Review of Economics & Finance

Submitted on 10/03/2017

Article ID: 1923-7529-2017-03-29-13 Chirag Shekhar, and Mark Trede

Portfolio Optimization Using Multivariate t-Copulas with Conditionally Skewed Margins

Chirag Shekhar

Birla Institute of Technology and Science, Pilani - K.K. Birla Goa Campus, INDIA

E-mail: chirag.shekhar@gmail.com

Prof. Dr. ***Mark Trede***

Department of Economics, Westfälische Wilhelms-Universität Münster

Am Stadtgraben 9, 48161 Münster, GERMANY

Tel: +49-251-8322951 E-mail: mark.trede@uni-muenster.de

Abstract: Over the last few decades, copulas have consistently gained significance in finance research, due to their usefulness in risk modeling. However, the idea of implicitly representing dependencies between multiple assets in a single mathematical entity is extremely useful in portfolio allocation models as well. While Church (2012) and many others have exploited these benefits, the efficiency of such frameworks in capturing the most essential features of financial data can still be enhanced. An obvious improvement would be to incorporate the fact that financial returns are generally asymmetric and skewed in nature, and therefore asymmetric (or skewed) margins can be used to describe them in a suitable copula framework. In this paper, we consolidate this idea with a GARCH(1,1) pre-whitening method that takes into account inter-temporal dependencies of returns, and use a utility maximization approach to find optimal portfolio allocation schemes. We show that the gains of optimal weighting, in terms of certainty equivalent returns, can be substantial for utility functions with reasonable risk aversion.

Keywords: Utility maximization; Portfolio management; GARCH process; Multivariate return distributions; Copula

JEL Classifications: G11, C58, C61

1. Introduction

When Markowitz (1952) posited what we now consider the modern portfolio theory, not a lot of work had been put into figuring out better ways to understand the price data that was used to inform the mean-variance framework he proposed. Since then, several advances in financial mathematics have helped "decipher" the nature of these data series. These advancements have in turn helped formulate models that emulate the real world better. Multiple approaches to Bayesian estimation of parameters obtained from historical data, notably the ones suggested by Garlappi *et al.* (2005), Avramov and Zhou (2010) and Lai *et al.* (2011) have been widely cited and applied in recent literature. Factor models proposed by Fama and French (1993), have also served as major improvements over conventional pricing models rooted in the mean-variance optimization framework.

Over the same period, copulas, originally used in Sklar(1959) and initially limited to applications in engineering and other sciences, were found to be very useful in risk analysis and rapidly gained traction with financial researchers. A copula is, simply put, a multivariate probability distribution whose marginal distributions are uniformly distributed on the interval [0,1]. This allows us to analyze non-linear dependencies between assets, and so copula based models can be used to implicitly incorporate complex interdependencies between assets, and thus offer an improvement over isolated simulation methods.

Such a process must develop on a justifiable premise, which includes the choice of the copula and the marginals used to represent individual return series. The approach we illustrate in this paper is loosely inspired by those implemented by Church (2012) and Kakouris and Rustem (2014). A robust optimization procedure suggested in the latter deploys an Archimedian copula for a cVaR based portfolio optimization approach, citing the effectiveness of this particular copula framework in modeling non-linear asset dependencies. In the latter, a two-layered simulation method is used instead, which models dependencies between asset returns using an asymmetric skewed t copula with symmetric student's t-distributed marginals. For this study, we sought to use a symmetric copula in a setup closer to the latter, considering the fact that the dependence structure between the financial series, as is illustrated later, is rarely asymmetric. At the same time, marginal distributions, which represent each return series do exhibit a typical negative skewness, and therefore, a skewed t-distribution is used to capture that effect.

Another important aspect of the model we purpose here is that the data is filtered to remove noise and selectively retain only the information that is relevant to the simulation procedure. A particularly important attribute of financial return series is volatility clustering Cont (2007) and several GARCH and stochastic volatility models have been developed specifically to model this phenomenon. We use a GARCH(1,1) process, as introduced in Breidt *et al.* (1998) to preserve this quality within the simulated series. Similar procedures are explained in Jondeau and Rockinger (2006) and Liu and Luger (2009), however, these models are focused specifically on estimating risk, and that is a narrower application than what follows.

The improvements proposed here would, however, fail to produce efficient results unless an optimization framework compliant with these assumptions is used. The traditional Markowitz approach is not only inconsistent with our assumptions about the marginal distributions of return series, but its inefficiencies have been exposed by multiple studies, most significantly by the success of another optimization approach based on cVaR (conditional value at risk) computation, proposed by Krokhmal *et al.* (2002). An obvious competing approach, which must be considered at this point, is expected utility maximization. Sharpe (2007) suggested a straightforward iterative portfolio optimization algorithm that allows bounds on the portfolio weights, and is the method of choice in our model. This approach assumes that an investor would always try to maximize the utility they derive at the end of their planning horizon, rather than blindly minimizing variance for maximal returns. Various paradigms and modular segments of such an approach are then open to discussion. Risk preferences of the investor can be implicitly incorporated into the utility function, subtly factoring in the utility of other market participants (Buccola and French, 1978). This approach is thus, a versatile one and if correctly configured, can produce very efficient results.

