

Quantitative Risk Management

Spring Assignment

John Dodson

February 7, 2018

This assignment is not a regular homework. It is worth half of the module grade for the spring term. Please share your solution with me through Google Drive before midnight on Sunday, February 25.

Please write up your results into a report and submit a PDF version. (My edition of LibreOffice can recognize most word processing formats; but sometimes the formatting, equations, and exhibits are altered or corrupted.) You may be able to covert directly to PDF from your word processor or type setter, or you may be able to “print to PDF” using a print driver.

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.

Risk Parity

The Markowitz efficient frontier was immediately a compelling concept on the buy-side from its origins in 1952; and it introduced the notions of passive investment and capitalization-weighted indexes, index tracking mutual funds, and ultimately exchange-traded funds. But there has always been a tension between the equilibrium side and the optimization side of Markowitz’ theory, because the optimization problem seems so intractable in practice.

Two generations later, in 1996, the concept we now know of as “risk parity” started to be popularized in this environment. The idea here was to strip away the hard part of the optimization, the forecasting of relative returns, and instead focus on diversification on its own merits.

In addition to improved econometrics, such as models for conditional heteroskedasticity and tail dependence in asset returns, theoretical innovations that facilitated this development included a recognition that the Markowitz utility-based risk measure (certainty-equivalent exponential utility) lacked some key features, including positive-homogeneity and concavity which facilitate allocation by gradient and convex programming, that are offered by more modern coherent risk measures such as expected shortfall.

The defining characteristic of a risk-parity allocation is that the marginal risk of each asset is identical.

Constructing an Expected Shortfall Risk Parity Portfolio

The goal of this project is to identify a risk-parity allocation of QQQ constituent stocks on January 19, 2018, based on one-day 95% expected shortfall. Use the simulations based on the Student’s- t_4 copula draws to measure this.

The marginal decomposition of expected shortfall based on simulations is quite easy: Say there are 10,000 simulations. For a given portfolio, you can easily calculate the net asset value in each simulation $\omega = 1, 2, \dots$. And that net asset value is simply a sum of ($\text{quantity}_i \times \text{price}_{i,\omega}$) for each asset $i = 1, 2, \dots, N$. We can readily identify the scenarios that correspond to the lowest 500 portfolio net asset value simulations. Call this set Ω . The expected shortfall is the average loss in these 500 scenarios. This average is also the sum of the average loss of each asset in these scenarios, $\frac{1}{|\Omega|} \sum_{\omega \in \Omega} \text{quantity}_i \times \text{loss}_{i,\omega}$.

So, in principle you could identify a vector of (quantity)_{*i*=1,...,*N*} such that the average loss on each asset in the set Ω is equal to some fixed value. The problem is, the set Ω also depends on that vector! So you need to search for a joint solution such that the marginal expected shortfalls are equal and add up to the total expected shortfall.

The risk parity portfolio generally includes leverage: the weights in the portfolio will not add to one. To the extent that the weights are less than one, there is a (risk-free) cash component. To the extent that the weights are greater than one, some of the assets have been pledged as loan collateral. The cash position will effect the expected return, but let's assume that it does not affect the expected shortfall.

Problems

1. Calculate the one-day 95% expected shortfall on the QQQ creation unit portfolio as of January 19, 2018. This will be our risk target, which we would like to allocate equally among the QQQ constituent stocks. It will also seed our set Ω_0 of loss scenarios¹. **(30 points)**
2. Determine the weights of a portfolio such that average loss in Ω_0 is equal for each constituent and sums to the target from (1.). **(30 points)**
3. Re-calculate the expected shortfall and the loss scenarios Ω_1 , and scale the weights to equalize the expected shortfall in Ω_1 back to the target from (1.). Repeat (2.)–(3.). Does it seem to converge? **(10 points)**
4. If it converges, calculate the leverage of the resulting portfolio (net asset value long plus net asset value short divided by net asset value) and list the holdings with the largest overweights and underweights relative to the QQQ creation unit. **(10 points)**

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 and includes appropriate citations. **(10 points)**

Solution

I will present my solution as a narrative along with my Julia² implementation.

