FINANCIAL ECONOMETRICS

The GARCH(1,1) model and volatility dynamics

How a three-parameter recursion reproduces the volatility clustering of financial returns, when it is stationary, and how it is estimated by quasi-maximum likelihood.

Intermediate · 9 min · GARCH · ARCH · volatility · conditional variance · Example — not reviewed by the author

Sheet
by Simon-Pierre Boucher

sheet.spboucher.ai/s/the-garch-1-1-model-and-volatility-dynamics
November 17, 2026

Financial returns are close to unpredictable in level, but their **magnitude** is not: turbulent periods follow turbulent periods. Engle's ARCH model [1], generalised to GARCH by Bollerslev [2], captures this by making today's variance a function of yesterday's shock and yesterday's variance.

01 Three parameters, one recursion

ƒ RETURN AND INNOVATION

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

The shock ε_t is a standardised noise z_t scaled by the conditional standard deviation σ_t, which is known at date $t-1$.

r_t	(log) return over period t	% per period
μ	conditional mean, constant here	
ε_t	innovation: unexpected return	
σ_t	conditional standard deviation (volatility), a function of information at $t-1$	
z_t	i.i.d. noise with zero mean and unit variance	

ƒ GARCH(1,1) CONDITIONAL VARIANCE

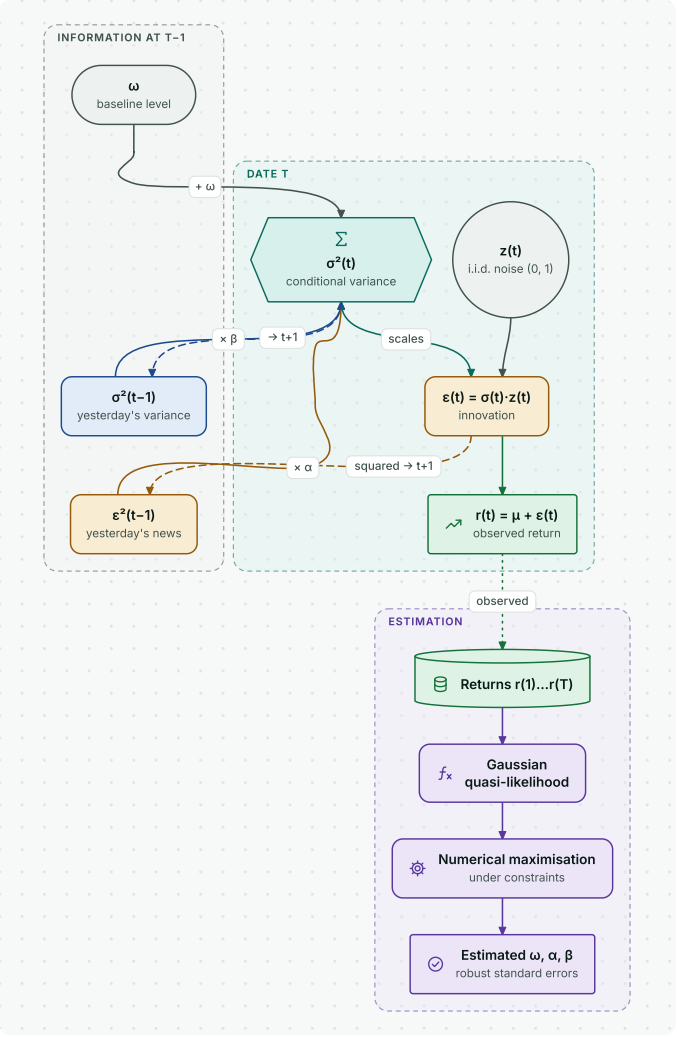
$$\sigma_t^2 = \omega + \alpha \, \varepsilon_{t-1}^2 + \beta \, \sigma_{t-1}^2$$

Today's variance combines a baseline level, yesterday's **news** (squared shock) and the **memory** of yesterday's variance. Setting $\beta = 0$ gives back ARCH(1).

