Financial Econometrics | Lecture 10

Volatility, Systems and the Long Run: GARCH, VAR and Cointegration

In Class Group Activity

· SJTU, Summer School 2026 · Fatih Kansoy

Lesson of the day. You cannot forecast the direction of a return, but you can forecast its risk, predict several series jointly, and tie two wandering prices into one long-run relationship. Getting the volatility input, the system and the equilibrium right is what separates a working risk desk from a blown-up one.

Phase 1 – Concept checks

ANSWERS

ConcepTest 1. **Answer.** Option **(b)**. Persistence $\alpha + \beta = 0.97$ is the decay rate of a shock: it fades slowly (half-life about 23 trading days), but the model is stable because $\alpha + \beta < 1$ (explosive needs ≥ 1). Option (a) confuses persistence with instability; (c) confuses forecasting the variance with forecasting the return (the recursion never forecasts the sign); (d) confuses the decay rate with the long-run level.

ConcepTest 2. **Answer.** Option **(c)**. VaR $= 2.326\,\sigma_t$ inherits the time variation of the conditional volatility, so freezing it at the calm-day level badly understates risk in a crisis. Option (a) treats VaR as a fixed property; (b) inverts the direction; (d) throws away the conditional model.

In the lab the same 99% formula ran 1.2% on a calm day, 2.99% on the latest day, and 17.7% at the COVID peak. A constant limit is blind exactly when you need it most.

ConcepTest 3. **Answer.** Option **(b)**. The levels regression is valid only if the residual is stationary, which is what cointegration means; here the ADF on the residual is -3.69, beyond the 5% Engle-Granger value -3.34, so they are cointegrated and the long-run relation (10yr $=$ $2.20 + 0.843 \times$ 3mo) is real. Option (a) trusts R^2 and t that are unreliable for I(1) series; (c) over-corrects into the blanket ban; (d) is the "both trending means cointegrated" error: trending together is necessary, not sufficient.

Phase 2 – The desk problem

WORKED SOLUTIONS

(b) Persistence $\alpha + \beta = 0.13 + 0.84 = 0.97$. Long-run variance $\dfrac{\omega}{1-\alpha-\beta} = \dfrac{0.05}{0.03} = 1.667$; long-run volatility $\sqrt{1.667} = 1.29\%$. Half-life $\dfrac{\ln 0.5}{\ln 0.97} = \dfrac{-0.6931}{-0.03046} = 22.8 \approx \mathbf{23}$ trading days. Today's $\sigma_t = 1.72\%$ is *above* the long-run 1.29% and will drift back down over the coming weeks.

(c) VaR $= 2.326\,\sigma_t$. Calm day: $2.326 \times 1.29\% = \mathbf{3.00}\%$ of the book. Turbulent day: $2.326 \times 1.72\% = \mathbf{4.00}\%$. The VaR rises by 1.00 percentage point, about a third, from a single bad day. (In the real S&P lab the same formula ran from 1.2% to 17.7% at the COVID peak, a roughly fifteen-fold jump.)

(d) (1) Correction term $= -0.05 \times z_{t-1} = -0.05 \times 0.80 = \mathbf{-0.040}$ points: the 10-year is pulled *down* by about 0.04 points next period, shrinking the gap. (2) About **5**% of the gap closes per period, which is *slow*: at 5% a period the half-life of a gap is $\ln(0.5)/\ln(0.95) \approx 13.5$ periods. Cointegration says the gap *will* revert; it does not promise speed. (3) The two rates are cointegrated: each is I(1), but the residual $z_t = \text{Rate10} - 1.50 - 0.90\,\text{Rate3}$ is stationary (you would confirm with an Engle-Granger residual test, as in the lab where the ADF was $-3.69 < -3.34$). A stationary residual makes the levels regression valid, recovering a genuine long-run equilibrium. For two *unrelated* random walks the residual would have a unit root, the regression would be spurious, and the high R^2 and big t would mean nothing. The residual's stationarity is the whole difference.

Phase 3 – The case

REVEAL AND ANSWERS

1. $\sigma_t^2 = 0.04 + 0.165(19.4) + 0.805(0.70) = 0.04 + 3.20 + 0.56 = \mathbf{3.80}$, so $\sigma_t = \sqrt{3.80} = 1.95\%$ and the conditional 99% VaR is $2.326 \times 1.95\% \approx \mathbf{4.5}\%$, already nearly four times the frozen 1.2%. The frozen limit is protecting no one: the true one-day risk has already left it far behind.

2. Half-life $= \ln(0.5)/\ln(0.97) \approx \mathbf{23}$ trading days, about a month. The spike is forecast to persist for weeks, not a day; this is not a blip.

3. A constant VaR limit is blind exactly in the crisis it exists to survive, because true risk scales with the conditional volatility, and in late February that volatility was already multiplying while the frozen January limit stood still.

4. The defensible answers are (ii) override and cut the book hard, or (iii) freeze new risk and escalate immediately; (i) "keep trading because the rule says you are compliant" is the wrong answer, because the rule is measuring the wrong thing.

Instructor reveal. The frozen-limit desk was the kind that got carried out in March 2020. Funds that sized risk off calm-period (unconditional) volatility were running far more risk than they thought: their true conditional VaR had multiplied roughly fifteen-fold (in the lab, 1.2% to 17.7%) while their reported limit had not moved at all. Several leveraged and risk-parity-style funds were forced into fire sales at the bottom, because their actual loss potential had exploded while their risk system still reported the January number. The desks that survived best ran *conditional* volatility models that flagged the danger in late February, cut exposure early, and treated the spike as persistent (half-life about a month) rather than a one-off. VaR is only as good as its volatility input, and volatility is not a constant. This is Lecture 8's lesson, that the size of returns is forecastable even when the direction is not, turned into a risk-management failure.

Reserve alternate. A second desk used the conditional GARCH VaR but assumed the spike was a one-off and did not cut. With persistence 0.97 the elevated volatility persisted for weeks, the limit was breached day after day, and the slow mean-reversion meant losses compounded first. Having the conditional model is not enough; you must believe its persistence.

End of worked solutions, and of the course.