

Week 4: Volatility Modelling

From Stylised Facts to GARCH

Part I: Why Volatility Matters

Volatility: The Central Concept

Volatility : the tendency of asset prices to fluctuate : matters for:

- ▶ **Risk management**: How much can we lose?
- ▶ **Option pricing**: What is uncertainty worth?
- ▶ **Portfolio construction**: How do risks combine?
- ▶ **Regulatory capital**: How much buffer is needed?

Yet volatility is **not directly observable** : we must estimate it from price data.

Learning Objectives

By the end of this session, you should be able to:

1. **Explain** why volatility is a latent variable and what that implies for estimation
2. **Describe** the stylised facts of financial volatility and the feedback mechanisms that amplify them
3. **Derive** how ARCH and GARCH models capture volatility clustering
4. **Interpret** GARCH parameters and their economic implications
5. **Apply** asymmetric extensions to capture the leverage effect
6. **Evaluate** where GARCH-based risk models succeed and where they fail under stress

Volatility as a Latent Variable

Volatility is not merely *hard to measure* : it is fundamentally **unobservable**.

We have a sequence of daily returns. From these, we must infer the hidden variance process that generated them.

- ▶ We observe r_t ; we never observe σ_t directly
- ▶ Each return is one draw from a distribution with **unknown variance**
- ▶ GARCH acts as a **signal filter**, not an accounting formula
- ▶ Estimation uncertainty *compounds* forecasting uncertainty
- ▶ The hidden state evolves : what we infer today is already changing

Every GARCH estimate carries irreducible inference error. This is not a limitation of the model : it is the nature of the problem.

Not All Measurement Problems Are Equal

GDP is measured with error : but it *exists*. In principle, sum every transaction and you have it.

Volatility σ_t is different. No amount of better data reveals it.

Concept	Type of Problem
GDP	<i>Measurement error</i> : the true value exists; our instruments are imperfect
Inflation	<i>Index construction</i> : definition matters, but price changes are observable
σ_t	<i>Latent state</i> : the concept itself is model-dependent; no true value exists independently

Every GARCH estimate is a **belief**, not a fact. Two economists with identical data but different models produce different estimates : and neither is wrong.

Endogenous Risk: When Volatility Creates Itself

Here the textbook picture breaks down. GARCH measures volatility : but in modern financial systems, it also *causes* it.

Step	What Happens
1. Shock	Adverse move raises GARCH-estimated volatility
2. Signal	VaR limits are breached; positions must be cut
3. Deleveraging	Institutions sell simultaneously : all running the same models
4. Price impact	Coordinated selling drives prices lower
5. Amplification	Lower prices generate new shocks → back to Step 1

Danielsson (2002) called this the emperor's dilemma: **the map changes the territory**.

When all institutions share the same risk model, the model *becomes* the systemic risk (Danielsson, Shin, and Zigrand 2012).

Three Episodes Where the Feedback Was Visible

Each episode shares the same amplification logic : different costumes, same mechanism.

2008 Global Financial Crisis

Bank VaR models simultaneously breached limits → mass deleveraging → forced selling of CDOs, MBS, and equities → further price falls → VaR breach again. The system tightened precisely when it needed to breathe.

February 2018 : “Volmageddon”

Short-volatility ETNs (notably XIV) were *contractually required* to buy VIX futures when volatility spiked : a death spiral engineered into the product's own hedging rules. XIV lost 93% of its value in a single session.

March 2020 : COVID Crash

Margin calls forced simultaneous selling across equities, credit, and gold : assets with no fundamental correlation. The cross-asset contagion was a liquidity crisis, not a fundamental one. Correlations went to one.

Why Doesn't Knowing It Stop It?

- ▶ **Regulatory rules** *require* institutions to respond to VaR breaches, regardless of whether they recognise a feedback episode in progress
- ▶ **Incentive asymmetry** : the cost of breaching a limit is immediate and personal; the benefit of holding on is diffuse and speculative
- ▶ **Coordination failure** : one institution staying put helps no-one if twenty others are selling
- ▶ **Model homogeneity** : Basel III standardised risk models amplify synchronisation (Danielsson, Shin, and Zigrand 2012)

The financial system has embedded a regulatory structure that *forces* procyclical behaviour during the very crises it is designed to contain.

i Note

Macroprudential tools such as the Countercyclical Capital Buffer (CCyB) are a partial answer : releasing capital requirements in a downturn to reverse the deleveraging signal. The Bank of England deployed exactly this in March 2020.