The performance of a copula estimated mean-variance portfolio was consistently better in the context of US and Latvian securities, according to one such study conducted by Jansons *et al.*

(2005). In the expected utility maximization framework that allows to incorporate investors' expectations, coupled with an extensive estimation model for enhanced accuracy in simulation, we produce a more realistic and efficient model. We show that the gains from adopting the optimal weighting scheme are substantial for an investor having a utility function with a reasonable amount of risk aversion.

We now proceed to elaborate on each of the aspects indicated above, the arrangement of which is as follows. In the first section, we introduce the approach we use to model multivariate returns, describing first the copula and then the GARCH filtering approach adopted in the procedure. Next, we develop the estimation procedure, introducing the iterative algorithm used to consolidate the two. In the following section, we introduce the utility maximization framework we have used in this model, and ultimately, we present an illustrative portfolio to describe the empirical performance of the model.

2. Model of Multivariate Returns

Pre-whitening and data filtering methods are a range of procedures used in data science and statistics to isolate systemic information about inter-temporal dependencies which is not relevant to estimating the contemporary dependence between data series. In the case of financial return series, we particularly seek to retain information like 'volatility clusters' for each individual asset. As usual in a time series setting, the return of asset *i* in period *t* is defined as

$$r_{it} = \ln(K_{it}/K_{i,t-1})$$

where K_{it} is the asset price.

Volatility clustering is a property of data that is often exhibited by financial return series, where large changes in prices tend to cluster together. This causes a persistence in the absolute (or squared) magnitude of the return series (Cont, 2007). Several GARCH and stochastic volatility models have been developed specifically to model this phenomenon (Breidt *et al.*, 1998).

In figure 1, we see a clear example of how often this phenomenon is observed in the case of financial return series. The returns show extended stretches of high or low volatility, and the volatility clusters are clearly discernible.

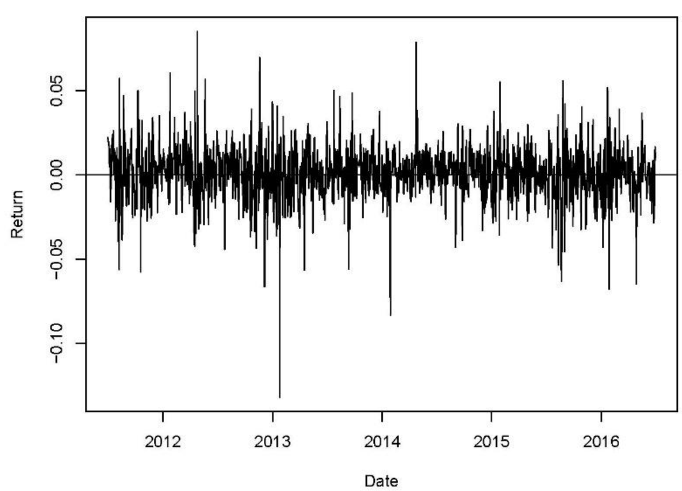

Figure 1. Daily stock returns of Apple Inc., from 1 July 2011 to 30 June 2016

2.1 GARCH filtering

One of the first models to account for volatility clusters in time series were ARCH and GARCH models, introduced by Engle (1982) and Bollerslev (1986). In a GARCH(1,1) model, volatility depends on the preceding period's volatility and squared return. The process is described as:

$$r_t = \mu + \sigma_t \varepsilon_t \qquad (1)$$

$$\sigma_t^2 = \omega + \alpha(r_{t-1} - \mu)^2 + \beta\sigma_{t-1}^2 \qquad (2)$$

where r_t is the return series, μ is the mean of this return series, and σ_t is the time-varying volatility. ε_t is the innovation variable, which has a pre-defined distribution, typically a normal or t-distribution, with zero expectation and unit variance. Equation (2) implies that there will be a positive autocorrelation in the volatility process, σ_t, with a rate of decay governed by the parameters $\alpha > 0$ and $\beta > 0$. If $\alpha + \beta < 1$ we have a stationary process, and generally, GARCH estimates on financial return series do conform to these bounds.

Given a time series r_t, we can fit a GARCH(1,1) process with a pre-defined conditional distribution for the innovation variable ε_t. The fitted model can then be used to simulate return series that exhibit the same inter-temporal dependence structure and hence volatility clusters as the original time series. We will now discuss the details of such a procedure and how the simulation and estimation method is configured.

