Quantitative Risk Management
Fall Assignment

John Dodson

October 4, 2017

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 11.

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.

You may discuss approaches to the problems with your classmates, but not share code or answers. I expect you to be the sole author of your report and scripts. Cite your sources, including your classmates as appropriate.

Please remind yourself of the University's definition of scholastic dishonesty.

Introduction

We have already worked with the Breeden-Lizenberger model for implied density and the Engle model for conditional heteroskedasticity. The former facilities study of the $\mathbb{Q}$ measure of future prices while the latter facilitates study of the $\mathbb{P}$ measure.

Risk-Neutral Measure

Based on the prices of options expiring on Friday, September 29, the implied forward for QQQ on Monday, September 25, was $F \approx \$142.949$ and the implied discount was $d \approx 0.999452$. The implied volatility curve was approximately

$$\sigma_{BSM}(K)\sqrt{\tau} \approx \sqrt{1.529 \times 10^{-9} + \sqrt{2.028 \times 10^{-8} + 1.625 \times 10^{-4} \log^2(K/F)}}$$

for strike price K and tenor $\tau = 4$ days.

The implied density of the QQQ price is $f^{\mathbb{Q}}_{S_\tau|\mathcal{F}_0}(K) = c''(K)/d$, where the subscript on the density denotes the random variable and the filtration and the superscript denotes the measure.

Since $S_T > 0$, the lower tail index of the price is not very interesting or useful. So let us instead consider log-returns through the (four-day) horizon. Let us define the random variable $U = \log S_\tau$. The density transforms as $f_{U|\mathcal{F}_0}(u) = e^u f_{S_\tau|\mathcal{F}_0}(e^u)$. From this we can calculate $\operatorname{var}^{\mathbb{Q}}[\log S_\tau|\mathcal{F}_0]$ and the tail parameter $\xi^{\mathbb{Q}}_L$.

For the tail parameter, it is useful to use the result in the text based on Karamata's theorem, that

$$\lim_{\eta \to -\infty} \mathrm{E}\left[\log\left(X/\eta\right)|X < \eta\right] = \xi_L$$

for any random variable X with left tail parameter $\xi_L \geq 0$.

Objective Measure

As we did with the last assignment, let's use $n = 1000$ days of historical (log) total returns for QQQ share holders to represent a sample under the objective measure. Let us assume that the mean of the daily returns is zero and that the conditional variance following a GARCH(1,1) process,

$$\sigma_i^2 = \alpha_0 + \alpha_1 \epsilon_{i-1}^2 + \beta_1 \sigma_{i-1}^2$$

for $\alpha_0 > 0$, $\alpha_1 > 0$, $\beta_1 > 0$, and $\alpha_1 + \beta_1 < 1$; and that ϵ_i/σ_i is standard white noise.

Since residuals are independent, the variance we are interested in is

$$\mathrm{var}^{\mathbb{P}}\left[\log S_\tau|\mathcal{F}_0\right] = \mathrm{E}_n\left[\sigma_{n+1}^2 + \sigma_{n+2}^2 + \sigma_{n+3}^2 + \sigma_{n+4}^2\right]$$

Since the standardized residuals are i.i.d., we can fit a Generalized Pareto to, say, the bottom $\theta = 1\%$ to estimate $\xi_L^{\mathbb{P}}$.

Problem Statement

For this project, please estimate the variance and lower tail parameter for the logarithm of a PowerShares QQQ trust share on September 29, 2017 (Friday), based on information available as of the close of trading on September 25, 2017 (Monday). **(20 points each)**.

Grading Rubric

Twenty 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

Options on equity index ETFs are some of the more liquid, and there over-the-counter variance swaps that can be statically hedged by listed options. So market forces tend to keep the $\mathbb{Q}$-measure (log) variance close to the $\mathbb{P}$-measure variance. This is the domain of "volatility traders". Typically we see option prices imply a slightly higher value than GARCH forecasts, probably because of the slippage and transaction costs associated with the dynamic hedge. And this was the case here.

Estimating the tail dependence from log-returns faced some theoretical and technical hurdles. I expect most of you were able to address the technical hurdles, but the theoretical challenges are more significant and less well-known: Namely, the result *must* be zero under at least the $\mathbb{Q}$-measure. This is because the mere existence of derivatives means that the risk-neutral expected value of the underlying value (the forward price) is finite. This means that the expected value of the exponential of the logarithm of the price is finite. And this means that all of the moments of the logarithm of the price are finite. Therefore the normalized n-block maxima are asymptotically Gumbel, hence $\xi_L^{\mathbb{Q}} = 0$.