Data Model

Please review the solution to the assignment from January 24, 2018, for notes and recommendations about the data management for this project.

For our purposes, assume that we have a long-form (multi-key) table with the columns listed in Table 1 populated with two years of daily data through January 19, 2018 related to the January 22, 2018 QQQ creation unit constituents.

We will collect results for the January 19, 2018 analysis in a wide-form table with columns listed in Table 2.

¹The constituent scenarios are the same throughout.

²<https://julialang.org/>

Ticker	(key) exchange symbol
Date	(key) trading session date
AdjClose	close price, adjusted for corporate actions
PrevDate	previous trading session date
PrevAdjClose	previous adjusted close price
LogReturn	daily log-return

Table 1: columns in the `tser` data frame

Ticker	(key) exchange symbol
GARCHalpha	GARCH(1,1) ARCH coefficient
GARCHbeta	GARCH(1,1) GARCH coefficient
GARCHVar	GARCH(1,1) conditional variance forecast
SampleVar	sample variance
NRIGg	symmetric NRIG shape parameter
AdjClose	close price on January 19, 2018
CUQuantity	QQQ creation unit share quantity for January 22, 2018
CUWeight	QQQ creation unit fraction of net asset value
RPWeight	risk parity portfolio weight
ActiveWeight	risk parity portfolio weight relative to QQQ creation unit
MarginalES95	marginal risk for risk parity portfolio position

Table 2: columns in the `rslts` data frame

Econometrics

As in the previous exercise involving QQQ creation unit, we fit GARCH(1,1) with symmetric NRIG residuals to daily log returns for each stock price timeseries and store results in the `rslts` table. Along the way, we store Gaussianized residuals into a design matrix `z`, from which we estimate pseudo-correlations `P`.

`using` DataFrames, Optim, SprAssign

```

wtser=sort!(unstack(tser, :Date, :Ticker, :LogReturn)) # cross-tabulation by date/ticker
wtser=delete!(wtser, :Date) # observations are assumed to be at a regular frequency
z=zeros(size(wtser,1), size(rslts,1)) # allocate for Gaussianized residuals
for i=1:size(rslts,1)
    ticker=rslts[i, :Ticker]
    ε=Array(wtser[ticker])
    θmle=optimize(θ->H_GARCH(ε, GARCH(ε, θ)), [.1, .7]).minimizer # use variance targeting
    (h, h1)=GARCH(ε, θmle, true)
    x=ε./sqrt.(h)
    gmle=optimize(g->H_NRIG(x, g), 0., 4.).minimizer
    z[:, i]=quan_Normal.(dist_NRIG.(x, gmle))
    rslts[i, :GARCHalpha]=θmle[1]
    rslts[i, :GARCHbeta]=θmle[2]
    rslts[i, :GARCHVar]=h1
    rslts[i, :SampleVar]=var(ε)
    rslts[i, :NRIGg]=gmle
end
Σ=cov(z)
u=diag(Σ).^(-1/2)
P=Σ.*(u*u') # pseudo-correlations

```

Scenarios

As before, we generate ten thousand scenarios for each constituent stocks.

```
N=10_000
scen=exp.(
    sqrt.(rslts[:GARCHVar]).*quan_NRIG.(
        dist_t4.(chol(P)'*randn(size(P,1),N)./sqrt.(mean(randn(4,N).^2,1))),
        rslts[:NRIGg]
    )
)
```

Note that we exponentiate the log-returns in the last step so that (simple-return) scenarios for a portfolio are simply the dot product with the portfolio's weights.

This scenario set remains fixed in the succeeding.

Risk Measurement

First, we calculate the 95% expected shortfall on the QQQ creation unit. This is simply the average of the worst 500 ($\text{div}(N, 20)$) simulated losses.

```
CU_scen=scen'*rslts[:CUWeight]
tail95=sortperm(CU_scen[:])[1:div(N,20)]
CU_es95=1-mean(CU_scen[tail95])
```

resulting in 1.8%.

