sell-off.

Implied vs. Realized Volatility

Variance Risk Premium:

$$\text{VRP} = \text{Implied Vol (VIX)} - \text{Realized Vol} > 0 \quad \text{(on average)}$$

- Option sellers earn a premium for bearing volatility risk
- One of the most robust findings in empirical finance
- VIX typically **leads** realized volatility during crises

Volatility Smile

Out-of-the-money put options show higher implied vol than calls — evidence the market prices in fat tails and negative skewness (Week 1 connection).

6

Python Live Demo

Step 1: Rolling Window Volatility

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

sp500 = yf.download("^GSPC",
    start="2018-01-01", end="2024-12-31")
prices = sp500["Close"].squeeze()
log_ret = np.log(prices / prices.shift(1)).dropna()

vol_20  = log_ret.rolling(20).std() * np.sqrt(252)
vol_60  = log_ret.rolling(60).std() * np.sqrt(252)
vol_252 = log_ret.rolling(252).std() * np.sqrt(252)
```

Step 2: EWMA Volatility

```python
def ewma_volatility(returns, lam=0.94):
    n = len(returns)
    var = np.zeros(n)
    var[0] = returns.iloc[0] ** 2
    for t in range(1, n):
        var[t] = (lam * var[t-1]
                  + (1-lam) * returns.iloc[t-1]**2)
    return pd.Series(
        np.sqrt(var) * np.sqrt(252),
        index=returns.index)

vol_ewma = ewma_volatility(log_ret, lam=0.94)
```

Step 3: Volatility Dashboard

```python
vix = yf.download("^VIX",
    start="2018-01-01", end="2024-12-31")
vix_close = vix["Close"].squeeze()

fig, ax = plt.subplots(figsize=(14, 6))
ax.plot(vol_20.index, vol_20,
    label="Rolling 20d", color="#1B3A5C")
ax.plot(vol_ewma.index, vol_ewma,
    label="EWMA (0.94)", color="#2D7D5E")
ax.plot(vix_close.index, vix_close/100,
    label="VIX/100", color="#C0392B")
ax.axvspan("2020-02-20", "2020-04-01",
    alpha=0.1, color="red")
```

7

Application: The Speed of Fear

COVID-19 Volatility Race

Measure	Feb 19	Mar 16	Response Time
Rolling 20-day	8.2%	89.7%	$\sim$2 weeks
Rolling 60-day	10.4%	54.1%	$\sim$4 weeks
Rolling 252-day	14.1%	23.8%	$\sim$8+ weeks
EWMA ($\lambda = 0.94$)	7.9%	83.4%	$\sim$1 week
VIX (implied)	14.4%	82.7%	Immediate

VIX led everything. EWMA was fastest backward-looking. 252-day was nearly useless.

Five Stylized Facts of Volatility

1. **Clustering**: large returns $\rightarrow$ large returns (GARCH captures this)
2. **Mean reversion**: extreme levels are temporary (GARCH adds this)
3. **Asymmetry**: bad news $\uparrow$ vol more than good news (GJR-GARCH, Week 4)
4. **Long memory**: autocorrelation decays slowly
5. **Co-movement**: global volatility spikes together

Design Specs

These five facts are the "specification sheet" for any good volatility model. Rolling windows capture none. EWMA captures clustering partially. GARCH captures more. Each week we add a layer.

8

Key Takeaways

Six Things to Remember

1. **Volatility is not constant**: it clusters, mean-reverts, and responds asymmetrically to news
2. **Rolling window**: simple but slow; equal weights and ghosting artifacts
3. **EWMA**: exponential decay, $\lambda = 0.94$, responsive and smooth
4. **Realized volatility**: gold standard when intraday data is available
5. **VIX**: forward-looking from option prices; often leads realized vol
6. **No one-size-fits-all**: match the measure to the application

Mission 2: Volatility Dashboard

Deliverables

1. Download Nikkei 225, S&P 500, FTSE 100 (2018–2024)
2. Compute rolling 20d, 60d, 252d + EWMA for each
3. Volatility dashboard plot per index (COVID-19 highlighted)
4. Summary table: pre-COVID, peak, response time
5. 200-word memo to Kenji recommending a measure

Bonus: $\lambda = 0.90, 0.94, 0.97$ comparison for Nikkei 225

Next Week Preview

Week 3: GARCH — Volatility Has a Memory

Priya asks: "EWMA reacts fast, but does it actually *predict* tomorrow's volatility?"

GARCH adds what EWMA lacks: **mean reversion** and a proper statistical framework.

Topics: GARCH(1,1), maximum likelihood estimation, conditional vs. unconditional variance