The story under the $\mathbb{P}$-measure may be a little more nuanced. Our textbook makes reference to returns being heavy-tailed. It is one of the stylized facts about financial data. Most asset managers would concede that the expected value (under any objective or subjective measure) of an asset value is probably finite. Certainly it is assumed in many allocation models, such as the Capital Asset Pricing Model (CAPM). If this is true, clearly for log-returns $\xi_L^{\mathbb{P}} = 0$. If, alternatively, we interpret returns to be simple returns, in fact any asset's simple return is bound below by -100% (assuming limited liability); hence $\xi_L^{\mathbb{P}} < 0$. So we need to reconcile the notion of heavy-tailed with the likely reality that $\xi \not> 0$.

There is a clearly some merit to the notion that not all entropy is manifest through variance (the normal case). In the spring we will look at a generalization of the normal distribution, termed the generalized hyperbolic family. We will find examples of distributions that have relatively low variances for a given level of entropy, and yet have $\xi = 0$ and thus are suitable for log-returns.

Risk-Neutral Measure

Back in the second homework, we evaluated the density of $U = \log\left(S_\tau/F\right)$:

$$f^{\mathbb{Q}}(u) = \Phi'\left(\frac{u}{\tilde{\sigma}(u)} + \tfrac{1}{2}\tilde{\sigma}(u)\right)\left(\frac{1}{\tilde{\sigma}(u)} - \frac{2u\tilde{\sigma}'(u)}{\tilde{\sigma}(u)^2} + \frac{u^2\tilde{\sigma}'(u)^2}{\tilde{\sigma}(u)^3} - \frac{\tilde{\sigma}(u)\tilde{\sigma}'(u)^2}{4} + \tilde{\sigma}''(u)\right)$$

where $\Phi'(\cdot)$ is the standard normal density and $\tilde{\sigma}(u) = \sigma_{BSM}(S_\tau)\sqrt{\tau} = \sqrt{c_1 + \sqrt{c_3 + c_5 u^2}}$. It is convenient to keep F in this definition of u. Clearly $\mathrm{var}\left[\log S_\tau\right] = \mathrm{var}\left[\log\left(S_\tau/F\right)\right]$.

After some symbolic manipulation, I implemented this in Julia as

```
"Gartheral parameters for QQQ weeklies"
c1=1.529E-9;c3=2.028E-8;c5=1.625E-4

"3-factor Gatheral form"
function isd(u)
        return sqrt(c1+sqrt(c3+c5*u^2))
end

"normal density"
φ=z->exp(-z^2/2)/sqrt(2π)

"density for u=log(S/F)"
function f(u)
        σ=isd(u)
        return φ(u/σ+σ/2)*(1/σ+c5/(4σ^3)*
                (2c1*c3/(σ^2-c1)^3-4u^2/(σ^2-c1)-1
                +(3c3+c5*u^2*(u^2/σ^2-σ^2/4))/(σ^2-c1)^2))
end
```

As you might expect, this distribution is very concentrated. In fact, my implementation with IEEE 754 64-bit doubles, we get underflow for $|u| > 20$ as you can see from the results in Table 1. This is not a problem *per se* for calculating the variance, but it is clearly a problem for calculating excess expectations for the second part.

u	f(u)
-20	0.0
-15	3.74734 E-260
-10	8.37474 E-174
-5	2.15969 E-87
0	52.5029
+5	3.20527 E-85
+10	1.84466 E-169
+15	1.22501 E-253
+20	0.0

Table 1: Implemented $f^{\mathbb{Q}}(\cdot)$ exhibiting floating point underflow.

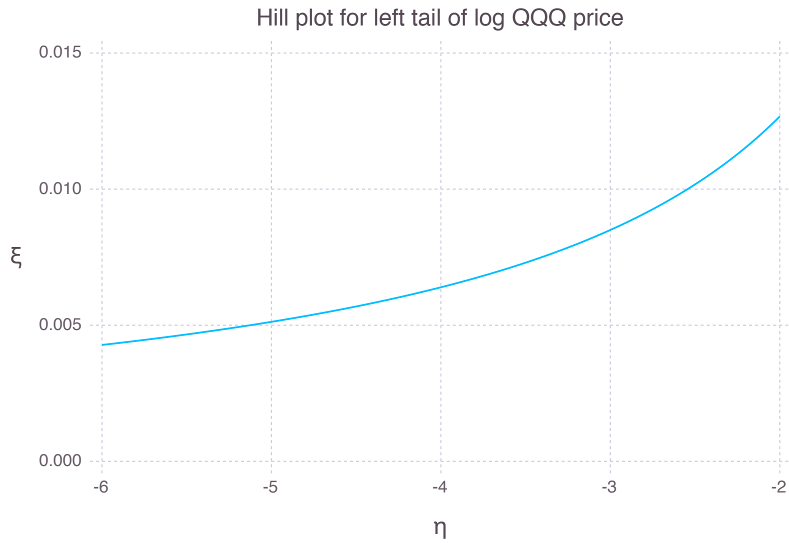

Figure 1: Hill plot for $\xi_L^{\mathbb{Q}}$ of $U = \log(S_\tau/F)$ conditional on the filtration $\mathcal{F}_0$.