Since there are about 21 trading days per month, 95% one-day expected shortfall is approximately the size of the loss one should expect if tomorrow were the worst day of the month and there were no material changes to the portfolio in the meantime. That is a quantity that most investment managers understand intuitively³.

Portfolio Construction

We are looking for a portfolio that has the same expected shortfall, but with marginal risk evenly spread across each holding. In my case, the investment universe has 96 stocks since we are missing historical data on eight of the 104 positions in the creation unit portfolio⁴.

```
M=96
RP_wgt=rslts[:CUWeight]
RP_scen=scen'*RP_wgt
for i=1:100
    RP_wgt=CU_es95/M./(1-mean(scen[:,tail95],2))
    RP_scen=scen'*RP_wgt
    tail95=sortperm(RP_scen[:])[1:div(N,20)]
end
```

I tried one hundred and one thousand iterations and got similar results for fixed scenarios. I also got broadly similar results with different versions of scenarios, but it is notable that with ten thousand scenarios the risk-parity goals is not exactly achieved: With fixed scenarios there are probably several results that are very similar, and iterations cycle between them.

Once we have settled on the risk parity loss scenarios, we can re-normalize the weights

```
RP_es95=sum(RP_wgt)-mean(RP_scen[tail95])
RP_wgt=RP_wgt*CU_es95/RP_es95
```

³Higher confidence levels are common in banking risk measurement, but those results tend to be more abstract since exposures and risk factors can evolve significantly between exceedance experiences.

⁴Some of the stocks in our set are closely related, such as the two Alphabet (Google) share classes and the cluster of Liberty Media shares, so in practice we may choose to divide the risk budget for a particular issuer between the various listed classes.

and record the results.

```

rslts[:RPWeight]=RP_wgt[:]
rslts[:ActiveWeight]=rslts[:RPWeight]-rslts[:CUWeight]
rslts[:MarginalES95]=(RP_wgt.*(1-mean(scen[:,tail95],2))[:])

```

The result has only long positions, but the sum of the weights is about 108% so implicitly there is a short cash position (a collateralized loan) of about 8% of the net asset value. In practice, one could scale down all of the positions to manufacture a total weight of 100% and no cash position. The expected shortfall of this “fully-invested” portfolio would be about 8% less than that of the creation unit, or about 1.7%.

Some results are presented in Table 3.

Ticker	Weight			Marginal Risk
	Risk Parity	Creation Unit	Active	
AAPL	1.32%	11.96%	-10.64%	0.0181%
ADSK	0.64%	0.33%	+0.31%	0.0190%
AMZN	1.19%	8.14%	-6.95%	0.0181%
CMCSA	1.27%	2.59%	-1.31%	0.0191%
COST	2.33%	1.10%	+1.23%	0.0204%
CSCO	1.32%	2.66%	-1.34%	0.0191%
FB	1.02%	5.64%	-4.62%	0.0182%
FISV	1.63%	0.37%	+1.26%	0.0194%
GOOG	1.32%	5.18%	-3.86%	0.0182%
GOOGL	1.30%	4.45%	-3.15%	0.0182%
ILMN	0.63%	0.46%	+0.17%	0.0190%
INTC	1.30%	2.74%	-1.44%	0.0191%
KHC	1.96%	1.27%	+0.70%	0.0193%
LRCX	0.67%	0.44%	+0.23%	0.0188%
MSFT	1.12%	9.06%	-7.95%	0.0187%
MU	0.56%	0.64%	-0.08%	0.0189%
NVDA	0.59%	1.82%	-1.23%	0.0183%
SBUX	2.23%	1.14%	+1.09%	0.0194%
STX	0.65%	0.20%	+0.46%	0.0197%
WBA	1.79%	0.99%	+0.80%	0.0201%

Table 3: selections from the 95% one-day expected shortfall risk parity solution for the QQQ constituents on January 19, 2018

Discussion

When I wrote this assignment, I thought it would be interesting to compare the risk parity portfolio to the creation unit portfolio in terms of active weights, that is the risk parity weights minus the creation unit weights by position. This is conventionally how one thinks about the source of outperformance and “tracking error” relative to a benchmark. But there is nothing “active” in this construction: it is just two different versions of a passive portfolio. It is not even clear which one should be the benchmark. In retrospect, I maybe should have asked for a discussion of the overweights and underweights relative to an equal weight allocation.