From Why to How

We have now established three foundational arguments:

- ▶ Volatility is **latent** : every estimate is an inference about a hidden state, not a fact about the world
- ▶ Risk models can **amplify** the volatility they measure, creating self-reinforcing feedback under stress
- ▶ This feedback is **empirically visible** in real crisis data, with a distinct signature in each episode

The rest of this session builds the tools : ARCH, GARCH, and their asymmetric extensions : that allow us to model, estimate, and forecast that hidden state.

Understanding where these tools work well, and where the endogenous risk argument says they will fall short, is the animating question throughout.

Connection to Week 3

Last week we modelled the **mean** of returns using ARIMA.

This week we model the **variance** : which turns out to be predictable even when returns are not.

Week 3 Concept	Week 4 Extension
Stationarity	Conditional vs unconditional variance
ACF of returns	ACF of squared returns
ARMA for mean	GARCH for variance
$AR(p)$ = lagged values	$ARCH(q)$ = lagged squared errors

Key insight: Returns show no autocorrelation (efficient market), but volatility clusters.

Part II: Stylised Facts of Financial Volatility

What Markets Give Us

Before building models, understand what we're trying to capture:

1. **Volatility clustering** : calm follows calm, storms follow storms
2. **Fat tails** : extreme events more frequent than Normal predicts
3. **The leverage effect** : negative returns increase volatility more

These are **empirical facts**, not assumptions.

Volatility Clustering

Tsay (2010): “Volatility is not constant over time. There are periods of high volatility alternating with periods of relative calm.”

The ACF Tells the Story

Returns show **no autocorrelation** (consistent with efficiency). Squared returns show **strong autocorrelation** (volatility persistence).

Fat Tails (Leptokurtosis)

Financial returns have **fatter tails** than Normal at every horizon examined.

- ▶ Kurtosis typically 5–10 for daily returns (Normal = 3)
- ▶ The probability of a move exceeding 3σ is roughly 5–10 $\times$ higher than Normal predicts
- ▶ VaR at 99% confidence, assuming Normality, systematically underestimates the true loss quantile
- ▶ **Critically: GARCH with Normal errors improves but does not solve this :**
standardised GARCH residuals still exhibit excess kurtosis of 3–5

The practitioner fix is to assume Student- t or GED (Generalised Error Distribution) innovations. The t -distribution with 5–8 degrees of freedom fits most equity series well. This is the default in Bloomberg's GARCH estimation.

The Leverage Effect

Negative returns increase volatility **more** than positive returns of the same magnitude.

First documented by Black (1976).

Why the Leverage Effect?

Theory	Mechanism
Leverage hypothesis	Price falls → debt/equity rises → firm riskier
Volatility feedback	Expected volatility up → required return up → price falls
Risk premium	Higher expected volatility → higher risk premium
Behavioural	Investors react more strongly to losses
Margin constraints	Downturns trigger margin calls, forced selling

All likely operate simultaneously.

Part III: The ARCH Model

Engle's Key Insight

Engle (1982) asked: What if variance depends on recent shocks?

$$r_t = \mu + \varepsilon_t, \quad \varepsilon_t = \sigma_t z_t, \quad z_t \sim N(0, 1)$$

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

Interpretation: Today's volatility depends on recent *surprises*.

ARCH: The Building Block

Component	Meaning
σ_t^2	Conditional variance at time t
α_0	Baseline variance
α_i	Weight on shock from i periods ago
ε_{t-i}^2	Past squared shocks

Problem: Need many lags to capture persistence $\rightarrow$ many parameters.

Part IV: The GARCH Model

From ARCH to GARCH

Bollerslev (1986) extended ARCH by including lagged conditional variances:

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

This is **GARCH(1,1)** : one ARCH term, one GARCH term.

Brooks (2019): “A GARCH(1,1) model will be sufficient to capture the volatility clustering in the data.”

