Filtered Dispersion

CP 25: Risk Models for Securities and Securities Options
SIAM FM19 Toronto

John Dodson

June 7, 2019

Motivation

- Conditional heteroskedasticity in returns is a common stylized fact of timeseries for prices and other financial risk factors.
- A common approach to dealing with this is to extract standard white noise through the application of a model in the generalized auto-regressive framework such as GARCH fit using quasi-maximum likelihood.
- This can fail to be robust, both in fitting and in simulation, because it assumes that the sample entropy is solely determined by the conditional variance.

Contribution

We propose an alternate model for extracting standard white noise from financial timeseries. There are two separate aspects of this proposal:

1. a method for updating the model parameters from a stream of observations, and
2. a particular version of the skewed Student's-t for the residual which has useful properties.

The framework we propose yields models which are strictly distinct from GARCH but are contained within the larger class of models from the GAS (Generalized Autoregressive Score) framework of Creal, Koopman, & Lucas.

Digression: Updating by Perturbation

Say we have an i.i.d. N-sample $\{x_1, x_2, \ldots, x_N\}$ of a random variable, X, whose distribution is known except for the variance $h > 0$, which we assume is finite

- Since we know the mean, wlog assume $\mathrm{E}[X] = 0$

The probability density of the (univariate) X is

$$f_X(x; h) = f_Z\left(\frac{x}{\sqrt{h}}\right)\frac{1}{\sqrt{h}}$$

for some known random variable, Z, with zero mean and unit variance.

Assuming $f_Z(\cdot)$ is smooth, the maximum likelihood estimate, $\hat{h}$, is a root of $\mathcal{L}'(\cdot)$ where

$$\mathcal{L}(h) = -\frac{N}{2}\log h + \sum_{i=1}^{N}\log f_Z\left(\frac{x_i}{\sqrt{h}}\right)$$

Updating by Perturbation

Say we already have an *a priori* estimate of the variance $\hat{h}^-$ and we want to update it with this new sample. Denote the updated estimate by $\hat{h}$. If we can consider this update as a perturbation, the MLE can be approximated by Newton's method, which is based on the series expansion

$$0 \approx \mathcal{L}'(\hat{h}^-) + \mathcal{L}''(\hat{h}^-)(\hat{h} - \hat{h}^-)$$

for $\mathcal{L}''(\hat{h}^-) \neq 0$. Unfortunately this condition is not guaranteed for all h in the domain and x in the support, $\mathcal{X}$. But it <u>is</u> guaranteed for the expected value:

$$\begin{aligned} \mathrm{E}\left[\mathcal{L}''(h;X)\right] &= \int_{\mathcal{X}} \left(\frac{\partial_h^2 f_X(x)}{f_X(x)} - \frac{(\partial_h f_X(x))^2}{f_X(x)^2} \right) f_X(x)\, dx \\ &= -\int_{\mathcal{X}} \frac{(\partial_h f_X(x))^2}{f_X(x)}\, dx = -\operatorname{var}\left[\mathcal{L}'(h;X)\right] < 0 \end{aligned}$$

Updating by Perturbation

So if $\mathcal{L}''(h) \approx \mathrm{E}\left[\mathcal{L}''(h;X)\right]$, a robust updated estimate is

$$\hat{h} = \hat{h}^{-} + \frac{\mathcal{L}'(\hat{h}^{-};\{x_1,\ldots,x_N\})}{N \operatorname{var}\left[\mathcal{L}'(\hat{h}^{-};X)\right]}$$

and since

$$\mathcal{L}'(h;\{x_1,\ldots,x_N\}) = -\frac{1}{2h}\sum_{i=1}^{N}\left(1 + \frac{\left(x_i/\sqrt{h}\right) f_Z'\left(x_i/\sqrt{h}\right)}{f_Z\left(x_i/\sqrt{h}\right)}\right)$$

and

$$\operatorname{var}\left[\mathcal{L}'(h;X)\right] = \frac{1}{4h^2}\left(\int_{\mathcal{Z}} \frac{z^2 f_Z'(z)^2}{f_Z(z)}\,dz - 1\right) \triangleq \frac{\kappa_Z - 1}{4h^2}$$

for a particular κ_Z that does not depend on the variance.

Updating by Perturbation

Our perturbed estimate is

$$\hat{h} = \left(1 - \frac{2}{\kappa_Z - 1}\right)\hat{h}^- + \frac{2}{\kappa_Z - 1}\frac{1}{N}\sum_{i=1}^{N}\frac{-f_Z'\left(\frac{x_i}{\sqrt{\hat{h}^-}}\right)}{\frac{x_i}{\sqrt{\hat{h}^-}}f_Z\left(\frac{x_i}{\sqrt{\hat{h}^-}}\right)}x_i^2$$

- Note that for $Z \sim \mathcal{N}(0,1)$, $f_Z'(z) = -z\, f_Z(z)$ so $\kappa_Z = 3$ and the estimator reverts to the usual MLE for the variance of a normal sample.

Now let's adapt this for the problem at hand, estimating a timeseries of conditional variances $\hat{h}_k^-$ from a timeseries of independent but heteroskedastic residuals ε_k.

- In particular, rather than having an N-sample, we have to content ourselves with a sequence of 1-samples(!)

Digression: Kalman Filter

In the Kalman filter framework, we are interesting in estimating the value of a latent dynamic "state variable" x from a regular discrete "measurement" z and a "control" u,

$$x_{k+1} = Ax_k + Bu_k + \sqrt{S}w_k$$
$$z_k = Hx_k + \sqrt{R}v_k$$

where w_k and v_k are each independent standard (mean zero, unit variance) random variables driving the process and measurement noise, and S and R are positive-definite.

