Quantitative Risk Management
Fall Assignment

John Dodson

October 8, 2015

This assignment is not a regular homework. It is worth half of the module grade for the fall term. Please share your solution with me through Google Drive before the beginning of the classroom session on Wednesday, October 21.

Your solution should include source code for programs and instructions for how to run them with the software available on Math Department computers, for example 'shelby' in VinH 314. I recommend you use `tar` or `zip` if you will be submitting several files.

I expect you to work alone and cite your sources. Please remind yourself of the University's definition of scholastic dishonesty.

Problems

This problem will be based on the timeseries of daily log-returns (adjusted for splits and dividends) of the S&P 500 index, ticker ^GSPC, for two years up to October 14, 2015.

An important variant of GARCH for equity factors is the "leveraged" model of Glosten, Jagannathan, and Runkle (at the University of Minnesota!) in 1993. We can parameterize this as

$$\sigma_t^2 = \alpha_0 + \alpha_1 \left(\varepsilon_{t-1} + \delta \left|\varepsilon_{t-1}\right|\right)^2 + \beta_1 \sigma_{t-1}^2$$

for $|\delta| < 1$. Existence of an unconditional variance requires $\alpha_1 \left(1 + \delta^2\right) + \beta_1 < 1$.

The point of this form is that there are two versions of the α_1 ARCH term: one in reaction to ups and one in reaction to downs. $\delta < 0$ expressed the leverage phenomenon.

1. Estimate the values of the parameters of a leveraged GARCH(1,1) process with NRIG residuals for the log-returns of the index. **(30 points)**
2. Simulate the level of the index at the close of October 21, 2015. **(30 points)**
3. Estimate the 99% value-at-risk and expected shortfall on a long position in the index using the simulations as an empirical loss distribution. **(10 points)**

N.B.: I will provide some assistance in the form of pseudo-code and a MATLAB M-file function for simulating NRIG variates.

Grading Rubric

Thirty out of one hundred points will be based on the follow criteria:

- You follow the instructions. **(5 points)**
- I can reproduce your results with the code and documentation you provide. **(10 points)**
- You include adequate citations. **(10 points)**
- Your write-up is clear and professional. **(5 points)**

Solution

The S&P Index had daily returns less than -2% on eleven days and positive returns greater than 2% on five days in the two years to October 14, 2015 (505 daily log-returns). This would seem to suggest that negative residuals are scaled somewhat more than positive residuals, and we see that in the model fit in Table 1.

$\hat{\alpha}_0$	0.433×10^{-5}
$\hat{\beta}_1$	0.738
$\hat{\alpha}_1$	0.110
$\hat{\delta}$	-1.00
$\hat{g}$	4.23

Table 1: MLE model parameter fits. The log-likelihood was -3.53 nats per observation.

As we have seen before, the intercept (α_0) is essentially zero and the GARCH terms (β_1) dominate the ARCH terms. We also see somewhat fat tails in the residuals, with the fitted NRIG having an excess kurtosis of about $+0.68$ (relative to the Gaussian). Interestingly, we also see that the asymmetry term δ is floored. This means that the ARCH contribution for a positive innovations is effectively zero, while the ARCH contribution for negative innovations is a sizable 0.438 $(4\,\alpha_1)$.

It is worth noting that the market declined by about 12% over six consecutive sessions in late August, followed immediately by the largest positive return in the period. This sequence seems to have heavily influenced the fit.

Let T denote October 14, 2015 in units of trading days. The MLE fit gives

forecast $\hat{\sigma}^2_{T+1}$	4.22×10^{-5}
unconditional $\hat{\sigma}^2$	10.1×10^{-5}

so it seems that volatility is relatively low at the end of the period but is expected to rise.

One final observation about the fit is that the sample variance of the log-returns is 6.76×10^{-5}, which is somewhat different from the unconditional variance estimate based on the MLE here (without variance targeting), so variance targeting may give somewhat different results.

We are interested in the distribution of the index five sessions after the last observation. Since our model is not based on i.i.d. innovations, we need to use Monte Carlo simulation. I implemented this in the following MATLAB script.

```
%% simulate
n=1E6; % simulation size
cum=zeros(n,1); % allocation for cumulative log-returns
epsi=nan(n,1); % allocation for simulated residual (over-written)
sigi=repmat(sqrt(foreMLE),n,1); % alloc. and init. for vols (over-written)
for i=1:5
    epsi=sigi.*rand_nrig(n,thetaMLE(5));
    sigi=sqrt(thetaMLE(1)+thetaMLE(3)*sigi.^2 ...
            +thetaMLE(2)*(epsi+thetaMLE(4)*abs(epsi)).^2);
    cum=cum+epsi;
end
mult=1-sort(exp(cum)); % empirical quantiles for simple return loss
VaR=mult(0.01*n); % 99% value-at-risk
ES=mean(mult(1:0.01*n)); % 99% expected shortfall
```

It is notable that the GARCH simulation has a significantly higher leptokurtosis than the fitted residuals. The sample excess kurtosis I got from this simulation of cumulative simple return was about $+3.4$. The skewness was also significant at 0.86 (skewed toward larger losses than gains). So our risk measure estimates will be much higher than one would get assuming normal returns.

The results I got are in Table 2.

value-at-risk	$0.0448 \times 1994.24 \approx 89$ points
expected shortfall	$0.0590 \times 1994.24 \approx 117$ points

Table 2: Five-day 99%-confidence risk metrics on October 14, 2015.

One final note: on the five trading days following August 18, 2015, the index lost 229 points. According to our empirical distribution, the (current) probability of five-day losses at least this great is about 0.01%. Presumably it would have been even less probable based on data strictly prior to this episode.