

Recovering Time-Varying Correlation with DCC-GARCH: A Controlled Study of Correlation Tracking and Portfolio-Risk Misstatement

Eugen Soloviov

July 2026

Abstract

Sample correlation is an average over a window during which the true dependence structure was moving. We ask, under known ground truth, how much a dynamic model recovers that a static or rolling estimate throws away. We build seeded synthetic data-generating processes in which the conditional correlation follows a prescribed calm/crisis/recovery path (equicorrelation 0.3 in calm, ramping to 0.9 in the crisis, partially recovering to 0.5) on top of GARCH(1,1) volatilities, implement the two-step Dynamic Conditional Correlation (DCC) estimator of Engle from scratch on top of the univariate `arch` library, and measure recovery directly. First, DCC tracks the correlation path with a crisis-window mean absolute error of 0.028 on a two-asset process, against 0.208 for the static full-sample correlation — a factor of roughly 7 — while beating every fixed rolling window on crisis-window error and on overall root-mean-square error; no single rolling window is simultaneously best overall and in the crisis, and DCC is. On a stationary DCC simulation the estimator recovers the true parameters $(a, b) = (0.05, 0.90)$ as $(0.049, 0.914)$. Second, a one-step 95% Value-at-Risk backtest on an equal-weight portfolio shows the cost of ignoring correlation dynamics: static covariance breaches at 0.111 inside the crisis — 2.2 times the nominal 0.05 — while its full-sample breach rate of 0.049 looks perfectly calibrated, hiding the failure exactly when it matters; DCC holds the crisis breach rate at 0.063, close to the 0.060 delivered by the true covariance. Third, a DCC dynamic hedge ratio cuts hedged-spread variance by a fraction 0.208 inside the crisis relative to a static ordinary-least-squares hedge, and tracks the true time-varying beta with a crisis error of 0.035 against 0.248. Finally, a parameter-count comparison shows why DCC scales where full multivariate GARCH cannot (8 versus 11 versus 21 free parameters at $d = 2$; 152 versus 6275 versus 3.25×10^6 at $d = 50$ for DCC, BEKK, and VECH), and the estimated (a, b) stay stable as the cross-section grows to $d = 10$. The deliverable is the calibrated method and the quantified cost of static correlation, not a market claim: the processes are synthetic by design.

1 Introduction

Ask for the correlation between two assets and you typically get a single number computed over some window nobody remembers choosing. That number averages over a period in which the true dependence structure was moving constantly. In calm markets assets drift apart; in a stress event they lock together, and the diversification a portfolio was built on evaporates at the moment it is most needed. Correlation is not a constant that occasionally gets estimated badly — it is a time series with its own clustering and regimes, and treating it as a scalar is the multivariate cousin of assuming constant volatility.

The econometrics literature offers a ladder of answers. Modeling the full conditional covariance matrix directly — VECH (Bollerslev et al., 1988) or BEKK (Engle and Kroner, 1995) — is general but its parameter count grows quadratically or quartically in the number of assets d , confining it to a handful of series in practice. The Constant Conditional Correlation model of Bollerslev (1990)

factors the covariance into univariate volatilities and a correlation matrix, buying tractability at the price of freezing the correlation. Engle’s Dynamic Conditional Correlation (DCC) model (Engle, 2002) keeps that factorization but lets the correlation matrix evolve through a two-parameter recursion, estimated in two steps (Engle and Sheppard, 2001); Cappiello et al. (2006) add downside asymmetry. DCC is attractive precisely because it scales: the volatility side is d independent univariate fits and the correlation side has only two free parameters, whatever d is.

What is easy to assert but harder to demonstrate is that this machinery actually recovers a moving correlation better than the naive alternatives, and that the difference matters for a downstream decision. Demonstrating it requires knowing the truth. That is the gap this paper fills. We do not claim a new estimator and we make no claim about real markets; the contribution is calibration evidence under controlled ground truth — synthetic, seeded data-generating processes whose conditional correlation path we prescribe by construction — together with a self-contained DCC implementation built on the univariate `arch` library (Sheppard et al., 2024), which fits univariate GARCH (Engle, 1982; Bollerslev, 1986) but has no multivariate layer. Implementing the DCC layer is part of the contribution.