σ_t^2	conditional variance of ε_t given the past	%² per period
ω	intercept, $\omega > 0$	
α	reaction to the previous shock (ARCH effect), $\alpha \geq 0$	
β	inertia of the variance (GARCH effect), $\beta \geq 0$	

Sources: 2

The volatility feedback loop



Solid arrow: computation at date t

Dashed arrow: carried over to the next date

Amber: shock channel (α)

Blue: memory channel (β)

Violet: estimation from the data

At each date the variance is computed from the past, then a shock is drawn; that shock and that variance feed the next date.

LEVEL

ω

Sets, together with $\alpha + \beta$, the average level of the variance. Must be strictly positive.

REACTION

α

How strongly a recent shock lifts the variance. A high α gives a jumpy volatility.

MEMORY

β

Share of yesterday's variance carried into today. A high β gives a smooth volatility.

PERSISTENCE

α + β

How fast a variance shock dies out. Close to 1: turbulent episodes last.

02 Persistence, long-run variance and mean reversion

ƒ UNCONDITIONAL VARIANCE

$$\bar{\sigma}^2 = \mathbb{E}\left[\varepsilon_t^2\right] = \frac{\omega}{1-\alpha-\beta}, \qquad \alpha+\beta<1$$

Take expectations of the recursion using $\mathbb{E}[\varepsilon_{t-1}^2] = \mathbb{E}[\sigma_{t-1}^2] = \bar{\sigma}^2$. The formula only makes sense when $\alpha + \beta < 1$.

$\bar{\sigma}^2$	long-run variance towards which σ_t^2 reverts	%² per period
$\alpha+\beta$	persistence of the variance	

Sources: 2

ƒ FORECAST AT HORIZON H

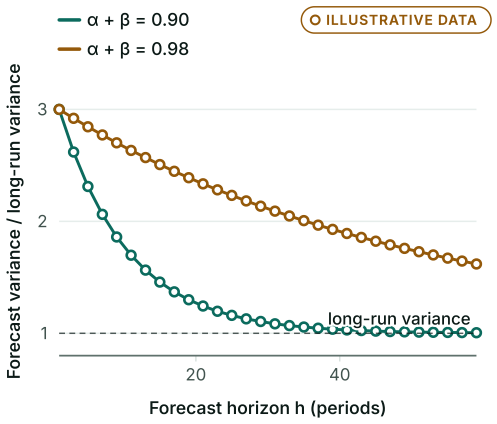
$$\mathbb{E}_t\left[\sigma_{t+h}^2\right] = \bar{\sigma}^2 + (\alpha+\beta)^{h-1}\left(\sigma_{t+1}^2 - \bar{\sigma}^2\right), \qquad h_{1/2} = \frac{\ln(1/2)}{\ln(\alpha+\beta)}$$

ƒx Result derived in this sheet

Obtained by iterating $\mathbb{E}_t[\sigma_{t+h}^2] = \omega + (\alpha+\beta)\,\mathbb{E}_t[\sigma_{t+h-1}^2]$. The gap to the long-run variance decays geometrically; $h_{1/2}$ is the number of periods after which it is halved.

$\mathbb{E}_t[\cdot]$	expectation conditional on information available at t	
σ_{t+1}^2	next-period variance, known at t	
h	forecast horizon	periods
$h_{1/2}$	half-life of a variance shock	periods

Forecast variance reverting to its long-run level



Curves computed from the forecast formula, starting from a variance three times the long-run variance (normalised to 1). Values chosen to show the shape, not estimated from data.