Parameter Interpretation

Parameter	Name	Interpretation
α_0	Constant	Long-run variance floor
α_1	ARCH term	Reaction to recent shocks
β_1	GARCH term	Persistence of volatility
$\alpha_1 + \beta_1$	Persistence	How long shocks affect volatility



Tip

For most financial assets, $\alpha_1 + \beta_1$ is close to (but less than) 1. Values above 0.9 are typical.

What Practitioners Actually See

Typical GARCH(1,1) estimates on daily returns from major markets:

Asset	$\hat{\alpha}$ (ARCH)	$\hat{\beta}$ (GARCH)	$\hat{\alpha} + \hat{\beta}$
S&P 500	0.07–0.09	0.88–0.91	~0.97
FTSE 100	0.07–0.10	0.87–0.90	~0.97
EUR/USD	0.04–0.06	0.92–0.94	~0.98
Gold	0.04–0.06	0.91–0.93	~0.97

High persistence ($\hat{\alpha} + \hat{\beta} \approx 0.97$) means a shock today still explains roughly 40% of excess volatility three weeks later.

During the 2020 COVID crisis, S&P 500 persistence estimates briefly approached 0.999 : not because the true process changed, but because extreme observations dominated the likelihood. This is the IGARCH warning signal in real data.

Why GARCH(1,1) Often Suffices

1. **Parsimony:** Only 3 parameters
2. **Memory:** Recursive structure uses entire history
3. **Mean reversion:** Volatility returns to long-run level

The unconditional variance (when $\alpha_1 + \beta_1 < 1$):

$$\sigma^2 = \frac{\alpha_0}{1 - \alpha_1 - \beta_1}$$

Forecasting with GARCH

Once fitted, GARCH(1,1) produces h -step ahead forecasts through a simple mean-reversion recursion.

Let $\bar{\sigma}^2 = \alpha_0 / (1 - \alpha_1 - \beta_1)$ be the long-run variance. Then:

$$E_t [\sigma_{t+h}^2] = \bar{\sigma}^2 + (\alpha_1 + \beta_1)^h (\sigma_t^2 - \bar{\sigma}^2)$$

With persistence $\alpha_1 + \beta_1 = 0.97$ and today's variance doubled:

Horizon	Excess variance remaining
1 day	97%
1 week (5 days)	86%
1 month (21 days)	53%
3 months (63 days)	15%
6 months (126 days)	2%

A doubling in volatility is not a doubling in variance. This is not a doubling in variance. This is not a doubling in variance.

GARCH(1,1) in Action

Part V: Estimating GARCH

How GARCH is Estimated: Maximum Likelihood

GARCH parameters are not computed from a formula : they are *searched for* numerically. The approach is **Maximum Likelihood Estimation**: find the parameters that make the observed return sequence most probable.

For GARCH with Normal innovations, the log-likelihood is:

$$\ell(\theta) = -\frac{1}{2} \sum_{t=1}^T \left[\ln(2\pi) + \ln(\sigma_t^2(\theta)) + \frac{\varepsilon_t^2}{\sigma_t^2(\theta)} \right]$$

At each t , the model asks: *how surprising was this return, given what I estimated the variance to be?*

A -5% return when $\hat{\sigma}_t = 1\%$ contributes enormous negative log-likelihood. The same return when $\hat{\sigma}_t = 5\%$ contributes far less : the model was already expecting turbulence.

Estimation in Practice: What Can Go Wrong

MLE is clean in theory. On real financial data, practitioners encounter:

- ▶ **Local optima** : the likelihood surface has multiple peaks; different starting values can yield materially different “optimal” parameters
- ▶ **Convergence to the boundary** : the optimiser finds $\hat{\alpha} + \hat{\beta} \rightarrow 1$, suggesting IGARCH; this is usually a warning sign, not a true finding
- ▶ **Scaling sensitivity** : working in decimal returns vs percentage returns shifts $\hat{\omega}$ by a factor of 10^4 ; always use percentage returns for numerical stability
- ▶ **Flat likelihood in crises** : when a crisis dominates the sample, the likelihood surface becomes very flat and optimisers stall

Practitioner discipline: always try multiple starting values; constrain $\alpha + \beta < 1$ explicitly; compare estimates across sub-samples. Two practitioners with the same data but different software can report different parameters : and both may be at local optima.