Concretely, we run four experiments (Section 4) against the estimator of Section 3:

1. **Correlation tracking.** On two- and three-asset processes with a known calm/crisis/recovery correlation path, DCC tracks the path with lower crisis-window error than any fixed rolling window and far lower than static correlation, and on a stationary DCC simulation it recovers the true parameters (a, b) .
2. **Static-risk misstatement.** A one-step portfolio Value-at-Risk backtest shows that static and rolling covariance under-state risk precisely during the correlation spike, breaching far above the nominal rate inside the crisis, while DCC stays close to nominal.
3. **Dynamic hedge ratio.** A DCC time-varying hedge ratio reduces hedged-spread variance relative to a static hedge and tracks the true time-varying beta, most visibly inside the crisis.
4. **Dimensionality.** A parameter-count comparison quantifies DCC’s linear scaling against BEKK’s quadratic and VECH’s quartic growth, and an empirical note confirms (a, b) estimation stays stable as d grows.

This study accompanies a `marketmaker.cc` blog post. All numbers derive from one seeded run of a public harness; a companion script verifies every numeric claim in this manuscript against the run’s JSON output.

2 The data-generating processes and ground truth

Every experiment draws from a synthetic process with a KNOWN conditional covariance path, so that estimation error can be measured rather than argued about. Let $r_t \in \mathbb{R}^d$ be the return vector at time t . We generate

$$r_t = D_t L_t e_t, \quad e_t \sim \mathcal{N}(0, I_d), \quad L_t L_t' = R_t, \quad (1)$$

where $D_t = \text{diag}(\sigma_{1,t}, \dots, \sigma_{d,t})$ collects per-asset conditional volatilities and L_t is the Cholesky factor of the true conditional correlation matrix R_t . The true conditional covariance is therefore $H_t = D_t R_t D_t$, known exactly at every step.

Correlation path. R_t is an equicorrelation matrix with a single time-varying off-diagonal ρ_t following a piecewise path over $T = 2000$ observations: a calm regime at $\rho = 0.3$, a linear ramp up to a crisis level of $\rho = 0.9$, a crisis plateau, a recovery ramp down, and a final calm regime at a partially recovered $\rho = 0.5$. The crisis window — the stress episode of elevated correlation — spans observations 600 to 1300. This is the ground-truth path every estimator in Section 5.1 is scored against.

Volatility path. Each $\sigma_{i,t}^2$ follows a GARCH(1,1) recursion $\sigma_{i,t}^2 = \omega_t + \alpha \epsilon_{i,t-1}^2 + \beta \sigma_{i,t-1}^2$ with $\alpha = 0.10$ and $\beta = 0.85$, sharing a variance intercept ω_t that rises from a calm 0.05 to a crisis 0.20 over the same schedule as the correlation. Volatility and correlation therefore spike together, as they do in real drawdowns: the crisis is a high-volatility, high-correlation regime, and the resulting elevation of true portfolio variance is what the risk experiment probes.

Stationary DCC process. To test parameter recovery we also generate standardized residuals directly from the DCC recursion itself (Section 3) with fixed, known parameters (a, b) and a fixed target correlation, over $T = 4000$ observations. Here the ground truth is the pair (a, b) , and the question is whether the estimator recovers it.

All random streams use `numpy`'s seeded generator, so every number is bit-reproducible by one command.

3 The DCC-GARCH estimator

3.1 Factorization

Following Engle (2002), the conditional covariance is factored as

$$H_t = D_t R_t D_t, \quad (2)$$

with D_t the diagonal matrix of conditional standard deviations from d univariate GARCH models and R_t the conditional correlation matrix. Elementwise, $h_{ij,t} = \rho_{ij,t} \sigma_{i,t} \sigma_{j,t}$: the conditional covariance of two assets is their dynamic correlation times each of their dynamic volatilities. Bollerslev (1990) froze R_t to a constant; DCC lets it move.

3.2 Step 1: univariate volatilities

For each asset i we fit a univariate GARCH(1,1) with the `arch` library (Sheppard et al., 2024), obtaining the conditional volatility $\sigma_{i,t}$ and the standardized residual

$$z_{i,t} = \frac{\epsilon_{i,t}}{\sigma_{i,t}}. \quad (3)$$

Each $z_{i,t}$ has approximately unit conditional variance; whatever co-movement remains in the vector $z_t = (z_{1,t}, \dots, z_{d,t})'$ is dependence, not a volatility artifact.

