Risk & Asset Allocation
Homework for Week 6

John A. Dodson

October 12, 2011

Problem

A solution to this problem is due at the beginning of the next session, which is 5:30 PM on Wednesday, October 19.

- Estimate the parameters of the NGARCH(1,1) model to continuous daily total returns on GOOG common equity shares over the two years through September 2011 using variance targeting. **(10 points)**

Solution

The model is $\epsilon_i = Y_i - m_i$ where Y is a timeseries of total returns, m is the conditional expectation, and h is the conditional variance where

$$h_i = \omega + \alpha \left(\epsilon_{i-1} + \gamma \sqrt{h_{i-1}} \right)^2 + \beta h_{i-1}$$

Below is the M-file script that I used to solve this problem, but first we have to define a function to return an array of conditional variances given an array of unexpected innovations and candidate values for the parameters.

```
function [h omega forecast]=NGARCH(epsilon,params)
% NGARCH(1,1) conditional variances
% params = [alpha beta gamma]
% assume that epsilon(1) is the oldest residual
sigma_sqr=mean(epsilon.^2);
omega=sigma_sqr*(1-params(1)*(1+params(3)^2)-params(2)); % variance targeting
epsilon=[epsilon;NaN]; % add on the next observation for forecasting
h=nan(size(epsilon));
h(1)=omega/(1-params(1)*params(3)^2-params(2)); % seed 'h(1)=h(0)'
for i=2:length(h)
    h(i)=omega+params(2)*h(i-1)...
        +params(1)*(epsilon(i-1)+params(3)*sqrt(h(i-1)))^2;
end
forecast=h(end);
h(end)=[]; % remove forecast
```

There are several reasonable approaches to setting the initial value for the conditional variance, Here, I have chosen to set it such that $h_0 = h_1$.

Notice that I have set-up this routine to return ω and the one-step forecast for h for subsequent convenience.

The script below includes several assumptions, including the conditional expectations of the log total returns and the starting values for the numerical search for the MLE parameter estimates. We must avoid negative conditional variances, so this means $\omega \geq 0$, $\alpha \geq 0$, and $\beta \geq 0$. Furthermore, in order to avoid a negative unconditional variance, we require $\alpha(1+\gamma^2)+\beta < 1$.

We could (and perhaps should) implement a constrained search. I found that this was not necessary in this case as long as I initiated the numerical search near to the solution. Generally for financial timeseries, $\omega \approx 0$, $\alpha \approx 0$, and $\beta \approx 1$.

```
%% get data
tsc=yahoo_prices({'GOOG'},'30-Sep-2009','30-Sep-2011');
prices=tsc.GOOG.Data(¬isnan(tsc.GOOG.Data));
%% calculate residuals 'epsilon'
y=log(prices(2:end)./prices(1:end-1)); % log total returns
m=zeros(size(y)); % conditional means (zero)
epsilon=y-m; % residuals
%% estimate NGARCH parameters
params=nan(4,1); % [omega alpha beta gamma]
obj=@(h,epsilon)sum(log(h)+epsilon.^2./h); % quasi-MLE objective
params(2:4)=fminsearch(@(params)...
    obj(NGARCH(epsilon,params),epsilon),[.01 .9 0.]);
[h params(1) forecast]=NGARCH(epsilon,params(2:4));
%% de-volatility returns
z=epsilon./sqrt(h);
```

The values I got for GOOG between October 2009 and September 2011 are in the table below.

ω	4.0×10^{-6}
α	0.052
β	0.98
γ	0.0020