If we pick an effective support of $[-15, 15]$, the native `quadgk()` gives a total density within `eps(1.)` of unity.

Based on this implementation, I calculated $\boxed{\mathrm{var}^{\mathbb{Q}}[\log S_\tau|\mathcal{F}_0] \approx 2.629 \times 10^{-4}}$.

We already argued that the lower tail parameter is zero; but we can attempt to inspect it directly using a Hill plot. I implemented this in the following:

```
ll=15. # this is an arbitrary large value for u
"tail parameter"
function ξ(η)
        return quadgk(u->log(u/η)*f(u),-ll,η)[1]/
                quadgk(f,-ll,η)[1]
end
```

The plot seems to suggest that the lower tail parameter, represent by the limit to the left, is positive and less than 0.005.

Objective Measure

The results for the objective $\mathbb{P}$ measure are largely recycled from previous assignments.

I set up my GARCH quasi-ML objective and forecaster here:

```
"GARCH(1,1) conditional variance for θ=[ω,α,β]"
function garch(ε,θ)
        σ²=fill(NaN,length(ε))
        if minimum(θ)>=0 && θ[2]+θ[3]<=1
                σ²[1]=θ[1]/(1-θ[2]-θ[3])
                for i=2:length(ε)
                        σ²[i]=θ[1]+θ[2]*ε[i-1]^2+θ[3]*σ²[i-1]
                end
        end
        return σ²
end

"quasi-MLE entropy for GARCH"
function H_GARCH(ε,σ²)
        return mean(log.(2π*σ²)+ε.^2 ./σ²)/2
end

"forecast daily conditional variances"
function forecast(m,ε,θ)
        σ²=θ[1]/(1-θ[2]-θ[3])
        σ²₀=garch(ε,θ)[end]
        σ²₁=θ[1]+θ[2]*ε[end]^2+θ[3]*σ²₀
        φ=θ[2]+θ[3]
        return σ²₁*φ.^(m-1)+σ²*(1-φ.^(m-1))
end
```

and my Generalized Pareto ML objective here:

```
"Generalized Pareto entropy for θ=[β,ξ] and sample x"
function H_GP(θ,x)
        η=maximum(x)
        if θ[1]<=0. || θ[2]<=-1. ||
                (θ[2]<0. && θ[1]/θ[2]>=minimum(x)-η)
                return NaN
        end
        if abs(θ[2])<eps()
                return log(θ[1])+(η-mean(x))/θ[1]
        end
        return log(θ[1])+(1+1/θ[2])*
                mean(log(1+θ[2]*(η-x)/θ[1]))
end
```

The variance estimate is implemented here:

```
θ₀=[var(ε)*(1-1/3-1/3),1/3,1/3] # initial value for Nelder-Mead
mle_GARCH=optimize(θ->H_GARCH(ε,garch(ε,θ)),θ₀).minimizer
fore=forecast([1:4;],ε,mle_GARCH)
varP=sum(fore)
```

which yields a value $\boxed{\mathrm{var}^{\mathbb{P}}\left[\log S_\tau|\mathcal{F}_0\right] \approx 2.441\times 10^{-4}}$.

The tail parameter estimate is implemented here:

```
z=sort!(ε./sqrt.(garch(ε,mle_GARCH)))
x=z[1:10]
θ₀=[std(x),0.] # initial value for Nelder-Mead
mle_GP=optimize(θ->H_GP(θ,x),θ₀).minimizer
```

which yields a value $\hat{\xi}_L^{\mathbb{P}} \approx -0.137$.

Based on the Cramér-Rao result, the standard error of this estimate is at least 0.273. So we certainly cannot reject a statistical hypothesis that $\xi_L^{\mathbb{P}} = 0$.

Summary

The result we got are summarized in Table 2.

	variance	tail parameter
risk-neutral measure	$\text{var}^{\mathbb{Q}}[\log S_\tau \vert \mathcal{F}_0] \approx (0.0162)^2$	$0 \leq \xi_L^{\mathbb{Q}} < 0.005$
objective measure	$\text{var}^{\mathbb{P}}[\log S_\tau \vert \mathcal{F}_0] \approx (0.0156)^2$	$\xi_L^{\mathbb{P}} = -0.137 \pm 0.546$ (two se)

Table 2: Summary of results for the logarithm on the QQQ price at the close of September 29, 2017, based on data from the close of September 25, 2017.

It is notable that the standard deviation under the risk-neutral measure is about 4% higher than the standard deviation under the objective measure. As I stated in the introduction to the solution, this is a violation of arbitrage theory, but it is also typical. It is possible that the market uses a more sophisticated estimator for the objective measure. More likely the gap represents the cost friction of implementing the arbitrage trade (selling options and dynamically hedging).