3.3 Step 2: the correlation recursion

We model an auxiliary process Q_t , a symmetric positive-definite matrix, with a scalar GARCH-like recursion driven by the outer products of the standardized residuals:

$$Q_t = (1 - a - b) \bar{Q} + a z_{t-1} z_{t-1}' + b Q_{t-1}, \quad (4)$$

where $\bar{Q}$ is the unconditional correlation matrix of the standardized residuals — plugged in as a sample estimate rather than optimized, the *correlation-targeting* device that confines the dimensionality to $\bar{Q}$ and leaves only a and b to estimate. The constraints $a, b > 0$ and $a + b < 1$ keep Q_t positive definite and mean-reverting, exactly analogous to a univariate GARCH(1, 1). Because Q_t has a diagonal that is not exactly one, we normalize it to a proper correlation matrix,

$$R_t = \text{diag}(Q_t)^{-1/2} Q_t \text{diag}(Q_t)^{-1/2}, \quad \rho_{ij,t} = \frac{q_{ij,t}}{\sqrt{q_{ii,t} q_{jj,t}}}, \quad (5)$$

which is symmetric positive definite with unit diagonal at every step by construction. The whole correlation layer has only two parameters, regardless of whether $d = 2$ or $d = 50$.

3.4 The quasi-log-likelihood

Under $\epsilon_t | \mathcal{F}_{t-1} \sim \mathcal{N}(0, H_t)$ the Gaussian log-likelihood separates into a volatility part depending only on D_t and a correlation part depending on R_t (Engle, 2002). Maximizing the volatility part is exactly the d univariate fits of Step 1; the correlation part, given the standardized residuals, is the two-parameter objective

$$\mathcal{L}^C(a, b) = -\frac{1}{2} \sum_{t=1}^T \left(\log |R_t| + z_t' R_t^{-1} z_t \right), \quad (6)$$

which we maximize over (a, b) with a derivative-free multi-start search. It is a *quasi*-likelihood: the two-step estimator is consistent but not fully efficient, and correct standard errors need the Engle and Sheppard (2001) correction; for recovering the correlation path the point estimates are what matter. Timing is strictly causal: in the recursion R_t is scored against z_t and only then is Q advanced with $z_t z_t'$ for step $t + 1$, so R_t uses information through $t - 1$ only — the off-by-one that silently inflates in-sample fit is avoided by construction.

3.5 Baselines

Against DCC we run the estimators a practitioner reaches for first. The *static* baseline is the full-sample Pearson correlation (or covariance), a single matrix held fixed for all t . The *rolling* baselines recompute the Pearson correlation (or covariance) over a trailing window of the last w observations, causal by construction, for $w \in \{30, 60, 120, 250\}$. These are exactly the estimators DCC must beat to justify its extra machinery.

4 Experimental design

All experiments are generated and analyzed by one seeded Python harness (`scripts/run_all.py`, estimator in `scripts/dcc.py`) under Python 3.14.6 with NumPy 2.5.1. A verification script (`scripts/check_paper_numbers.py`) checks every numeric claim below against `results/results.json` and fails on any mismatch; quantities are rounded from that file.

Experiment 1: correlation tracking. On the two- and three-asset regime processes we compare the average pairwise correlation implied by DCC, by each rolling window, and by the static estimate against the true path ρ_t , reporting mean absolute error (MAE) and root-mean-square error (RMSE) both overall and restricted to the crisis window. Separately, on the stationary DCC process we fit (a, b) and compare to the known truth.

Table 1: Experiment 1, two-asset process ($T = 2000$): tracking error of each estimator against the true correlation path, overall and inside the crisis window (observations 600 to 1300). DCC has the lowest crisis-window error of all and the lowest overall RMSE; no single rolling window is best both overall and in the crisis. The static estimate is the full-sample correlation 0.673.

Estimator	MAE overall	RMSE overall	MAE crisis	RMSE crisis
Static (full sample)	0.238	0.259	0.208	0.217
Rolling $w = 30$	0.092	0.129	0.037	0.048
Rolling $w = 60$	0.069	0.100	0.031	0.040
Rolling $w = 120$	0.060	0.087	0.037	0.054
Rolling $w = 250$	0.071	0.102	0.074	0.101
DCC-GARCH	0.061	0.083	0.028	0.039

Experiment 2: static-risk misstatement. For an equal-weight portfolio we compute one-step 95% Gaussian Value-at-Risk, $\text{VaR}_t = z_{0.95} \sqrt{w' \Sigma_t w}$ with $z_{0.95} \approx 1.645$, using in turn the static, rolling, DCC, and true covariance for Σ_t , each built from information through $t - 1$. A breach is a realized portfolio loss exceeding VaR_t . We report the breach rate overall, in the crisis window, and in the calm periods; the nominal rate is 0.05.