Diagnosing and Selecting a GARCH Model

A fitted GARCH model must pass two tests before being trusted.

Ljung-Box test on standardised residuals

If GARCH has captured all the volatility clustering, the standardised residuals $\hat{z}_t = \hat{\varepsilon}_t / \hat{\sigma}_t$ should be approximately i.i.d. Test both $\hat{z}_t$ (mean dynamics) and $\hat{z}_t^2$ (variance dynamics). Significant autocorrelation in $\hat{z}_t^2$ means the model has not captured all clustering : consider a higher-order specification or an asymmetric extension.

AIC and BIC for model comparison

When choosing between competing specifications:

$$\text{AIC} = -2\hat{\ell} + 2k \quad \text{BIC} = -2\hat{\ell} + k \ln(T)$$

where $\hat{\ell}$ is the maximised log-likelihood and k is the number of parameters. BIC penalises complexity more heavily : for most daily equity series, GARCH(1,1) wins the BIC competition against higher-order alternatives.

Part VI: Asymmetric GARCH Models

Capturing the Leverage Effect

Standard GARCH treats positive and negative shocks symmetrically.

But the leverage effect says this is wrong.

Solutions:

- ▶ **GJR-GARCH**: Add indicator for negative shocks
- ▶ **EGARCH**: Use logarithms and signed shocks

GJR-GARCH

Glosten, Jagannathan, and Runkle (1993) add an indicator for negative shocks:

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

Where $I_{t-1} = 1$ if $\varepsilon_{t-1} < 0$.

Parameter	Interpretation
γ	Additional impact of bad news
$\gamma > 0$	Leverage effect present
$\gamma = 0$	Symmetric GARCH

EGARCH

Nelson (1991) use logarithms:

$$\ln(\sigma_t^2) = \alpha_0 + \alpha_1 \left| \frac{\varepsilon_{t-1}}{\sigma_{t-1}} \right| + \gamma \frac{\varepsilon_{t-1}}{\sigma_{t-1}} + \beta_1 \ln(\sigma_{t-1}^2)$$

Advantages:

- ▶ No positivity constraints needed
- ▶ γ directly captures asymmetry
- ▶ Logarithmic specification often fits better

Part VII: Advanced Topics

GARCH-M: Risk-Return Relationship

The **GARCH-in-Mean** model embeds a direct test of the risk-return trade-off:

$$r_t = \mu + \delta \sigma_{t-1} + \varepsilon_t$$

If $\delta > 0$: higher expected volatility commands higher expected returns. The empirical evidence is genuinely mixed : and the *reason* it is mixed is instructive:

- ▶ At **daily horizons**: $\hat{\delta}$ is typically insignificant; short-run return noise swamps any premium signal
- ▶ At **monthly horizons**: evidence for $\delta > 0$ strengthens considerably : consistent with institutional investors pricing risk over planning horizons, not day by day
- ▶ In **high-volatility regimes**: the premium becomes more detectable, consistent with the 2008 and 2020 episodes where risk-premium logic re-emerged sharply

A more powerful approach uses the *variance risk premium* : the spread between implied (VIX^2) and expected realised variance. Bollerslev, Tauchen, and Zhou (2009) show this predicts future excess returns with R^2 around 4–7% at quarterly horizons, substantially stronger than GARCH-M's δ

Long Memory: IGARCH

When $\alpha_1 + \beta_1 = 1$, we have **Integrated GARCH** : shocks persist forever, and the unconditional variance no longer exists.

This sounds alarming. In practice, it is almost always a diagnostic signal rather than a true finding:

- ▶ **2008**: Full-sample S&P 500 estimates jumped toward 0.999 after Lehman, reflecting the structural break in the financial sector : not infinite volatility memory
- ▶ **2020**: COVID pushed persistence estimates above 0.999 for every major equity index; rolling 1-year estimation showed clear reversion by Q3 2020
- ▶ **The structural break interpretation**: a series with two regimes : calm and crisis : will produce a full-sample estimate with very high persistence even if each regime is genuinely stationary (Hamilton 1989)

If your estimated $\hat{\alpha} + \hat{\beta} \geq 0.999$, the first question is not “is this IGARCH?” but “is there a structural break in my sample?”

