Value-at-Risk

Week 6 — Financial Management: Volatility, Risk, and AI

Yi-Hao Lai

2026 Spring

Today's Roadmap

1

The Story: One Number to Rule Them All

The VolTech Challenge

Scenario: Kenji's board demands a single risk number for the FSA regulatory report.

Dr. Lin: "This is the number every bank calculates every single day."

Can we compress all risk into one number?

The Core Question

What is the **maximum** we can lose tomorrow, with 99% confidence?

Today's Learning Objectives

By the end of this session, you will:

1. Define VaR and its three key ingredients
2. Compute VaR using three different methods
3. Explain why VaR has a critical blind spot
4. Understand Expected Shortfall as a superior alternative
5. Implement VaR and ES in Python with GARCH

2

Understanding Value-at-Risk

VaR: The Definition

Value-at-Risk

$$P(L_t > \text{VaR}_\alpha) = 1 - \alpha$$

Portfolio VaR

$$\text{VaR}_\alpha = -V \cdot q_{1-\alpha}(r_t)$$

where $q_{1-\alpha}$ is the $(1-\alpha)$-quantile of returns

Three Ingredients of VaR

Every VaR number requires:

1. **Confidence level** (α):
   - 99% for regulatory (Basel III)
   - 95% for internal risk management
2. **Holding period**:
   - 1 day for trading desks
   - 10 days for regulatory capital
3. **Return distribution**:
   - This is where modeling choices matter
   - This is where Weeks 1–5 pay off

Interpreting VaR

Example

"The 1-day 99% VaR is $10 million" means:

On **99 out of 100** trading days, the loss will be less than $10M.

On that **1 bad day**, you could lose $11M, $50M, or $100M — VaR doesn't tell you.

The Blind Spot

VaR tells you the **door** to the bad room. It doesn't tell you what's **inside**.

3

Method 1: Historical Simulation

Historical Simulation VaR

Procedure:

1. Collect the last W daily returns (e.g., $W = 500$)
2. Sort from worst to best
3. VaR $=$ the $(1 - \alpha) \times W$-th worst return
4. For 99% VaR with $W = 500$: the **5th worst** return

Advantages

- No distributional assumptions
- Captures fat tails naturally
- Easy to explain

Disadvantages

- Slow to adapt
- Past $\neq$ future
- Optimistic before crises

4

Method 2: Parametric VaR

Parametric VaR: Normal Distribution

Normal Parametric VaR

$$\text{VaR}_\alpha = -(\mu_t + z_{1-\alpha} \cdot \sigma_t)$$

- σ_t from GARCH model — **updates daily**
- $z_{0.01} = -2.326$ for 99% VaR
- $z_{0.05} = -1.645$ for 95% VaR

Key insight: VaR adapts dynamically because σ_t rises when volatility spikes and falls when markets calm.

Parametric VaR: Student-t Distribution

Student-t Parametric VaR

$$\text{VaR}_\alpha = -(\mu_t + t_{1-\alpha,\nu} \cdot s_t), \quad s_t = \sigma_t \sqrt{\frac{\nu - 2}{\nu}}$$

Distribution	$q_{0.01}$	VaR with $\sigma = 1\%$
Normal	-2.326	$\approx 2.33\%$
t $(\nu = 7)$	-2.998	$\approx 2.53\%$
t $(\nu = 5)$	-3.365	$\approx 2.60\%$

The t-distribution produces 8–12% higher VaR — real capital implications.

5

Method 3: Monte Carlo VaR

Monte Carlo Simulation VaR

Procedure:

1. Fit GARCH model and obtain conditional parameters
2. Simulate $N = 10{,}000$ one-step-ahead returns
3. Each simulated return $=$ one possible "tomorrow"
4. VaR $= (1 - \alpha)$ quantile of simulated distribution

Advantages

- Any distribution
- Path-dependent payoffs
- Gold standard for derivatives

Disadvantages

- Computationally expensive
- Results vary across runs
- "Garbage in, garbage out"

Three Methods Compared

Criterion	Historical	Parametric	Monte Carlo
Assumptions	None	Distribution	Distribution
Adaptability	Slow	Fast (GARCH)	Fast (GARCH)
Flexibility	Low	Medium	High
Computation	Fast	Fastest	Slowest
Regulatory use	Common	Common	Derivatives
Crisis response	Poor	Good	Good

Recommendation

For daily risk reporting: **Parametric VaR with GJR-GARCH + *t*-distribution.** Best balance of accuracy, speed, and adaptability.

6

Beyond VaR: Expected Shortfall

The Problem with VaR

VaR's Blind Spot

Two portfolios with the same VaR can have very different tail risk:

	Portfolio A	**Portfolio B**
99% VaR	$10M	$10M
Avg loss if VaR breached	$11M	$50M

VaR says these are equally risky. They are not.

Expected Shortfall (CVaR)

Expected Shortfall

$$\text{ES}_\alpha = -E[r_t \mid r_t < -\text{VaR}_\alpha]$$

Average loss *given that* the loss exceeds VaR

- Always $\geq$ VaR (more conservative)
- Basel III FRTB: ES at 97.5% replaces VaR for market risk
- **Coherent risk measure**: satisfies subadditivity

Subadditivity

$$\text{ES}(A + B) \leq \text{ES}(A) + \text{ES}(B)$$

ES/VaR Ratio: A Tail Diagnostic

Distribution	**ES/VaR (99%)**	**Interpretation**
Normal	≈ 1.15	Light tail
t $(\nu = 7)$	≈ 1.25	Moderate tail
t $(\nu = 5)$	≈ 1.35	Heavy tail
Empirical (crises)	1.5–$2.0+$	Very heavy tail

A high ES/VaR ratio means: when things go wrong, they go *really* wrong. The tail beyond VaR is severe.