Experiment 3: dynamic hedge ratio. On the two-asset process we form the minimum-variance hedge ratio of asset one against asset two, $\beta_t = \rho_t \sigma_{1,t} / \sigma_{2,t}$, from DCC outputs, and compare it to a static ordinary-least-squares hedge ratio (the full-sample covariance ratio). Hedging with the lagged β_{t-1} we measure the variance of the hedged spread and its reduction relative to the static hedge, plus the tracking error of the estimated β_t against the true β_t .

Experiment 4: dimensionality. We tabulate the free-parameter count of full VEC, full BEKK, and DCC as d grows, and re-run the parameter-recovery experiment at $d \in \{2, 3, 5, 10\}$ to check that the estimated (a, b) stay stable as the cross-section widens.

5 Results

5.1 Correlation tracking and parameter recovery

Table 1 reports the tracking errors on the two-asset process. The static full-sample correlation is a flat line at 0.673, and it is hopeless against a path that moves between 0.3 and 0.9: its crisis-window MAE is 0.208. DCC drives that down to 0.028, a factor of roughly 7. The comparison against rolling windows is where the honest structure of the result lives. No single fixed window is best everywhere: the window that tracks best overall ($w = 120$, overall MAE 0.060) lags badly through the crisis (crisis MAE 0.037), while the window that tracks the crisis best ($w = 60$, crisis MAE 0.031) is noisier overall. DCC is at or below the best window on every slice at once — crisis MAE 0.028 (below the best window’s 0.031), overall RMSE 0.083 (below every window’s), and overall MAE 0.061 (indistinguishable from the best window’s 0.060). A short window is responsive but noisy; a long window is smooth but stale; DCC’s exponential, adaptively weighted memory is both, which is exactly the tradeoff the two parameters (a, b) are there to resolve.

The three-asset process tells the same story: DCC’s crisis-window MAE is 0.066, below every rolling window (best 0.070 at $w = 60$) and far below the static 0.202, and it fits $(a, b) = (0.031, 0.966)$ — the near-unit persistence and small shock loading that is the canonical DCC fingerprint.

Table 2: Experiment 2: one-step 95% VaR breach rates for an equal-weight portfolio (nominal 0.05), by covariance estimator, overall and split into crisis and calm periods. Static covariance looks calibrated overall (0.049) but breaches at 0.111 — $2.2\times$ nominal — inside the crisis; DCC stays close to the true-covariance floor.

Covariance estimator	Breach overall	Breach crisis	Breach calm
Static (full sample)	0.049	0.111	0.015
Rolling ($w = 60$)	0.062	0.069	0.058
DCC-GARCH	0.056	0.063	0.052
True covariance	0.060	0.060	0.060

Parameter recovery is exact to within sampling noise. On the stationary two-asset DCC process with true $(a, b) = (0.05, 0.90)$, the estimator returns $(0.049, 0.914)$: an absolute error of 0.0015 in a and 0.014 in b . On a three-asset process with true $(0.04, 0.92)$ it returns $(0.040, 0.914)$. The self-contained two-step estimator recovers the data-generating parameters, so the tracking performance above is not an artifact of a mis-specified fit.

5.2 Static covariance under-states portfolio risk in the spike

Table 2 is the practical payoff. For the equal-weight two-asset portfolio, the static full-sample covariance produces a 95% VaR whose overall breach rate is 0.049 — essentially perfect calibration, and exactly the false comfort that makes static covariance dangerous. That average hides a violent split: inside the crisis the static VaR is breached at 0.111, 2.2 times the nominal 0.05, while in the calm periods it breaches at only 0.015. A risk manager watching the full-sample number would see a well-calibrated model while carrying more than double the intended tail risk through every stress episode. The rolling covariance ($w = 60$) improves the crisis breach rate to 0.069 but still runs hot, because a trailing window that straddles the regime boundary is a blend of calm and crisis. DCC holds the crisis breach rate to 0.063, within touching distance of the 0.060 that the *true* covariance itself delivers in finite sample — the irreducible floor. The three-asset portfolio repeats the pattern: static breaches at 0.103 in the crisis versus DCC’s 0.070 and the truth’s 0.060.

