Quantitative Risk Management
Fall Assignment

John Dodson

October 12, 2016

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 19.

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 2,500 business days (about ten years) through October 7, 2016.

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 modification is that there are two versions of the α_1 ARCH term: one in reaction to ups and one in reaction to downs. $\delta < 0$ expresses the leverage phenomenon, where declines in asset values provoke more volatility then increases do.

1. Estimate the values of the parameters of a GJR-GARCH(1,1) process for the log-returns of the index using quasi-maximum likelihood. **(30 points)**

2. Use the fit above to standardize the historical residuals and use maximum likelihood to fit a generalized Pareto to the 2% left tail. **(20 points)**

3. Based on the conditional variance forecast and the characterization of the residuals, calculate the 99% one-day value-at-risk (in index points) for a long position. **(30 points)**

Grading Rubric

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

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

Solution

The S&P Index had daily returns less than -5% on fourteen days and positive returns greater than +5% on nine days between November 2, 2006 and October 7, 2016. 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.

$\hat{\alpha}_0$	3.21×10^{-6}
$\hat{\beta}_1$	0.871
$\hat{\alpha}_1$	0.0504
$\hat{\delta}$	-1.000

Table 1: QMLE model parameter fits with a log-likelihood of 7998.6

It is interesting to compare this result with simpler models. If we restrict the fit to a constant volatility ($\alpha_1 = 0$), the maximum log-likelihood is about 2.91 nats per day. For regular GARCH(1,1) ($\delta = 0$), this increases to 3.17 nats per day. Finally, letting δ float, we get 3.20 nats per day. At all stages, the increase in model complexity is justified by the increase in the fit quality.

As we have seen before, the intercept (α_0) is essentially zero and the GARCH terms (β_1) dominate the ARCH terms. 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 0.202 ($4\,\alpha_1$).

Let T denote October 7, 2016 in units of trading days. The QMLE fit gives

forecast σ^2_{T+1}	57.6×10^{-6}
unconditional σ^2	$114. \times 10^{-6}$

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 175.1×10^{-6}, which is somewhat different from the unconditional variance estimate based on the MLE here (without variance targeting), so variance targeting may give different results.

Let's check the standardized residuals for clustering. In the table below are the dates of the 25 worst daily returns over the period. It does seem by eye that there are more extreme losses in the earlier part of the period than the later part of the period, but the distribution of the intervals is not too far from the expected exponential distribution. In fact, the mean of the interval is 100.1 and the standard deviation is 92.3, which are not far from the theoretical values of 100 business days for each under the independence assumption. Without standardizing, the mean length of the interval between 1% exceedances is 30.7 and the standard deviation is 117.1, both of which would tend to reject the independence hypothesis. In fact, without standardizing, there is not a single loss exceedance at the 1% level between April 20, 2009 and August 4, 2011.

27-Feb-2007	10-Feb-2009	27-Apr-2010	1-Jun-2011	31-Jul-2014
19-Oct-2007	20-Apr-2009	16-Jul-2010	4-Aug-2011	29-Jun-2015
1-Nov-2007	17-Aug-2009	11-Aug-2010	8-Aug-2011	20-Aug-2015
6-Jun-2008	1-Oct-2009	28-Jan-2011	7-Nov-2012	24-Jun-2016
29-Sep-2008	16-Apr-2010	22-Feb-2011	15-Apr-2013	9-Sep-2016

Table 2: The dates of the twenty-five largest standardized losses in the historical period.

Let's move on to fitting a generalized Pareto to the 2% left tail of the standardized daily returns. Using the methods from previous assignments and the demonstration, I got the results in the table below.

θ	0.02
$\hat{\eta}$	-2.493
$\hat{\beta}$	0.322
$\hat{\xi}$	0.412

Table 3: MLE parameters for the 2% left tail of the standardized daily returns with a log-likelihood of -209.6.

Evaluating the quantile function,

$$F^{-1}(q) = \eta - \frac{\beta}{\xi}\left(\left(\frac{\theta}{q}\right)^{\xi} - 1\right) \text{ for } q \leq \theta$$

at $q = 0.01$ gives a 1% threshold loss of about -2.75 standard deviations.

This corresponds to a long value-at-risk of about

$$2153.74 \times \left(e^{-2.75\times\sqrt{57.6\times10^{-6}}} - 1\right)$$

or about **-44.5 index points**, or just over a 2% loss for Monday, October 10, 2016.

Compare this result to a simplistic version, based on the sample variance and the normal distribution. With $\Phi^{-1}(0.01) \approx -2.33$, we get a long 1% value-at-risk of about -65.3 index points.