2.2 Standard t-Copula

The standard t-copula is a generalization of the Gauss copula, much like the students t-distribution is a generalization of the normal distribution in probability theory. The standard d-dimensional t-copula is a function C defined as the following

$$C(\boldsymbol{u}; R,v) = t_{R,v}(t_v^{-1}(u_1),\ldots,t_v^{-1}(u_d)) \qquad (3)$$

where $\boldsymbol{u} \in (0,1)^d$ and t_v^{-1} is the inverse of the univariate standard Student-t cdf. $t_{R,v}$ is the multivariate standard Student-t cdf with the linear $(d \times d)$-correlation matrix R and degrees of freedom, v. For $v \to \infty$, we approach the Gauss copula. Figure 2 compares the bivariate Gauss copula to the t-copula with $v = 3$ degrees of freedom when both copulas have a correlation parameter of $\rho = 0.8$. In contrast to Gaussian copulas, t-copulas exhibit tail dependence both in the lower and upper tail. Besides, extreme values in the off-diagonal corners are more likely than under normality.

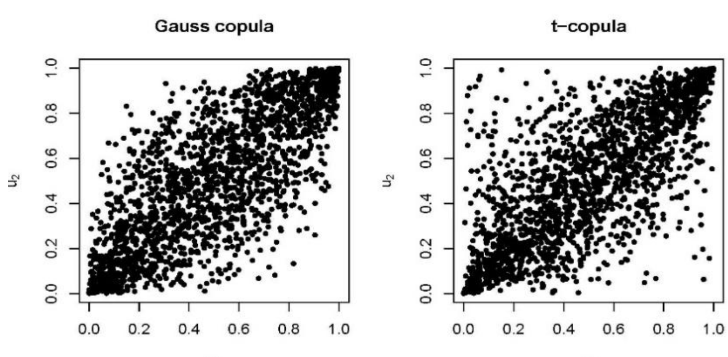

Figure 2. A comparison between bivariate Gauss copula with correlation parameter 0.8 and standard t-copula with correlation 0.8 and 3 degrees of freedom

3. Estimation and Simulation Procedure

The first step is to fit a GARCH(1,1) model to each return series $j = 1,...,N$. It is vital to choose a suitable conditional distribution for the innovations ε_t. This distribution conditionally describes the nature of the return series and captures its behavior. Since a financial series is typically closer to an asymmetric or skewed t-distribution, a good choice for this conditional distribution is a skewed student's t-distribution.

Upon fitting the model, we obtain the estimated parameters for equations (1) and (2), i.e., $\hat{\mu}$, $\hat{\omega}$, $\hat{\alpha}$ and $\hat{\beta}$, for each asset. Additionally, we estimate the parameters of the innovation distribution, which in our case is a skewed t-distribution. Therefore, we also obtain the parameters $\hat{\delta}$ (degrees of freedom) and $\hat{\gamma}$ (skewness) for each univariate return series. Note that the degrees of freedom parameter δ of the innovations is, in general, different from the copula's degrees of freedom parameter v. Apart from the estimated parameters we also compute the residuals $\hat{\varepsilon}_t$.

Once the fit is completed, we generate a filtered return series. To do so, we use the vector of estimated $\hat{\varepsilon}_t$, and the information regarding its distribution. Using a probability integral transform (PIT), we obtain a vector of observations of a uniformly distributed variable, which we call u_t. Let $F_{\hat{\delta},\hat{\gamma}}$ be the distribution function for the innovation parameter's skewed student's t-distribution, then for each asset

$$u_t = F_{\hat{\delta},\hat{\gamma}}(\hat{\varepsilon}_t).$$

We can now join these uniformly distributed variables by a copula. The copula parameters v and R can then be estimated by standard procedures (e.g., maximum likelihood method).

The estimated copula is used to simulate a matrix composed of vectors of uniformly distributed variables, $\tilde{u}_t$. We can then transform it into a matrix of innovation parameters using the quantile function $F^{-1}_{\hat{\delta},\hat{\gamma}}$,

$$\tilde{\varepsilon}_t = F^{-1}_{\hat{\mu},\hat{\delta},\hat{\gamma}}(\tilde{u}_t).$$

This results in a vector of $\tilde{\varepsilon}_t$for each return series. Now, we use the estimated values of $\hat{\omega}$, $\hat{\alpha}$ and $\hat{\beta}$ to obtain the simulated return series

$$\tilde{r}_t = \mu + \tilde{\sigma}_t\tilde{\varepsilon}_t$$

by obtaining a vector of simulated $\tilde{\sigma}_t$, given by the iterative process

$$\tilde{\sigma}_t^2 = \hat{\omega} + \hat{\alpha}(\tilde{r}_{t-1} - \hat{\mu})^2 + \hat{\beta}\tilde{\sigma}_{t-1}^2$$

where $\tilde{\sigma}_0^2 = \hat{\omega}/(1 - \hat{\alpha} - \hat{\beta})$.

