Value-at-Risk of coffee portfolios

S. Gugushvili J. Nowotarski C. W. Oosterlee
L. Ortiz-Garcia E. Verbitskiy*

Abstract

Coffee is the second most traded commodity in the world, and the coffee market can be very volatile even over very short periods of time. Nedcoffee was looking for better ways to access Value-at-Risk of their portfolio. Additional difficulty stems from the fact that portfolio comprises of contracts with delivery dates in January, March, May, July, September and November of each year. Our proposed solution was evaluated using historical market data.

KEYWORDS: coffee, market volatility, value-at-risk, modeling

1 Introduction

With sales volume equal to 106000 metric tons in the fiscal year 2012, Nedcoffee is a major coffee trader with headquarters located in Amsterdam, from which it trades and controls all its green coffee from its sourcing companies in Africa, Asia and South America. Coffee market is a volatile market with traders managing highly complex portfolios. A problem of paramount importance in risk management in general and for Nedcoffee in particular is estimation of a profit and loss distribution of a portfolio over a specified time horizon and the associated risk measures. Value-at-Risk (VaR) has become an important measure for estimating and managing portfolio market risk; see Jorion (2007) for a detailed exposition. VaR is defined as a certain quantile of the change in value of a portfolio during a specified holding period. While the basic concept of VaR is simple, many complications can arise in practical use. Of these part are statistical: VaR is not an absolute, but a model-dependent quantity. Choosing a right probability distribution for an adequate description of the profit and loss distribution is thus of great importance. However, models will typically depend on parameters, which have to be inferred from the data and uncertainty in which will propagate to estimates of VaR. A further complication is that when determining VaR, one is estimating a quantile far in the tail of the distribution, which is a notoriously difficult statistical task. There is also a problem of a different, conceptual nature, that is inherent in the definition of VaR: it is an incoherent risk measure, which in non-technical terms means that a diversified portfolio might have a higher VaR than

*Corresponding author.

prior to diversification, the fact which traders will be reluctant to accept; see Artner et al. (1999) for details. Despite this, the use of VaR is extremely widespread in practice.

Since trading decisions at Nedcoffee are to a considerable extent determined by VaR considerations, the company is greatly interested in 1) constructing better models to be used in VaR computations than currently used by the company, and 2) given a model, using statistically efficient tools for the actual computation of VaR.

The rest of the report is organized as follows: in Section 2 we recall the definition of VaR and introduce some notation. In Section 3 we briefly review the approach employed by Nedcoffee and indicate its shortcomings. Sections 4 and 5 outline some alternatives and present small scale simulation study results for one of them. Section 6 concludes with an outlook and some future work.

2 VaR

Assume a portfolio consists of positions in k different assets, and let $N_i(t)$ and $P_i(t)$ denote respectively the number of contracts and the price of one contract in the ith position, $i = 1, \ldots, k$, at time t. The price of portfolio at time t is then

$$S(t) = \sum_{i=1}^{k} N_i(t) P_i(t).$$

Let Δt be the the holding period of the portfolio, so that the portfolio composition remains constant over the time period $[t, t + \Delta t]$, i.e. $N_i(t) = N_i(t+1)$. The value of the portfolio at time $t + \Delta t$ is $S(t + \Delta t)$. The change in the portfolio value during the holding period is

$$\Delta S = S(t + \Delta t) - S(t) = \sum_{i=1}^{k} N_i(t) \Delta P_i(t) = \sum_{i=1}^{k} N_i(t) R_i(t) = \langle \mathbf{N}(t), \mathbf{R}(t) \rangle, \qquad (1)$$

where $\langle \mathbf{N}(t), \mathbf{R}(t) \rangle$ is the scalar product of vectors $\mathbf{N}(t)$ and $\mathbf{R}(t)$ with components $N_i(t)$'s and $R_i(t)$'s, respectively. The VaR_α risk measure, associated with a given level $0 < \alpha < 1$, is defined by the relation

$$\mathbb{P}(\Delta S < -\mathrm{VaR}_\alpha | \mathbf{N}(t)) = \alpha. \qquad (2)$$

Thus

$$\mathrm{VaR}_\alpha = |F^{-1}(\alpha)|,$$

where F is the distribution function of $\Delta S(t)$ given $N(t)$. In practice Δt typically ranges from one day to two weeks and $\alpha \leq 0.05$, often $\alpha = 0.01$.

In the case of Nedcoffee the portfolio consists of various types of coffee futures and some options written on them. In this paper we will for simplicity assume that the Nedcoffee portfolio consists of futures only. Ignoring options, in principle k can be as large as 10, which corresponds to two major coffee species, Arabica and Robusta, and five possible contract listings per species. Nedcoffee is primarily interested in estimating the 1-day VaR of their portfolio, so that $\Delta t = 1$ day. The level α they aim at is somewhat unrealistically set at $\alpha = 0.015$. In our argumentation we will use a general α.