We can never know the value of x_k. We can only estimate it. In the Kalman framework, we estimate it twice for each k: once at time $k-1$ (*a priori*), which we denote $\hat{x}_k^-$; and once again at time k (*a posteriori*) when we have the newest measurement z_k, which we denote $\hat{x}_k$.

Kalman Filter

The minimum-loss (for normal errors) one-step prediction of the state at time k is

$$\hat{x}_{k+1}^{-} = A\hat{x}_k + Bu_k$$

in terms of the *a posteriori* estimate

$$\hat{x}_k = \hat{x}_k^{-} + K_k\left(z_k - H\hat{x}_k^{-}\right)$$

where the "measurement innovation gain" K_k can be determined analytically through solving an optimization.

In the context of a regularly sampled stationary heteroskedastic timeseries, $(Y_{k\Delta t})_k$, we can apply the Kalman filter framework to predict the conditional variance $h_{k+1} \triangleq \mathrm{var}_{k\Delta t}\left[Y_{(k+1)\Delta t}\right]$, based on observations of the "residual" $\varepsilon_{k+1} \triangleq Y_{(k+1)\Delta t} - \mathrm{E}_{k\Delta t}\left[Y_{(k+1)\Delta t}\right]$.

Filtered Variance

Introducing the perturbation update and setting $u_k \equiv 1$ (required for stationarity) we get

$$\hat{h}_{k+1}^{-} = A\left(1 - \frac{2}{\kappa_Z - 1}\right)\hat{h}_k^{-} + A\frac{2}{\kappa_Z - 1}\frac{-f_Z'\left(\hat{z}_k\right)}{\hat{z}_k f_Z\left(\hat{z}_k\right)}\varepsilon_k^2 + B$$

where $\hat{z}_k = \frac{\varepsilon_k}{\sqrt{\hat{h}_k^{-}}}$.

And since

$$\int_{\mathcal{Z}} \frac{-z\, f_Z'(z)}{f_Z(z)} f_Z(z)\, dz = 1$$

we can evaluate the unconditional expectation on both sides to get $B = \sigma^2(1 - A)$ in terms of the unconditional one-period variance, so the recursion to update the *a priori* conditional variance prediction is

$$\boxed{\hat{h}_{k+1}^{-} = \sigma^2 + A\left(\left(1 + \frac{2}{\kappa_Z - 1}\left(\frac{-\hat{z}_k f_Z'\left(\hat{z}_k\right)}{f_Z\left(\hat{z}_k\right)} - 1\right)\right)h_k^{-} - \sigma^2\right)}$$

Filtered Variance

In GARCH, the response of the conditional variance forecast is proportional to the squared residual. In this framework, the response is more general. The quadratic response can be recovered by assuming $Z \sim \mathcal{N}(0,1)$. But financial returns are generally leptokurtotic.

Skewed Student's-t

An alternate specification for financial returns that has received some recent attention is the Generalized Hyperbolic version of skewed Student's-t, $Z|Q \sim \mathcal{N}\left(\beta Q - \beta \operatorname{E}[Q], Q\right)$ where Q is a reciprocal Gamma r.v. with density

$$f_Q(q) = \frac{\chi(\beta,\nu)^{\nu/2}}{2^{\nu/2}\Gamma\left(\nu/2\right)} q^{-\nu/2-1} e^{-\frac{\chi(\beta,\nu)}{2q}}$$

for $\nu > 4$.

Skewed Student's-t

Robustness

In contrast to the Gaussian case, where the response to a residual is quadratic, we can prove that in the GH skewed Student's-t case

$$\frac{-z\, f'_Z(z)}{f_Z(z)} \approx 1 + \frac{\nu}{2} + |\beta z| - \beta z$$

for $|z| \gg 0$.

- the contribution of a large measurement to the sum is asymptotically linear or constant depending on sign.

Log-Returns

Another useful feature of the GH skewed Student's-t is that, unlike the symmetric version, it has a finite moment generating function, $\mathrm{E}\left[e^{hZ}\right]$, for $\beta < 0$ and $0 \leq h \leq -2\beta$.

- So we can model log-returns!

Filtered Variance

Residual Responses

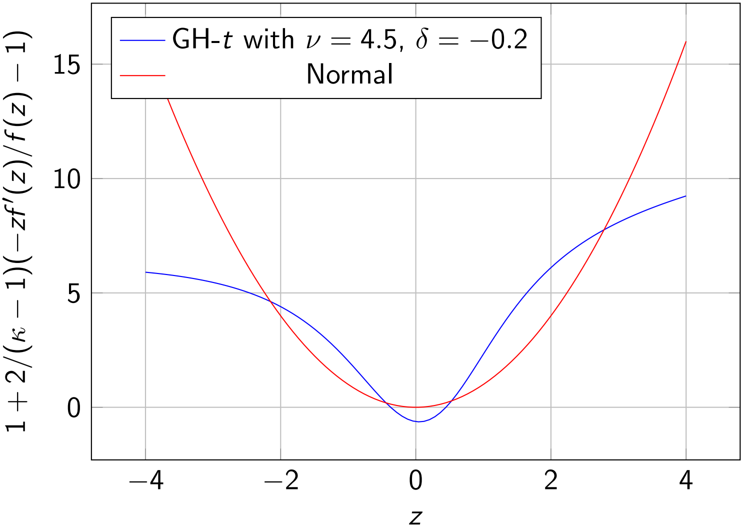

Filtered Variance

normalized Q-Q plot, ten years daily