DEFINITION

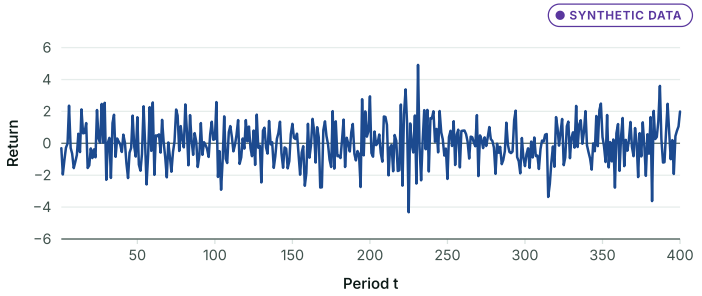
Volatility clustering

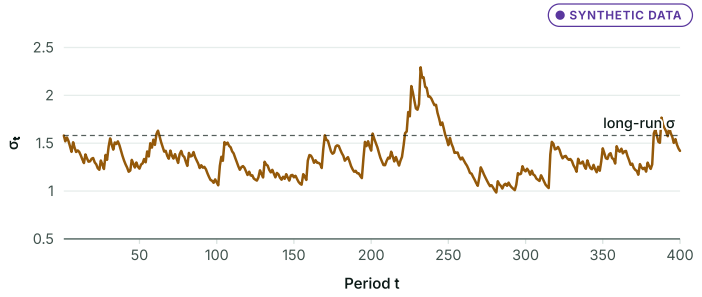
Volatility clustering means that large price changes tend to be followed by large changes, of either sign, and small ones by small ones. In a GARCH(1,1), returns are not autocorrelated but their squares are: ε_t^2 follows an ARMA(1,1) with autoregressive coefficient $\alpha + \beta$ [2].

Simulate a GARCH(1,1) path

Synthetic data — seed 7

$r_t = \sigma_t z_t$ with $z_t \sim \mathcal{N}(0,1)$ and $\sigma_t^2 = \omega + \alpha\, r_{t-1}^2 + \beta\, \sigma_{t-1}^2$. The long-run variance $\omega/(1-\alpha-\beta)$ exists only when $\alpha + \beta < 1$: beyond that, β is reduced automatically. Gaussian innovations, zero mean: real returns have fatter tails.





Persistence α + β	Unconditional variance ω/(1−α−β)	Unconditional standard deviation	Sample variance
0.98	2.5	1.5811	1.6963

Shock half-life (periods)
34.3

Parameters: Constant ω = 0.05; Shock reaction α = 0.08; Memory β = 0.9; Periods n = 400

Synthetic data: returns drawn from the model with the chosen parameters and a fixed seed. Move $\alpha + \beta$ towards 1 to see the clusters lengthen.

03 Positivity and stationarity

Parameter constraints of a GARCH(1,1) and what they guarantee.

PROPERTY	CONDITION	CONSEQUENCE
Positivity	$\omega > 0,\ \alpha \geq 0,\ \beta \geq 0$	$\sigma_t^2 > 0$ at every date, whatever the path of shocks.
Covariance stationarity	$\alpha + \beta < 1$	Finite unconditional variance $\omega/(1-\alpha-\beta)$; variance forecasts converge [2].
Strict stationarity	$\mathbb{E}\left[\ln(\beta + \alpha z_t^2)\right] < 0$	Weaker than the previous condition: it allows $\alpha + \beta \geq 1$ [3].
Finite fourth moment (Gaussian z_t)	$3\alpha^2 + 2\alpha\beta + \beta^2 < 1$	Kurtosis $3\,\dfrac{1-(\alpha+\beta)^2}{1-(\alpha+\beta)^2-2\alpha^2} > 3$: tails fatter than the normal distribution [2].

WARNING

The boundary case α + β = 1

When $\alpha + \beta = 1$ (integrated GARCH, or IGARCH), the unconditional variance is infinite and a variance shock never fades from the forecasts. Yet the process remains **strictly stationary** when $\omega > 0$ [3]: "infinite persistence" in expectation does not mean an explosive path.