5.3 A dynamic hedge ratio

On the two-asset process the static OLS hedge ratio is a constant 0.650, while the true minimum-variance hedge ratio averages 0.576 and moves with the regime. The DCC hedge ratio averages 0.575 and tracks the true path with a mean absolute error of 0.067 overall and 0.035 inside the crisis, against the static hedge’s 0.236 and 0.248. Trading the lagged hedge, the DCC-hedged spread has a fractional variance reduction of 0.120 relative to the static-hedged spread overall, rising to 0.208 inside the crisis — precisely where the static hedge, calibrated to the full-sample average, is most wrong. The dynamic hedge does not create an edge; it keeps the traded spread market-neutral through the regime shift instead of drifting into directional exposure exactly when correlation moves.

5.4 Dimensionality

Table 3 is the structural reason DCC is used at all. Counting free parameters, full VECM carries an intercept plus two coefficient matrices on the half-vectorization of the covariance, growing as $O(d^4)$; full BEKK carries a triangular intercept plus two $d \times d$ coefficient matrices, $O(d^2)$; DCC

Table 3: Experiment 4: free-parameter counts of full VEC, full BEKK, and DCC (three per univariate GARCH(1,1) plus the two correlation scalars) as the number of assets d grows. VEC is $O(d^4)$, BEKK is $O(d^2)$, DCC is $O(d)$ with a correlation layer fixed at two parameters.

d	VEC	BEKK	DCC
2	21	11	8
3	78	24	11
5	465	65	17
10	6105	255	32
20	88410	1010	62
50	3252525	6275	152

carries three parameters per univariate GARCH plus the two correlation scalars, $O(d)$. The gap is not asymptotic hand-waving: at $d = 5$ it is 17 parameters for DCC versus 65 for BEKK and 465 for VEC, and at $d = 50$ it is 152 versus 6275 versus 3.25×10^6 . VEC and BEKK are simply not estimable on the cross-sections DCC handles routinely. The correlation layer itself never grows past two parameters. Estimation stays well-behaved as the cross-section widens: re-running the recovery of true $(a, b) = (0.05, 0.90)$ at $d = 2, 3, 5, 10$ returns a within 0.006 and b within 0.015 of the truth at every width, with no drift as d grows.

6 Discussion

6.1 What the dynamic model buys

The three substantive experiments are three views of one fact: a correlation that moves must be tracked, not averaged. In tracking (Table 1) the static estimate’s error is an order of magnitude worse than DCC’s inside the crisis, and the rolling windows force a choice — responsiveness or smoothness — that DCC does not. In risk (Table 2) the same static average that looks calibrated in aggregate (0.049 against nominal 0.05) under-states the crisis tail by a factor of 2.2, which is the failure mode that matters: a risk model is not judged on its average behavior but on its behavior in the tail, and the static model’s tail arrives concentrated inside exactly the window it cannot see. DCC’s crisis breach rate of 0.063 against the true-covariance floor of 0.060 says the residual error is close to irreducible. In hedging the dynamic ratio’s crisis variance reduction of 0.208 is the same effect measured in a trading decision rather than a diagnostic.

6.2 Why correlation targeting and causal timing matter

Two implementation choices carry the result and are worth isolating. Correlation targeting — plugging $\bar{Q}$ in as a sample estimate rather than optimizing it — is what collapses the correlation layer to two parameters and makes the $d = 50$ column of Table 3 possible; without it the correlation step would itself be an $O(d^2)$ optimization. Its cost, invisible in this study by design, is that $\bar{Q}$ uses the full sample, so a strict walk-forward evaluation must re-estimate it on training data only. Causal timing in the recursion — scoring R_t before advancing Q — is what keeps the tracking numbers honest; the reverse convention would let R_t peek at the contemporaneous shock and flatter every error in Table 1.

6.3 The honest anatomy of the risk result

The risk experiment is engineered to make one point cleanly, and its limitations follow from that engineering. The crisis is a joint volatility and correlation spike, so the static under-statement mixes two effects; we do not decompose how much of the $2.2\times$ breach comes from stale volatility versus stale correlation, only that ignoring the joint dynamics is what fails. The true covariance breaches at 0.060 rather than exactly 0.05 because a Gaussian-quantile VaR on a GARCH-mixture return is mildly miscalibrated in finite sample even when the covariance is known — which is the correct floor to compare against, and DCC’s 0.063 essentially reaches it. These are single seeded processes, not Monte Carlo frequency statements; they demonstrate the mechanism and its magnitude under known truth, not its distribution across markets.

