Risk Everywhere: Modeling and Managing Volatility

Tim Bollerslev, Benjamin Hood, John Huss, Lasse Heje Pedersen

Discussion by: Liuren Wu

April 29th, 2016

Overview

- Examines volatility behaviors of 50 global macro series over 20 years.
  - Tons of high-frequency data, lots of details
  - Choices existing and enhanced specifications
  - Different estimation approaches
- Key strength lies in balance between
  - academic sophistication v. practical robustness
  - empirical examination v. economic common sense
- Findings will likely become benchmarks for future volatility modeling/forecasting efforts.

Commonality of *risk dynamics*

- Finding: Same *Risk* dynamics *Everywhere*
  - The *behaviors/dynamics* of all volatility series look similar up to a scale.
  - Pooled estimation generates more stable parameter estimates and better *out-of-sample* forecasting performance.
- Cannot agree more: This is an important practical direction to go, even if *in-sample* likelihood deteriorates somewhat.
  - RiskMetrics does similar things: $v_t = .97 v_{t-1} + .03 r_t^2$.
    The same persistence parameter (e.g., $\phi = 0.97$) is applied to all financial series.
- Go further: No estimation necessary: Parameters can be set as *controls*.
  - Set smoothing (ϕ) according to target horizon, very much like fixed-window estimates
    - One-month variance: $\phi = (.5)^{(1/20)} = 0.966$
    - One-year variance: $\phi = (.5)^{(1/251)} = 0.997$

Commonality of *risk dynamics*

- Finding: Same *Risk* dynamics *Everywhere*
  - The *behaviors/dynamics* of all volatility series look similar up to a scale.
  - Pooled estimation generates more stable parameter estimates and better *out-of-sample* forecasting performance.
- Cannot agree more: This is an important practical direction to go, even if *in-sample* likelihood deteriorates somewhat.
  - RiskMetrics does similar things: $v_t = .97 v_{t-1} + .03 r_t^2$.
    The same persistence parameter (e.g., $\phi = 0.97$) is applied to all financial series.
- Go further: No estimation necessary: Parameters can be set as *controls*.
  - Set smoothing (ϕ) according to target horizon, very much like fixed-window estimates
    - One-month variance: $\phi = (.5)^{(1/20)} = 0.966$
    - One-year variance: $\phi = (.5)^{(1/251)} = 0.997$

Multi-frequency dynamics

$$RV_{t+h} - RV_t^{LR} = \sum \beta_j \left(RV_t^j - RV_t^{LR} \right) + e, \quad j = Day, Week, Month$$

- A good starting point with the HAR model, which anchors prediction with *multiple frequency components* (Day, Week, Month...)
- Practical modifications (enhancements):
  - Each forecasting horizon is estimated separately.
  - Remove the mean parameter with a long-run moving average (RV_t^{LR}).

Why are these enhancements? — They make economic sense.

- Financial fluctuations are combinations of different cycles
  - Long-run debt cycles, with a duration of several decades/a century
  - Moderate-run productivity (business) cycles, once about every 5 years
  - Monetary policy cycles within/along each business cycle
  - Supply-demand shocks, daily news, emotions, ...
- Shorter cycles must fluctuate within the confines of slower (longer) cycles.

Multi-frequency dynamics

$$RV_{t+h} - RV_t^{LR} = \sum \beta_j \left(RV_t^j - RV_t^{LR} \right) + e, \quad j = Day, Week, Month$$

- A good starting point with the HAR model, which anchors prediction with *multiple frequency components* (Day, Week, Month...)
- Practical modifications (enhancements):
  - Each forecasting horizon is estimated separately.
  - Remove the mean parameter with a long-run moving average (RV_t^{LR}).

Why are these enhancements? — They make economic sense.

- Financial fluctuations are combinations of different cycles
  - Long-run debt cycles, with a duration of several decades/a century
  - Moderate-run productivity (business) cycles, once about every 5 years
  - Monetary policy cycles within/along each business cycle
  - Supply-demand shocks, daily news, emotions, ...
- Shorter cycles must fluctuate within the confines of slower (longer) cycles.

A cascade structure for multi-frequency dynamics

a cascade of infinite dimensions

$$
\begin{array}{rcll}
\ldots & = & \ldots & \text{super high-frequency pattern} \\
dv_t^j & = & \kappa_j(v_t^{j+1} - v_t^j)dt + \sigma_j dW_t, & j = S, H, D, W, M, Q, Y, C \\
\ldots & = & \ldots & \text{super long/slow cycles}
\end{array}
$$

- κ_j controls the frequency/duration of the cycle, $1/\kappa$ is in years.
- Roughly power law (geometric) scaling across cycles: e.g., $\kappa_j = 2\kappa_{j+1}$
  - One can either estimate κ or just set them as controls:
    ... 1/252 (day), 1/52 (week), 1/12 (month), 1/4 (quarter), 1 (year), 10 (decade), 100 (century), ...
- There is no long-run mean, but just increasingly longer cycles

Calvet, Fisher, Wu — use the cascade structure to model interest rate dynamics and term structure, but it is equally applicable to variance dynamics

Multi-frequency dynamics in practice

$$dv_t^j = \kappa_j(v_t^{j+1} - v_t^j)dt + \sigma_j dW_t, \quad j = S, H, D, W, M, Q, Y, C$$

- Practical predication is always *a local approximation* of the cascade, at the locale of your particular interest
- Separate estimation (and better yet, separate specification) for each horizon makes perfect sense, with *each focusing on a different block of the cascade:*
  - Minute-by-minute forecast focuses on intraday patterns — $RV^{LR} =$ 1-month average.
  - Daily variance prediction — $RV^{LR}=$ 1-year average.
  - Annual variance prediction — $RV^{LR} =$ 10-year average.
- The modifications are economically sensible enhancements:
  - Separate estimation for different target horizons — Yes, and maybe also use different frequency components.
  - No long-run mean — Use the next, slower cycle as the local "center."
  - Common risk dynamics — Same target horizon should focus on the same block of the cascade, with $1/\kappa$=horizon.

Go further

- Go beyond common dynamics: *Same risk everywhere*

  Different financial series also show strong co-movements.

  - Cross-sectional averaging of scaled volatility levels, at least within classes, can directly improve out-of-sample forecasts.
    - The most persistent component tends to be the common component.

- Go beyond autoregression: *Conditional variance* for scenario analysis/stress tests — What happens to the portfolio if

  - Fed raises rates
  - Stock market crashes (e.g., down 30% in one day)
  - Market experiences a protracted recession (e.g., down 30% in one year)

  Examples:

  - Dupire's local volatility $\sigma(t, S)$
  - Carr and Wu (2016): Option Realized Volatility (ORV) — not only time-weighting but also dollar gamma weighting