This procedure results in a matrix of simulated return series that exhibit both the same inter-temporal and the same between-assets dependence structure as the observed data. This matrix shows a realistic potential scenario of asset returns. By drawing not just a single matrix but many matrices we can approximate the joint distribution of asset returns closely and use it to find the optimal, utility-maximizing portfolio weights.

To summarize, the following steps in their particular order yield a matrix of simulated returns that can be used for portfolio optimization:

1. Obtain the returns for each of the N assets in question.
2. Fit a GARCH(1,1) model to each of the N return series to obtain the parameters of the conditional distribution and the volatility process.
3. Perform a PIT on the estimated innovations using the estimated parameters, and obtain N vectors of uniformly distributed variables, one for each asset.
4. Fit a t-copula to the matrix of these N vectors, and estimate the copula parameters.
5. Simulate a set of N vectors from the estimated copula.
6. Using the simulated uniformly distributed vectors, produce simulated innovations using the quantile function for the conditional distribution.
7. Use the estimated parameters of the GARCH process to simulate the returns from these innovations.

Steps 5 to 7 are replicated to produce a large number of draws from the joint asset return distribution. The results of this procedure for 'AAPL' (using the data in figure 1) are indicated in figure 3, where the retention of volatility clusters is immediately apparent. Several other benefits of this procedure are included in our discussion of an illustrative portfolio, which follows later.

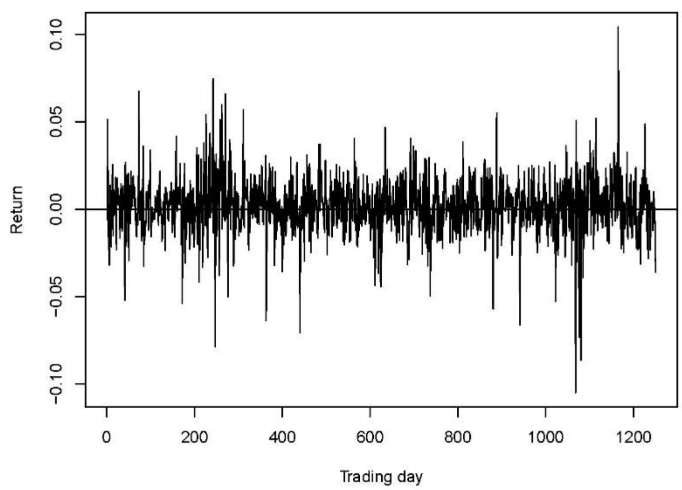

Figure 3. One simulated path of the daily stock returns for Apple Inc. over 1250 observations (The returns were drawn from the fitted model.)

4. Utility Maximization Framework

The fundamental assumption of the expected utility framework is that investors seek to maximize their utility, as they perceive it, out of their portfolio. Thus, with every state of portfolio returns there is an associated utility, or a measure of 'happiness' associated with the outcome (total return in that state). If the return of asset $i = 1,\ldots,N$ in period t is r_{it}, the gross return of a portfolio p with weights $w_{1t},\ldots,w_{Nt}$ is

$$R_t^{(p)} = \sum\nolimits_i w_{it} e^{r_{it}}.$$

The weights are the proportions of total wealth held in each asset, hence they add to unity. Let $f_t(r_1,...,r_N)$ denote the joint density function of returns. The expected utility for the portfolio is

$$E(u) = \int_{-\infty}^{\infty} ... \int_{-\infty}^{\infty} u\left(R_t^{(p)}\right) f_t(r_1, ..., r_N) dr_1 ... dr_N \qquad (4)$$

The optimization algorithm suggested by Sharpe (2007) is set up for discrete return distributions. We transform the approach such that it can be applied to continuous distributions. To keep the notation simple, we drop the time subscript t in the description of the optimization algorithm.

The first derivative of utility is marginal utility, which we will denote as $u'(R^{(p)})$. For risk averse investors, marginal utility decreases as $R^{(p)}$ increases. Assuming that individual assets can only have weights between an upper and a lower limit, i.e., $l_i \leq w_i \leq u_i$, and adding to 1, we have a constrained optimization problem, where $E(u)$, in equation (4) is maximized subject to these constraints.

4.1 Optimization algorithm

The marginal expected utility (MEU) of the portfolio with respect to weight w_i

$$\frac{\partial E(u)}{\partial w_i} = \int_{-\infty}^{\infty} ... \int_{-\infty}^{\infty} u'\left(R^{(p)}\right) \exp(r_i) f(r_1, ..., r_N) dr_1 ... dr_N.$$

Note that the restriction that the weight sum to unity has been ignored. It will turn out that this simplification is justified for the optimization algorithm.

Consider a portfolio p where $\partial E(u)/\partial w_i < \partial E(u)/\partial w_j$. Suppose that both asset i and j are tradable in the sense that their weights are between the lower and upper bounds. We have a possibility of improving the utility of the portfolio since asset j offers a higher marginal expected utility than asset i. So, asset i must be sold and the proceeds are used to buy asset j.