7 Limitations

- **Synthetic data, by design.** Both processes — a prescribed correlation path over GARCH volatilities, and a stationary DCC recursion — are chosen so that ground truth is known exactly, which is the point of a calibration study and also its boundary. Real returns have fat tails, asymmetric correlation responses, and structural breaks the Gaussian quasi-likelihood does not model; the experiments say what DCC recovers when its assumptions hold, not how it behaves when they do not.
- **Scalar DCC dynamics.** One a and one b govern every pair. The asymmetric DCC of Cappiello et al. (2006), which adds a downside-correlation term, is the natural extension for markets where crash-correlation dominates; we did not implement it.
- **Correlation targeting leaks the full sample.** The in-sample $\bar{Q}$ used here is a look-ahead device acceptable in a controlled study but not in a backtest; the tracking and risk numbers are in-sample filtering results, not walk-forward performance.
- **Gaussian quasi-likelihood.** We estimate under a Gaussian kernel; Student- t innovations and the Engle and Sheppard (2001) standard-error correction would be needed for inference on real data. For recovering the point path they are not.
- **Parameter counts are structural, not fitted.** Table 3 counts free parameters of the full VEC and BEKK specifications; the diagonal and scalar restrictions people actually use are smaller, but so is their generality, and none of them recover DCC’s $O(d)$ scaling with a moving full correlation matrix.
- **Single seeded runs.** The tracking, risk, and hedge experiments are worked examples with known truth, not frequency statements across seeds; the recovery experiment’s stability across d is the closest thing to a robustness sweep we report.

8 Conclusion

Under controlled conditions with a known moving correlation, DCC-GARCH does what it promises. It tracks a calm/crisis/recovery correlation path with a crisis-window error roughly a seventh of the static estimate’s and at or below the best fixed rolling window on every slice at once, and it recovers the data-generating (a, b) to within sampling noise. The cost of ignoring the dynamics is concrete and one-sided: a static covariance that looks perfectly calibrated in aggregate under-states

portfolio tail risk by a factor of 2.2 inside the correlation spike, while DCC tracks the true covariance to the finite-sample floor; a static hedge drifts out of neutrality exactly when correlation moves, where the dynamic hedge cuts spread variance by a fifth. And because correlation targeting fixes the correlation layer at two parameters, all of this scales to cross-sections where full multivariate GARCH cannot be estimated at all. None of this is a claim about any market — the processes are synthetic and seeded — but as a statement about the method, the dynamic model recovers what the static one throws away, and the discarded information is most valuable precisely when it is most costly to lack.

Reproducibility. All code, tests, and outputs accompany this paper: `scripts/run_all.py` re-generates `results/results.json` from fixed seeds (Python 3.14.6, NumPy 2.5.1); `scripts/dcc.py` contains the self-contained DCC estimator; `scripts/check_paper_numbers.py` verifies every numeric claim in this manuscript against that file and fails on any mismatch; `tests/` contains deterministic invariant tests, including that the estimator recovers a known (a, b) within tolerance.

References

- Tim Bollerslev. Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327, 1986. doi: 10.1016/0304-4076(86)90063-1.
- Tim Bollerslev. Modelling the coherence in short-run nominal exchange rates: A multivariate generalized ARCH model. *The Review of Economics and Statistics*, 72(3):498–505, 1990. doi: 10.2307/2109358.
- Tim Bollerslev, Robert F. Engle, and Jeffrey M. Wooldridge. A capital asset pricing model with time-varying covariances. *Journal of Political Economy*, 96(1):116–131, 1988. doi: 10.1086/261527.
- Lorenzo Cappiello, Robert F. Engle, and Kevin Sheppard. Asymmetric dynamics in the correlations of global equity and bond returns. *Journal of Financial Econometrics*, 4(4):537–572, 2006. doi: 10.1093/jjfinec/nbl005.
- Robert Engle. Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business & Economic Statistics*, 20(3):339–350, 2002. doi: 10.1198/073500102288618487.
- Robert Engle and Kevin Sheppard. Theoretical and empirical properties of dynamic conditional correlation multivariate GARCH. NBER Working Paper 8554, National Bureau of Economic Research, 2001.
- Robert F. Engle. Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007, 1982. doi: 10.2307/1912773.
- Robert F. Engle and Kenneth F. Kroner. Multivariate simultaneous generalized ARCH. *Econometric Theory*, 11(1):122–150, 1995. doi: 10.1017/S0266466600009063.
- Kevin Sheppard et al. arch: Autoregressive conditional heteroskedasticity (ARCH) and other tools for financial econometrics (Python). Software, Zenodo, 2024. <https://github.com/bashtage/arch>.