The largest holdings in the creation unit are Apple, Alphabet (Google), Microsoft, and Amazon. All tech stocks. Since the NASDAQ-100 index is so technology-focused, it is no surprise that these are somewhat higher contributors to marginal risk than typical stocks, so their weights are significantly lower in the risk parity portfolio. In fact, the largest holdings⁵ are Costco, Starbucks, and Kraft Heinz, none of which is significantly involved in selling technology.

⁵Although the combined weight of the two Alphabet (Google) share classes in fact makes it the largest holding in the risk parity portfolio.

Results are a little more interesting when applied to a wider class of investments, for example debt and equity, because it leads to “passive” results involving substantial leverage.

Appendix: Julia definitions

```

module SprAssign

using SpecialFunctions, QuadGK

export GARCH, H_GARCH,
       dens_NRIG, dist_NRIG, quan_NRIG, H_NRIG,
       dist_Normal, quan_Normal,
       dist_t4

"GARCH(1,1) conditional variance for  $\theta=[\alpha,\beta]$  or  $[\omega,\alpha,\beta]$  with an optional forecast"
function GARCH( $\epsilon$ ,  $\theta$ , forecastQ=false)
    if (length( $\theta$ )==2)
         $\theta$ =[var( $\epsilon$ )*(1- $\theta$ [1]- $\theta$ [2]); $\theta$ ] # optional variance targeting
    end
    h=fill(NaN,length( $\epsilon$ )+1)
    if length( $\theta$ )==3 && minimum( $\theta$ )>=0 &&  $\theta$ [2]+ $\theta$ [3]<1
        h[1]= $\theta$ [1]/(1- $\theta$ [2]- $\theta$ [3]) # initialize at unconditional variance
        for i=2:length(h)
            h[i]= $\theta$ [1]+ $\theta$ [2]* $\epsilon$ [i-1]^2+ $\theta$ [3]*h[i-1]
        end
    end
    if forecastQ
        return (h[1:end-1],h[end])
    else
        return h[1:end-1]
    end
end

"quasi-log-likelihood objective for GARCH"
function H_GARCH( $\epsilon$ ,h)
    return mean(log.(2 $\pi$ *h)+ $\epsilon$ .^2 ./h)/2
end

"distribution of a standard symmetric normal / reciprocal inverse Gaussian"
function dens_NRIG(x,g)
    return exp(g)sqrt(1+g)besselk(0,sqrt(g^2+(1+g)*x.^2))/ $\pi$ 
end

"log-likelihood objective for standard symmetric NRIG shape parameter"
function H_NRIG(x,g)
    # 'x' is an array of samples
    return -mean(log.(dens_NRIG.(x,g)))
end

"standard normal quantile function"
function quan_Normal(p)
    return -sqrt(2)erfcinv(2p)
end

"NRIG distribution function"

```

```

function dist_NRIG(x,g)
    return 1/2+quadgk(z->dens_NRIG(z,g),0,x)[1]
end

"distribution function of a standard Normal random variable"
function dist_Normal(x)
    return erfc(-x/sqrt(2))/2
end

"distribution function of a conventional Student's-t_4 random variable"
function dist_t4(x)
    return (1+(x*(6+x^2))/((4+x^2)^(3/2)))/2
end

"standard NRIG quantile function"
function quan_NRIG(p,g)
# Newton-Raphson method to invert dist_NRIG
    x0=quan_Normal(p) # initial guess based on Normal result
    x1=NaN
    for n=1:20
        d=dens_NRIG(x0,g)
        x1=(p-1/2+quadgk(x->d-dens_NRIG(x,g),0,x0)[1])/d
        if abs(x1-x0)<1e-12
            break
        end
        x0=x1
    end
    return x1
end

end # SprAssign

```