Thus, for a portfolio with N assets, we have a possibility to ascertain whether such an exchange can be made to improve the portfolio. First, we figure out the best buy and the best sell, by calculating the MEU for all assets. The one with the highest MEU and weight $w_j < u_j$, is the best buy, and the one with the lowest MEU and weight $w_i > l_i$ is the best sell. Notice that since the weights still sum up to one after shifting part of the portfolio value from i to j, we are consistent with the mathematical necessities of portfolio allocation. Therefore, the optimization algorithm is as follows:

1. Calculate the MEU for every asset. Label the one with lowest MEU the best sell, and the one with the highest MEU as best buy.
2. Compare the lower bound of the best sell with its current weight in the portfolio, and the upper bound of the best buy with its current weight in the portfolio. If they are both within limit, reduce the weight of the best sell by a small amount and allocate it to the weight of the best buy, then go to step 1. Else go to step 3.
3. Remove assets with weights already at the bounds from the collection of assets for consideration in the next step. If there are still 2 or more assets under consideration, repeat 1.

4.2 HARA utility function

Now that we have presented a universal optimization algorithm, we discuss the specifics of the utility function that describes the investor's preferences. A requirement that most common utility functions meet is that expected utility is additively separable in the states, and a class of functions consistent with it, and suitable for a wide range of investor preferences is the HARA utility function. The utility of gross return R is expressed as

$$u(R) = \frac{(R-b)^{1-c}}{1-c} \tag{5}$$

and the marginal utility function is

$$u'(R) = (R-b)^{-c} \tag{6}$$

The value of b can be considered to be the minimum required level of return because the marginal utility of returns at this level becomes infinite. The value of c, a parameter of how risk averse the investor is, is usually greater than one, but in the optimization algorithm described earlier we may also use $c \to 1$ (in the limit $u(R) = \ln(R-b)$).

A particular case arises when $b = 0$, when the investor has constant relative risk aversion (CRRA). A more general view arises either side of that situation. When $b < 0$, we have increasing relative risk aversion, but if $b > 0$ we have a decreasing relative risk aversion. Of course, the expected utility framework does not depend on this particular utility function but can be adapted to any kind of utility function. Note however, that the existence of the expectation of utility must be ensured. Depending on the particular form of the utility function, the existence might be compromised if the return distributions have heavy tails. In our calculations, we guarantee that expected utility is well-defined by setting b to a small negative value.

4.3 Calculation with simulated data

In our case, we approximate the joint distribution of asset returns by a high number of observations, simulated using the procedure described in the previous section, to calculate the marginal utility for each asset, and then proceed with the optimization algorithm.

To illustrate this procedure, we consider a simple case with only two assets, X and Y. A generalization to higher dimensions is straightforward. For returns r_x and r_y respectively, the gross return on a portfolio of these two assets with weights w_x and $w_y = 1 - w_x$ is given by

$$R^{(p)} = w_x \exp(r_x) + w_y \exp(r_y).$$

We approximate the integral (the expected utility)

$$\iint_{-\infty}^{\infty} u(R^{(p)}) f(r_x, r_y)\, dr_x\, dr_y$$

by Monte-Carlo integration. Let $(r_{x,1}, r_{y,1}), \ldots, (r_{x,n}, r_{y,n})$ be draws from the (estimated) joint return distribution. We calculate the expected utility as

$$E(u) = \frac{1}{n} \sum_{i=1}^{n} u\big(w_x \exp(r_{x,i}) + w_y \exp(r_{y,i})\big).$$

By the law of large numbers, this approximation has a considerable degree of accuracy for large n. The marginal expected utility of asset X with respect to weight w_x is given by

$$MEU_X = \frac{1}{n}\sum_{i=1}^{n} u'\left(w_x \exp\left(r_{x,i}\right) + w_y \exp\left(r_{y,i}\right)\right)\exp\left(r_{x,i}\right)$$

and similarly for Y. These expressions can then be used for the optimization algorithm.

5. Illustrative Portfolio

The first step in the formulation of the portfolio, naturally, is selection of assets. The only criterion used in determining the set of assets listed below was to include a diverse set of assets from different industrial sectors for practicality and ease of illustrating changes in composition. Since this portfolio is only for illustrative purposes, we have selected blue-chips from the American stock markets rather than SME scrips. Some of these scrips are listed on NASDAQ while others are listed on NYSE. The data is obtained from Thomson Reuters Eikon. All time series run from 1 July 2011 to 30 June 2016, the number of trading days is 1258. Table 1 shows the stocks and some descriptive statistics of their returns distributions. Table 2 shows the correlations (lower triangle) and rank correlations (upper triangle) between the returns.

