Discussion of *Time-Varying Volatility and the Power Law Distribution of Stock Returns*

Paper by Missaka Warusawitharana

Guillaume Roussellet[1]

[1] *McGill University*

FMA Conference

October 2017

$$\boxed{R_{t+1} = \mu + \sigma_t \varepsilon_{t+1}} \tag{1}$$

- This is the semi-strong representation of any returns, so ε_{t+1} is a martingale difference and σ_t^2 is the best variance forecast as of t.

$$\mathbb{P}(|R_{t+1} - \mu| > x) \propto x^{-\alpha} \iff \mathbb{P}(\sigma_t|\varepsilon_{t+1}| > x) \propto x^{-\alpha} \tag{2}$$

- Returns have fat tails which decay as a power function (power law).
- This can be linked to either σ_t or ε_{t+1}

$$\boxed{R_{t+1} = \mu + \sigma_t \varepsilon_{t+1}} \tag{1}$$

- This is the semi-strong representation of any returns, so ε_{t+1} is a martingale difference and σ_t^2 is the best variance forecast as of t.

$$\mathbb{P}(|R_{t+1} - \mu| > x) \propto x^{-\alpha} \Longleftrightarrow \mathbb{P}(\sigma_t |\varepsilon_{t+1}| > x) \propto x^{-\alpha} \tag{2}$$

- Returns have fat tails which decay as a power function (power law).
- This can be linked to either σ_t or ε_{t+1}

Question

- Which component drives the power-law **marginal** behavior of returns?

- Central empirical tool: linear regression

$$\log[\,\widehat{\mathbb{P}}(\sigma_t|\varepsilon_{t+1}| > x)\,] = \alpha_0 + \alpha \cdot \log(x) + \eta$$

- can we reproduce the α's we get on data with conditionally Gaussian models? ▸ histogram

The answer is **yes**

- Cute stats example: i.i.d. $\mathcal{N}(0, \sigma_t^2)$ with i.i.d. $\sigma_t^2 \sim \mathrm{I}\Gamma(\alpha, \beta)$

Complete procedure

- Build *model-free* intra-day variance series of DJI stocks (TSRV on HF data)
- Fit a model on the series **or** draw samples of them along with independent Gaussian draws of ε_t
- Compare the estimates of the tail index either closed-form **or** regression on simulated data
- Show that the results of data and models are non-significantly different.
- perform a bunch of robustness tests (block-bootstrap, TSRV at different window lengths)

The answer is **yes**

- Cute stats example: i.i.d. $\mathcal{N}(0, \sigma_t^2)$ with i.i.d. $\sigma_t^2 \sim \mathrm{I}\Gamma(\alpha, \beta)$

Complete procedure

- Build *model-free* intra-day variance series of DJI stocks (TSRV on HF data)
- Fit a model on the series **or** draw samples of them along with independent Gaussian draws of ε_t
- Compare the estimates of the tail index either closed-form **or** regression on simulated data
- Show that the results of data and models are non-significantly different.
- perform a bunch of robustness tests (block-bootstrap, TSRV at different window lengths)

My interpretation of the results

You can stop trying fancy distributions for standardized shocks as long as your volatility model is correctly specified.

- Paper is very clean and well-written
- The procedures are carried out carefully
- I learnt a lot from reading it!

Comments

#1 Enhancing the testing procedures

#2 What do I learn for my standard vol models?

My interpretation of the results

You can stop trying fancy distributions for standardized shocks as long as your volatility model is correctly specified.

- Paper is very clean and well-written
- The procedures are carried out carefully
- I learnt a lot from reading it!

Comments

#1 Enhancing the testing procedures

#2 What do I learn for my standard vol models?

Comment #1: Testing for results

- What you are actually searching is whether conditional normality of ε_t is in line with the tail behavior of returns.
- You are testing the α coefficients from the ECDF regression but you never look at the usual residual tests.

Tests I would like to see

- ACF of std residuals and their squares
- Normality tests (JB, SW, KS, etc)
- Testing independence through time (higher powers of residuals ACF)

Comment #1: Testing for results

- What you are actually searching is whether conditional normality of ε_t is in line with the tail behavior of returns.
- You are testing the α coefficients from the ECDF regression but you never look at the usual residual tests.

Tests I would like to see

- ACF of std residuals and their squares
- Normality tests (JB, SW, KS, etc)
- Testing independence through time (higher powers of residuals ACF)

Comment #1: Testing for results (cont'd)

- The main method for non-parametric variance modelling is the **two-scale realized variance** of Ait-Sahalia *et. al.*
- It is measured on data, so you have confidence bands around your series
- Paper of Mykland & Zhang (*forth E'ca*) treats this problem
- A simulation (on toy example) shows that they can be large.
- Can you correct standard deviations in the regressions to account for the error in variable problem?