Multivariate: DCC

For portfolios, we need **covariances**, not just variances.

Dynamic Conditional Correlation (DCC) (Engle 2002):

1. Estimate univariate GARCH for each asset's volatility
2. Standardise residuals and model their *correlation* through a separate dynamic process
3. The conditional covariance matrix combines the two

$$H_t = D_t R_t D_t$$

where D_t is a diagonal matrix of GARCH-estimated volatilities and R_t is the time-varying correlation matrix.

In your Bloomberg lab, you observed the practical implication directly: the correlation between equities and gold was negative in January 2020 and strongly positive at the peak of the COVID sell-off. A static correlation matrix used in portfolio optimisation would have been deeply misleading. DCC captures exactly this instability.

Where GARCH Falls Short

No model is universally reliable. GARCH(1,1) has well-documented failure modes that every practitioner should know before relying on it.

Failure Mode	The Problem	Practitioner Response
Regime changes	One parameter set cannot fit both calm and crisis	Rolling windows; Markov regime-switching
Jump risk	GARCH is a diffusion model; overnight gaps are not captured	Jump-GARCH extensions
Illiquid markets	Non-trading creates spurious return autocorrelation	Higher-frequency data; bid-ask correction
Long-sample contamination	Crises push persistence to near-1, dominating the likelihood	Sub-sample estimation; crisis dummies

Part VIII: Practical Application

VIX: Implied vs Realised Volatility

VIX = market's expectation of 30-day volatility (from options)

Realised volatility = actual volatility observed over 30 days

Comparing them reveals:

- ▶ **Volatility risk premium** (VIX typically $>$ realised)
- ▶ **Market fear gauge** (VIX spikes in crises)
- ▶ **Forecasting ability** (VIX is forward-looking)

Volatility Risk Premium

The premium is positive on average : but inverts sharply during crisis episodes. The simulated series below illustrates both regimes; the real data in your Bloomberg lab will show the same pattern in 2008, 2018, and 2020.

Part IX: Connection to Statistical Science

Volatility and Uncertainty

Volatility modelling is fundamentally about **quantifying uncertainty**:

- ▶ **Conditional variance**: Uncertainty given what we know now
- ▶ **Forecasting**: How uncertain will the future be?
- ▶ **Risk management**: What's the range of possible outcomes?

This connects directly to Week 1's theme: data science as the study of variation and uncertainty.

From GARCH to Machine Learning

GARCH models are **parametric** : they assume a specific functional form.

Modern extensions relax this:

Classical	ML Extension
GARCH	Neural network volatility models
AR for returns	RNNs, LSTMs for sequences
Regime switching	Hidden Markov Models
Linear time series	Transformer architectures

The principles (stationarity, persistence, asymmetry) remain the same.

Part X: Beyond GARCH

GARCH vs Sequence Learning

Classical GARCH assumes a specific functional form for how past shocks feed into future variance. Can sequence learning do better?

GARCH Approach	Sequence Learning Alternative
GARCH(1,1) with fixed parameters	LSTM that learns volatility dynamics
Regime-switching GARCH	Attention mechanisms for regime detection
Assumes specific functional form	Learns flexible non-linear mappings
Interpretable parameters (α , β)	Black-box but potentially more flexible
3–5 parameters	Hundreds to thousands of parameters

In rigorous out-of-sample forecasting comparisons, simple GARCH models regularly match or outperform LSTM-based alternatives : the Makridakis principle applies: simpler models generalise better on shorter samples with high noise-to-signal ratios.

The Trade-off

GARCH advantages:

- ▶ Interpretable parameters with direct economic meaning
- ▶ Works with as little as 2–3 years of daily data
- ▶ Grounded in well-understood stylised facts
- ▶ Fast to estimate; robust to overfitting

Sequence learning advantages:

- ▶ Can capture complex non-linear patterns without functional form assumptions
- ▶ Learns from raw data directly
- ▶ Potentially captures regime dynamics that GARCH must specify explicitly

The choice depends on your objectives: explanation vs pure prediction, data availability, whether interpretability is a regulatory or communication requirement. In most practitioner contexts, GARCH is the baseline that ML must demonstrably beat : and often does not.