Table 1. Descriptive statistics of the assets included in the illustrative portfolio

Asset	mean (ann.)	std. dev (ann.)	1st order autocorr	skewness
Apple (AAPL)	0.1339	0.2718	0.017	−0.3837
Bank of Am. (BAC)	0.0359	0.3742	−0.059	−0.5231
Boeing (BA)	0.1120	0.2390	0.022	−0.5214
Coca Cola (KO)	0.0574	0.1543	−0.058	−0.3258
General Electric (GE)	0.0991	0.2117	−0.005	0.2092
Microsoft (MSFT)	0.1356	0.2436	0.004	−0.1865

The historical price data we used to configure the model ranged from 1 July 2011 to 30 June 2016 (1258 observations). Apple and Microsoft are traded at NASDAQ, the other stocks at NYSE.

In equation (6), the risk aversion parameter c can vary from 0 to ∞ but a reasonable range for this parameter is above 1, which indicates risk aversion, where marginal utility decreases with increasing returns. If $c = 1$ the investor is risk neutral, values below 1 indicate risk affinity and are deemed unrealistic. We restricted the range of c from 1.1 to 15. In order to ensure the existence of expected utility in the presence of heavy tails in the return distributions, we set the shift parameter $b = -0.01$.

We consider two different optimization approaches. First, we derive the optimal portfolio for the unconditional joint return distribution. For a given utility function, this portfolio is optimal if the weights are kept constant. Of course, the optimal portfolio weights depend on the risk attitude of the investors. The more risk averse they are, the higher the weight of low volatility assets, even if that implies a lower expected rate of return.

Table 2. Correlations between the assets

Assets	AAPL	BAC	BA	KO	GE	MSFT
AAPL	1.0000	0.2663	0.2368	0.2040	0.3276	0.4099
BAC	0.2486	1.0000	0.3371	0.3114	0.4782	0.3334
BA	0.2295	0.3494	1.0000	0.2966	0.3995	0.2895
KO	0.2193	0.2862	0.3272	1.0000	0.4244	0.2998
GE	0.3346	0.5342	0.4424	0.4309	1.0000	0.4105
MSFT	0.3927	0.3203	0.3183	0.3408	0.4363	1.0000

The lower triangle shows the correlation coefficients, the upper triangle the rank correlations.

Figure 4 shows the changes in portfolio allocation as the risk aversion parameter is increased from (almost) risk neutrality, c = 1.1, to high risk aversion, c = 15, keeping the shift parameter b = −0.01 constant. The weights were restricted to the interval [−0.25,1.25] (indicated by the dashed lines), hence a moderate amount of short selling is allowed.

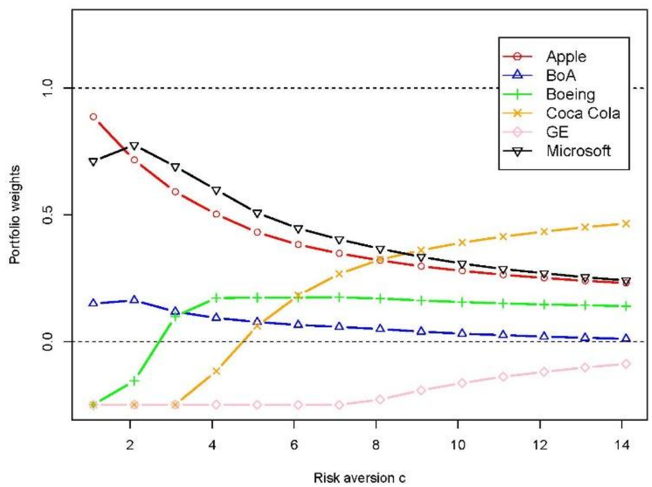

Figure 4. Changes in portfolio allocation, as the risk aversion parameter c changes. The portfolio weights are optimized for the unconditional distribution.

Second, the optimal weights are re-calculated each period, taking into consideration that recent return volatility bears information about current riskiness. If one particular asset becomes relatively more risky its weight will be reduced. Hence, rather than maximizing the unconditional expected utility $\mathrm{E}\left(u\left(R_t^{(p)}\right)\right)$ we maximize expected utility conditional on past returns. Conveniently, the estimated conditional variances $\hat{\sigma}_t^2$, i.e., the estimates of equation (2), are a by-product of the GARCH(1,1) estimation.

Figure 5 depicts the evolution of the optimal portfolio weights for an investor with risk aversion c = 5. It is evident that the weights change substantially and apparently erratically over time. All six assets are shorted a couple of times. Note that transactions costs have been neglected when calculating the optimal weights.