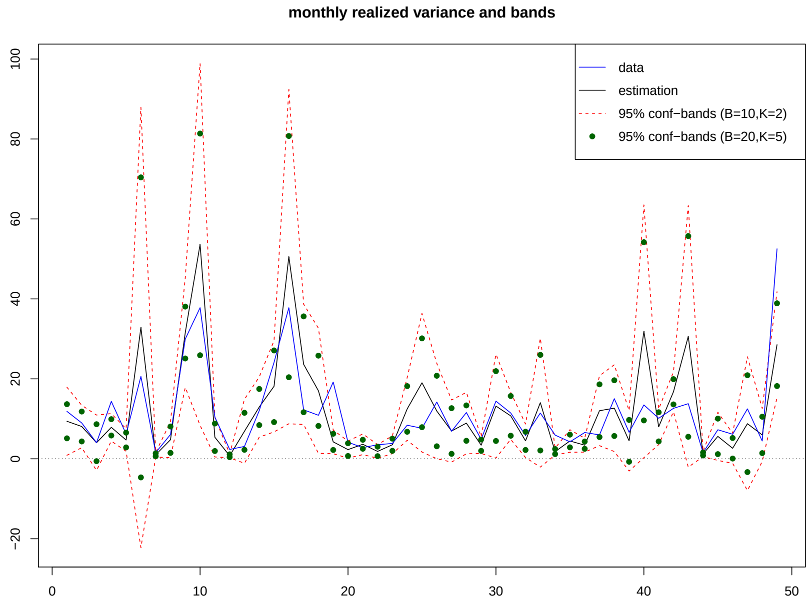
monthly realized variance and bands
data
estimation
95% conf−bands (B=10,K=2)
95% conf−bands (B=20,K=5)
−20
0
20
40
60
80
100
0
10
20
30
40
50

Comment #1: Testing for results (last)

- What is the robustness of the α measurement?
- Are both tails to be considered equally?

Following slide: Measure of α for SPX, starting from different thresholds (.9, .95, .99) and differencing left and right tails.

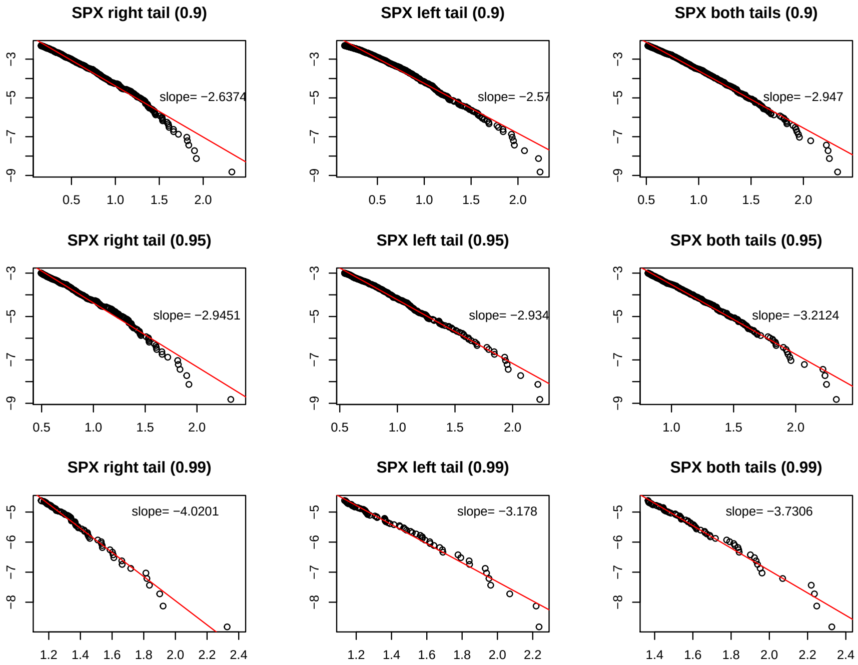
SPX right tail (0.9)
slope= −2.6374
SPX left tail (0.9)
slope= −2.57
SPX both tails (0.9)
slope= −2.947
SPX right tail (0.95)
slope= −2.9451
SPX left tail (0.95)
slope= −2.934
SPX both tails (0.95)
slope= −3.2124
SPX right tail (0.99)
slope= −4.0201
SPX left tail (0.99)
slope= −3.178
SPX both tails (0.99)
slope= −3.7306

- The typical exercise is to build volatility models at the daily frequency (GARCH models for instance) and forecast/compute VaR.
- Some people use RV and model its time-dependence at daily freq.

QUESTIONS

- What is the use of 15min returns here? Can't you do the same at daily? Needs some justification
- Asymmetries exist at the daily freq (AGARCH, GJR) but disappear at high freq. Can your model explain skewness in returns as well?
- How do you treat overnight returns? They might have a **very large** impact (intra-day vol seasonality)

Comment #2: What do I learn for volatility modelling (cont'd)

- This raises the problem of conditional vol modelling w.r.t. time aggregation.
- Is the fact that shocks to returns at 15min freq are conditionally Gaussian informative for my volatility model at the daily freq?

I did my own experiment: Estimate a GARCH(1,1) on SPX 1990-2016.