04 Quasi-maximum likelihood

GAUSSIAN QUASI-LOG-LIKELIHOOD

$$\hat{\theta} = \arg\max_{\theta} \sum_{t=1}^{T} \ell_t(\theta), \qquad \ell_t(\theta) = -\frac{1}{2}\left[\ln(2\pi) + \ln\sigma_t^2(\theta) + \frac{\varepsilon_t^2(\theta)}{\sigma_t^2(\theta)}\right]$$

The likelihood is maximised **as if** z_t were Gaussian. If the conditional mean and variance are correctly specified, the estimator remains consistent even when z_t is not; only the standard errors need correcting [4].

θ	parameter vector $(\mu, \omega, \alpha, \beta)$
$\ell_t(\theta)$	contribution of observation t to the log-likelihood
$\varepsilon_t(\theta)$	residual $r_t - \mu$
$\sigma_t^2(\theta)$	variance obtained by running the recursion from a starting value
T	number of observations

Sources: 4

Filter the mean

Compute the residuals $\varepsilon_t = r_t - \mu$ (or those of a richer mean model).

Initialise and run the recursion

Set σ_1^2, for instance to the sample variance of the residuals, then compute σ_t^2 recursively.

Maximise the quasi-likelihood

Numerical optimisation under the constraints $\omega > 0$, $\alpha, \beta \geq 0$.

Robust standard errors

Use the Bollerslev–Wooldridge "sandwich" covariance matrix, valid without normality [4].

Check

Standardised residuals $\hat{z}_t = \hat{\varepsilon}_t/\hat{\sigma}_t$ and their squares should no longer be autocorrelated.

ASSUMPTION

What GARCH(1,1) assumes

The variance reacts to the **squared** shock: a rise and a fall of the same size have the same effect. The baseline model therefore misses the asymmetry often observed in equities; extensions (EGARCH, GJR-GARCH, among others) were proposed for it [5].

05 Check your understanding

QUESTION

A GARCH(1,1) has parameters $\omega = 0.02$, $\alpha = 0.08$ and $\beta = 0.90$. What is its unconditional variance?

A. 1

B. 0.02

C. 0.98

D. It is infinite

SHOW THE ANSWER

A — Correct. $\omega/(1-\alpha-\beta) = 0.02/0.02 = 1$.

B — Incorrect. That is ω, the intercept of the recursion, not the long-run variance.

C — Incorrect. That is the persistence $\alpha + \beta$.

D — Incorrect. It would be infinite only if $\alpha + \beta \geq 1$.

Confusing conditional and unconditional variance

σ_t^2 changes at every date and depends on the past; $\bar{\sigma}^2$ is a constant. A stationary GARCH is **unconditionally homoskedastic** and conditionally heteroskedastic.

Believing the model predicts the sign of returns

GARCH describes the expected magnitude, not the direction: $\mathbb{E}_{t-1}[\varepsilon_t] = 0$ whatever σ_t.

Reading α + β ≥ 1 as an explosive path

The condition $\alpha + \beta < 1$ is about the existence of the variance. Strict stationarity obeys a weaker condition [3].

Using Gaussian maximum-likelihood standard errors uncorrected

If the standardised innovations are not Gaussian, the inverse Hessian no longer estimates the variance of the estimator correctly: the robust form is required [4].

Forgetting the time unit

Parameters depend on the sampling frequency: a GARCH estimated on daily returns does not apply as is to monthly returns.

References

[1] Robert F. Engle (1982). *Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation.* Econometrica, 50(4), 987. doi:10.2307/1912773

[2] Tim Bollerslev (1986). *Generalized autoregressive conditional heteroskedasticity.* Journal of Econometrics, 31(3), 307–327. doi:10.1016/0304-4076(86)90063-1

[3] Daniel B. Nelson (1990). *Stationarity and Persistence in the GARCH(1,1) Model.* Econometric Theory, 6(3), 318–334. doi:10.1017/S0266466600005296

[4] Tim Bollerslev, Jeffrey M. Wooldridge (1992). *Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances.* Econometric Reviews, 11(2), 143–172. doi:10.1080/07474939208800229

[5] Christian Francq, Jean-Michel Zakoian (2019). *GARCH Models.* Wiley. doi:10.1002/9781119313472