Part XI: Labs and Assessment

This Week's Lab Structure

Homework (Colab : complete before class):

Simulation-first approach: fit GARCH(1,1) to synthetic data where you know the true parameters, then apply the same workflow to real SPY returns from the shared Bloomberg database. Compare symmetric vs GJR-GARCH and evaluate model diagnostics.

In-Class Lab (Bloomberg Terminal : forensic investigation):

Working across three crisis episodes, you will use Bloomberg to ask whether the endogenous risk feedback discussed in the lecture is *detectable in real data*:

- ▶ **2017 baseline** : calibrate the normal VIX/realised relationship
- ▶ **2008 GFC** : trace the VaR-deleveraging spiral in VIX and sector returns
- ▶ **Volmageddon (Feb 2018)** : read the VIX futures term structure to distinguish mechanical from fundamental
- ▶ **COVID crash (Mar 2020)** : measure the cross-asset correlation breakdown directly

Key Takeaways

1. **Volatility is latent** : GARCH filters; it does not calculate. Every estimate is a belief about a hidden state
2. **Risk models can amplify risk** : the endogenous feedback loop is empirically visible in 2008, 2018, and 2020
3. **Stylised facts** constrain sensible models : clustering, fat tails, and the leverage effect are non-negotiable targets
4. **GARCH(1,1) is the workhorse** : three parameters, high persistence, rarely need more
5. **MLE is the engine** : but local optima, boundary solutions, and scaling issues are genuine practical hazards
6. **Know the failure modes** : regime breaks, jumps, illiquidity, and long-sample contamination each require specific responses
7. **Correlations are not stable** : DCC and the COVID evidence show diversification disappears exactly when it is needed most

Directed Learning

Core reading and tasks

Read Tsay (2010) Chapter 3 (volatility models) and Brooks (2019) Chapter 9 (GARCH in practice). Complete the homework lab before the Bloomberg session : the in-class investigation builds directly on it.

For the Bloomberg session

Before you arrive, form a view on each episode: which do you expect to show the highest persistence, the largest VRP sign reversal, and the clearest cross-asset contagion? Having a prior makes the data analysis active rather than passive.

Optional extension

Read Danielsson (2002) in full : it is short (25 pages) and more readable than most academic finance papers. Note where his 2002 critique did and did not anticipate the episodes you examined in the lab.

Exit Ticket

Answer these before you leave : one sentence each is sufficient:

1. **Latency:** Why is estimating σ_t fundamentally different from measuring GDP? What does “irreducible inference error” mean in practice?
2. **Persistence:** Your GARCH estimate gives $\hat{\alpha} + \hat{\beta} = 0.97$. A shock today doubles the conditional variance. Roughly how many trading days until variance returns to within 10% of its long-run level?
3. **The feedback loop:** Identify the single step in the endogenous risk spiral that regulatory capital reform could most plausibly interrupt. Justify your choice in one sentence.
4. **Failure modes:** You are fitting GARCH to a small-cap UK equity series from 2005 to 2024. Name the two failure modes from today's session most likely to affect your estimates.

References

- Bekaert, Geert, and Guojun Wu. 2000. "Asymmetric Volatility and Risk in Equity Markets." *Review of Financial Studies* 13 (1): 1–42.
<https://doi.org/10.1093/rfs/13.1.1>.
- Black, Fischer. 1976. "Studies of Stock Price Volatility Changes." *Proceedings of the 1976 Meetings of the American Statistical Association, Business and Economics Statistics Section*, 177–81.
- Bollerslev, Tim. 1986. "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* 31 (3): 307–27.
[https://doi.org/10.1016/0304-4076\(86\)90063-1](https://doi.org/10.1016/0304-4076(86)90063-1).
- Bollerslev, Tim, George Tauchen, and Hao Zhou. 2009. "Expected Stock Returns and Variance Risk Premia." *Review of Financial Studies* 22 (11): 4463–92.
<https://doi.org/10.1093/rfs/hhp008>.
- Brooks, Chris. 2019. *Introductory Econometrics for Finance*. 4th ed. Cambridge, UK: Cambridge University Press.