7

Python Live Demo

Step 1: Data and GARCH Model Setup

```python
import numpy as np, pandas as pd
import yfinance as yf
from arch import arch_model
from scipy import stats

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()

# Fit GJR-GARCH(1,1) with t-distribution
model = arch_model(returns, vol="Garch",
    p=1, o=1, q=1, dist="t")
res = model.fit(disp="off")
nu = res.params["nu"]
print(f"Degrees of freedom: {nu:.2f}")
```

Step 2: Historical Simulation VaR

```python
def var_historical(returns, alpha,
                   window=500):
    """Rolling Historical Simulation VaR."""
    var_hs = pd.Series(
        index=returns.index, dtype=float)
    for i in range(window, len(returns)):
        hist = returns.iloc[i-window:i]
        var_hs.iloc[i] = -np.percentile(
            hist, (1 - alpha) * 100)
    return var_hs

var_hs_99 = var_historical(returns, 0.99)
print(f"HS VaR (99%): "
      f"{var_hs_99.iloc[-1]:.4f}%")
```

Step 3: Parametric VaR (Normal vs. t)

```python
def var_parametric(res, alpha, dist="t"):
    cond_vol = res.conditional_volatility
    mu = res.params.get("mu", 0)
    nu = res.params.get("nu", 30)
    if dist == "t":
        q = stats.t.ppf(1 - alpha, df=nu)
        scale = np.sqrt((nu - 2) / nu)
        var_p = -(mu + q * scale * cond_vol)
    else:
        q = stats.norm.ppf(1 - alpha)
        var_p = -(mu + q * cond_vol)
    return var_p

var_t = var_parametric(res, 0.99, "t")
var_n = var_parametric(res, 0.99, "normal")
print(f"Parametric-t VaR:  {var_t.iloc[-1]:.4f}%")
print(f"Parametric-N VaR:  {var_n.iloc[-1]:.4f}%")
```

Step 4: Expected Shortfall

```
def es_historical(returns, alpha,
                  window=500):
    es = pd.Series(
        index=returns.index, dtype=float)
    for i in range(window, len(returns)):
        hist = returns.iloc[i-window:i].values
        var_level = np.percentile(
            hist, (1 - alpha) * 100)
        tail = hist[hist <= var_level]
        es.iloc[i] = -tail.mean() \
            if len(tail) > 0 else np.nan
    return es

es_99 = es_historical(returns, 0.99)
print(f"ES (99%): {es_99.iloc[-1]:.4f}%")
print(f"ES/VaR ratio: "
      f"{es_99.iloc[-1]/var_hs_99.iloc[-1]:.2f}")
```

Step 5: COVID-19 VaR Comparison Plot

```
fig, ax = plt.subplots(figsize=(14, 6))
ax.plot(returns.index, returns,
    color="gray", lw=0.4, alpha=0.5,
    label="Returns")
ax.plot(var_hs_99.index, -var_hs_99,
    color="#D4A843", lw=1,
    label="HS VaR (99%)")
ax.plot(var_t.index, -var_t,
    color="#C0392B", lw=1,
    label="Parametric-t VaR (99%)")
ax.axvspan("2020-02-20", "2020-04-01",
    alpha=0.1, color="red",
    label="COVID-19")
ax.legend(loc="lower left", fontsize=8)
ax.set_title("VaR During COVID-19 Crisis")
plt.tight_layout()
```

8

Application: COVID-19 Crash

VaR Failure During COVID-19

What Happened in March 2020

1. **Historical Simulation** was dangerously slow — still using calm 2019 data, VaR was far too low. 10+ breaches in a single month.
2. **Parametric VaR (GJR-GARCH)** adapted within days — GARCH σ_t spiked quickly, VaR rose. Fewer breaches.
3. **Normal vs.** t: Normal VaR had 30–40% more breaches than the t-version.

Lesson: Conditional VaR (GARCH-based) adapts to regime changes. Unconditional VaR (HS) fails when it matters most.

Recommendation for Kenji's Board

Alex's Recommendation

Use **GJR-GARCH + Student-*t*** Parametric VaR because:

1. Adapts to changing volatility (GARCH mechanism)
2. Accounts for fat tails (*t*-distribution)
3. Captures leverage effect (GJR asymmetry)

Report **both VaR and ES** at 99% to the FSA.

Next question: How do we *prove* our model works?

Answer: Backtesting (Week 7).

9

Key Takeaways

Six Things to Remember

1. **VaR answers one question**: Maximum expected loss at a given confidence level over a given holding period
2. **Three methods**: Historical Simulation (simple), Parametric (GARCH-based), Monte Carlo (flexible)
3. **GARCH-based VaR adapts fastest**: Updates conditional variance daily, responds to regime changes
4. **Fat tails matter**: t-distribution produces more conservative and more accurate VaR
5. **VaR has a blind spot**: Says nothing about losses beyond the threshold; ES fills this gap
6. **ES is the new standard**: Basel III FRTB replaces VaR with ES for regulatory capital

Mission 6: VaR for the Pension Fund

Deliverables

1. Compute daily 99% VaR using all three methods for S&P 500 and Nikkei 225
2. Compute 99% Expected Shortfall (parametric-t)
3. Count VaR breaches in 2020 for each method
4. Create comparison plot with returns and VaR lines
5. 200-word memo: recommend one method and explain why

Bonus: Two-asset portfolio VaR (60/40 S&P/Nikkei)

Next Week Preview

Week 7: VaR Backtesting and Stress Testing

The FSA demands *proof* that our VaR model works.

David: "How do we know our VaR model is actually reliable?"

Topics: Kupiec test, Christoffersen test, traffic light system, stress testing, extreme scenarios