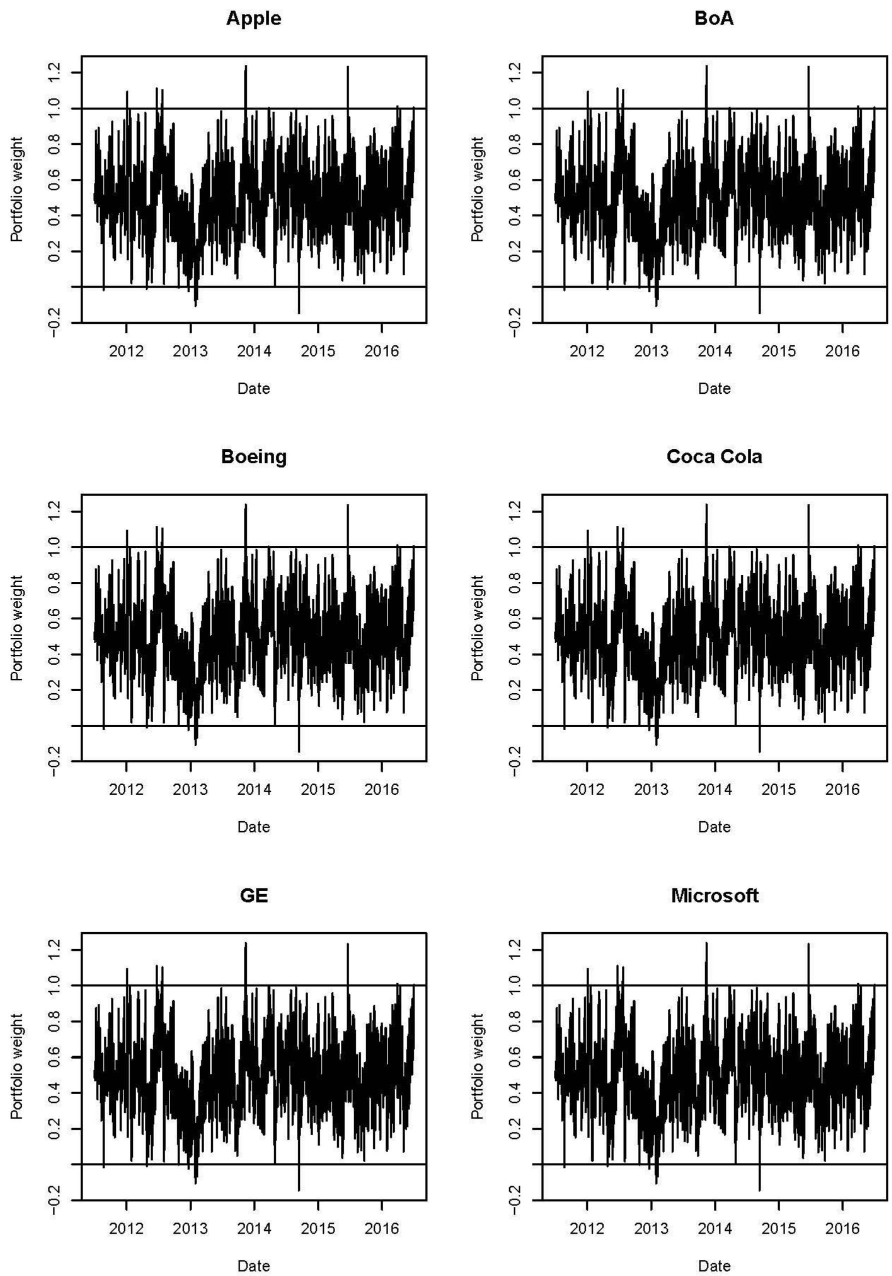

Figure 5. Time-varying utility maximizing portfolio weights for a risk aversion parameter $c = 5$ (and $b = 0$). In each period, the sum of the weights is 1. The weights are restricted to the interval $[-0.25, 1.25]$.

We compare the performance of the two optimization approaches (unconditional and conditional) to each other as well as to a simple "$1/n$" portfolio composition with equal weights for all assets. A natural metric for a portfolio p is its certainty equivalent return $\overline{CER}^{(p)}$, implicitly defined by

$$u\left(\overline{CER}^{(p)}\right) = E\left(u\left(R^{(p)}\right)\right).$$

Hence, the $\overline{CER}^{(p)}$ is the non-random constant gross return that generates the same utility as the actual random gross return distribution $R^{(p)}$. In case of HARA utility we obtain

$$\overline{CER}^{(p)} = \left[E((R^{(p)} - b)^{1-c})\right]^{1/(1-c)} + b.$$

The certainty equivalent return can be computed and meaningfully compared for different portfolio compositions. Table 3 reports the certainty equivalent returns for the optimally weighted unconditional and conditional portfolios as well as a simple $1/n$ portfolio rule. The calculations for the unconditional portfolio and the $1/n$ portfolio weighting schemes are based on n = 50000 simulated returns. Thus, the theoretical certainty equivalent returns are approximated very accurately (due to the law of large numbers). The gain of optimal unconditional weighting over $1/n$ weights is considerable: the net *CER* is more than 45% larger.

When the comparison is based not on the theoretical expectations (as approximated by a large number of simulation runs) but on the rather short historical return series the picture is different. Here, simple $1/n$ weighting would have performed better than the unconditional weighting scheme. Of course, over a relatively short period of time, random fluctuations have a large impact.