	ω	α	β
coef	0.0111	0.0754	0.9150

Comment #2: What do I learn for volatility modelling (cont'd)

- This raises the problem of conditional vol modelling w.r.t. time aggregation.
- Is the fact that shocks to returns at 15min freq are conditionally Gaussian informative for my volatility model at the daily freq?

I did my own experiment: Estimate a GARCH(1,1) on SPX 1990-2016.

	ω	α	β
coef	0.0111	0.0754	0.9150

Comment #2: What do I learn for volatility modelling (last)

- Normality is rejected because of skew
- ACF of squared residuals is OK nonetheless

Cute(er) theoretical result (see Refs)

Tail index for marginal returns equals 2κ, where κ is the only (strictly) positive solution of:

$$\mathbb{E}\left[\left(\alpha\varepsilon_t^2 + \beta\right)^\kappa\right] = 1$$

- $\varepsilon_t \sim \mathcal{N}(0,1) \Longrightarrow 2\kappa \simeq 5.06$
- $\varepsilon_t \sim \mathcal{T}(df = 4) \Longrightarrow 2\kappa \simeq 2.88$
- $\varepsilon_t \sim \mathcal{S} - Laplace(1.24, 1) \Longrightarrow 2\kappa \simeq 3.098$

Do your results invalidate this model? I can't tell.

Comment #2: What do I learn for volatility modelling (last)

- Normality is rejected because of skew
- ACF of squared residuals is OK nonetheless

Cute(er) theoretical result (see Refs)

Tail index for marginal returns equals 2κ, where κ is the only (strictly) positive solution of:

$$\mathbb{E}\left[\left(\alpha\varepsilon_t^2 + \beta\right)^\kappa\right] = 1$$

- $\varepsilon_t \sim \mathcal{N}(0,1) \Longrightarrow 2\kappa \simeq 5.06$
- $\varepsilon_t \sim \mathcal{T}(df = 4) \Longrightarrow 2\kappa \simeq 2.88$
- $\varepsilon_t \sim \mathcal{S} - Laplace(1.24, 1) \Longrightarrow 2\kappa \simeq 3.098$

Do your results invalidate this model? I can't tell.

Comment #2: What do I learn for volatility modelling (last)

- Normality is rejected because of skew
- ACF of squared residuals is OK nonetheless

Cute(er) theoretical result (see Refs)

Tail index for marginal returns equals 2κ, where κ is the only (strictly) positive solution of:

$$\mathbb{E}\left[\left(\alpha\varepsilon_t^2 + \beta\right)^{\kappa}\right] = 1$$

- $\varepsilon_t \sim \mathcal{N}(0,1) \Longrightarrow 2\kappa \simeq 5.06$
- $\varepsilon_t \sim \mathcal{T}(df = 4) \Longrightarrow 2\kappa \simeq 2.88$
- $\varepsilon_t \sim \mathcal{S} - Laplace(1.24, 1) \Longrightarrow 2\kappa \simeq 3.098$

Do your results invalidate this model? I can't tell.

- A nice and informative paper!
- More thorough statistical tests and more information about the usefulness of HF data would be even nicer
- Put it in perspective of volatility modelling in general: how does this change my way of thinking about my standard GARCH/SV models

Figure 2: Distribution of power law coefficients

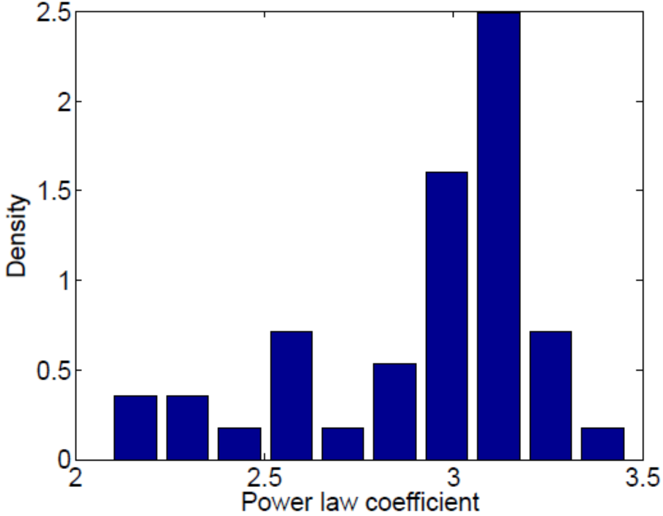

Tail properties of Garch Processes

- Mikosch, T. and Starica, C. (2000) *Limit theory for the sample autocorrelations and extremes of a GARCH(1,1) process*. Ann. Statist. 28, 1427-1451
- Basraka, B. and Davis, R. and Mikosch, T. (2002), *Regular variation of GARCH processes*, Sto. Proc. & App., 99, 95-115
- Liu, J-C. (2006) *On the Tail Behaviors of a Family of GARCH Processes, Econometric Theory*, 22-5, 852-862

Observed AVAR

- Mykland, P. and Zhang, L. (Forth) *Assessment of Uncertainty in High Frequency Data: The Observed Asymptotic Variance*, Econometrica