The bottom line in table 3 reports the *CER* of the conditional weighting scheme. It would have performed much better than the other two. Over the observation period of 5 years, it would have generated a net *CER* more than three times larger than $1/n$ weights. We conclude that re-weighting the portfolio optimally in each period has a substantial impact on an investor's expected utility.

Table 3. Certainty equivalent net returns for three portfolio weighting schemes: unconditionally optimal, conditionally optimal, and $1/n$ weighting

Portfolio weighting scheme	Return sample	Certainty equivalent return (net)
unconditional	simulated (n = 50000)	0.0005478
$1/n$	simulated (n = 50000)	0.0003759
unconditional	historical (n = 1258)	0.0001587
$1/n$	historical (n = 1258)	0.0001951
conditional	historical (n = 1258)	0.0006081

6. Conclusion

From the example we have illustrated in the previous section, it is evident that this model is not only capable of rendering simulations that retain the most essential statistical traits of the data used to configure the model, but also that the proposed optimization framework can cater to a broad range of risk preferences. Critically, the framework offers substantial advantages over contemporary portfolio optimization procedures, and the improvements the simulation procedure offers over the model configured only using historical data proves that using the filtered simulation mechanism may substantially improve predictive efficiency in practice.

Of course, determining the parameter c, which is crucial to accurately capture investor risk preferences, is by itself a nontrivial procedure, and therefore this model does not claim to be the ultimate source of information on allocation strategy, and it is important to couple it with an adequately precise preference determination methodology. The model is then capable of aiding the decision process by providing an accurate and efficient framework which can offer substantial monetary benefits in practice.

There is also scope to further enhance the model, as improvements motivated by advances in both the estimation methods we have deployed and alternative portfolio optimization frameworks are devised over time.

References

[1] Avramov, Doron, and Guofu Zhou (2010). "Bayesian portfolio analysis", *Annu. Rev. Financ. Econ.*, 2(1): 25–47.

[2] Bollerslev, Tim (1986). "Generalized autoregressive conditional heteroskedasticity", *Journal of Econometrics*, 31(3): 307–327.

[3] Breidt, F. Jay, Nuno Crato, and Pedro de Lima (1998). "The detection and estimation of long memory in stochastic volatility", *Journal of Econometrics*, 83(1-2): 325–348.

[4] Buccola, S.T., and B. C. French (1978). *Estimating Exponential Utility Functions*. U.S. Department of Agriculture, Economics, Statistics, and Cooperatives Service.

[5] Church, Christ (2012). "The asymmetric t-copula with individual degrees of freedom", Oxford, UK: *University of Oxford Master Thesis*, 2012.

[6] Cont, Rama (2007). *Volatility Clustering in Financial Markets: Empirical Facts and Agent Based Models*, pages 289–309. Springer, Berlin, Heidelberg.

[7] Engle, Robert F. (1982). "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation", *Econometrica*, 50(4): 987–1007.

[8] Fama, Eugene F., and Kenneth R. French (1993). "Common risk factors in the returns on stocks and bonds", *Journal of Financial Economics*, 33(1): 3–56.

[9] Garlappi, Lorenzo, Raman Uppal, and Tan Wang (2005). "Portfolio selection with parameter and model uncertainty: A multi-prior approach". *Sauder School of Business Working Paper*, 2005.

[10] Jansons, Vladimirs, Konstantins Kozlovskis, and Natalja Lace (2005). "Portfolio modeling using the theory of copula in Latvian and American equity market". *Proceedings 19th European Conference on Modeling and Simulation*, 2005.

[11] Jondeau, Eric, and Michael Rockinger (2006). "The copula-GARCH model of conditional dependencies: An international stock market application", *Journal of International Money and Finance*, 25(5): 827–853.

[12] Kakouris, Iakovos and Berc Rustem (2014). "Robust portfolio optimization with copulas", *European Journal of Operational Research*, 235(1): 28–37.

[13] Krokhmal, Pavlo, Jonas Palmquist, and Stanislav Uryasev (2002). "Portfolio optimization with conditional value-at-risk objective and constraints", *Journal of Risk*, 4(2): 43–68.

[14] Lai, Tze Leung, Haipeng Xing, and Zehao Chen(2011). "Mean-variance portfolio optimization when means and covariances are unknown", *Annals of Applied Statistics*, 5(2A): 798–823.

[15] Markowitz, Harry (1952). "Portfolio selection", *Journal of Finance*, 7(1):77–91.

[16] Sharpe, William F. (2007). "Expected utility asset allocation", *Financial Analysts Journal*, 63(5): 18–30.

[17] Sklar A. (1959). "Fonctions de repartition a *n* dimensions et leurs marges", *Publ. Inst. Statis. Univ. Paris*, 8: 229–231.