3 Nedcoffee approach

From formulae (1) and (2) it is obvious that VaR depends on the choice of the model for the futures price process $P(t) = (P_1(t), \dots, P_k(t))$. A number of possibilities are available here.

NEDCOFFEE employ currently an empirical formula for Value-at-Risk at confidence level $\alpha = 0.9985$ (0.15 percent) based on the assumption that underlying Arabica and Robusta coffee prices are separately normally distributed with constant covariance matrices over a period of 3 months (≈ 60 trading days). We are not going to discuss precise details of the method, but will use the NEDCOFFEE VaR estimate for benchmarking purposes.

Given the assumption of Gaussian distribution of prices holds, covariance matrices can be easily estimated using the available historical data on futures prices and then the VaR can be determined in a straightforward fashion. However, except for simplicity of computation, there is little empirical justification for assumptions made in this case. As an illustration of this, we produced a normal Q-Q plot based on returns of the 2-month futures from 15 November 1993 to 7 February 1994, which gives us in total 59 data points. Strong deviation from normality is visible in the plot. Also a formal test for normality, the Shapiro-Wilk test, performed on the same dataset yields the p-value equal to 0.002392, which is very strong evidence against the null hypothesis that the data originate from a certain normal distribution. We would reject the null hypothesis at level 0.05.

4 Possible alternatives

The results from the previous section indicate that one has to look for alternative models and VaR computation methods than those currently employed by Nedcoffee. Two natural options are: a continuous-time model, in which P is a solution to a (multidimensional) stochastic differential equation (SDE), or a time series model, such as a (multivariate) GARCH model. Models based on SDEs are attractive due to the fact that under suitable assumptions a streamlined theory for pricing financial derivatives (e.g. options) is available for them. Furthermore, they are capable of reproducing the mean reversion property one often sees in asset prices (this is achieved through appropriately choosing the drift coefficient of the equation), as well as fitting a wide range of return distributions (this is achieved by selecting a right diffusion coefficient, or by using a general Lévy process instead of the Brownian motion as a driving process of the equation). On the other hand a very fine level of detail provided by sample paths of

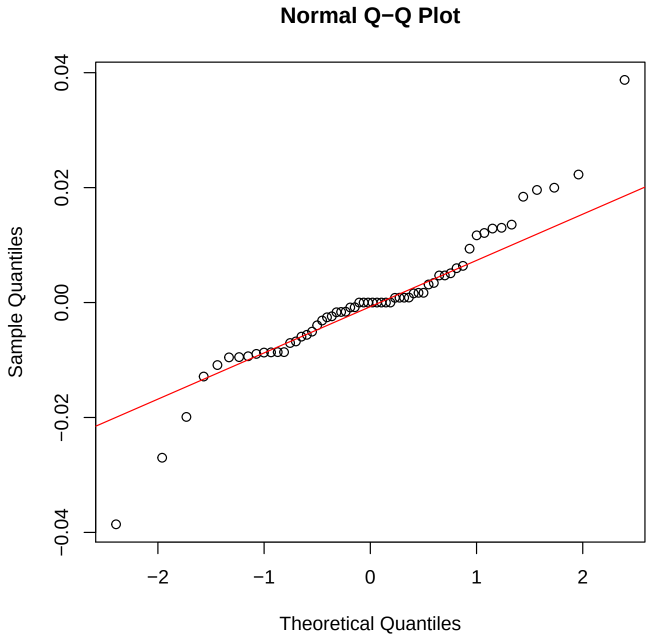

Figure 1: Normal Q-Q plot for the returns based on the 2-month futures data (15/11/1993 to 07/02/1994).

SDE models is not always warranted to be seen in actual financial time series; see e.g. Carr et al. (2002). When asset prices are observed at high frequency, microstructure noise becomes a problem. Moreover, parameter estimation in SDE models, especially in the high-dimensional case when both the dimension of the system of SDEs and of the parameter space are large, is computationally and statistically a very difficult task, unless one restricts attention to simple, but often not realistic models, such as e.g. the Black-Scholes model. In the case of the Black-Scholes asset price dynamics, provided the model parameters have been accurately estimated, VaR can be efficiently computed following the method described in Ortiz-Garcia and Oosterlee (2013) (an extra technical complication in our case would be the fact that we are dealing with futures prices). We refer to the same paper for additional references. As far as the time series models are concerned, multivariate generealisations of traditional univariate models are far from trivial due to the fact that the multivariate character of the model greatly increases the number of parameters required for its description, while a drastic cut of the number of parameters due to parsimony considerations might well render the model inadequate for data description purposes; see e.g. Silvennoinen and Teräsvirta (2009). GARCH process is not an only option here; one can e.g. also consider the AR processes (either the classical or the semiparametric ones), but the same remarks apply.

5 Present approach

Below we propose an approach to VaR computation that in our opinion strikes a good balance between being computationally easy and still beter than the one currently employed by Nedcoffee.

Assume that the underlying asset prices are jointly normally distributed with a constant covariance matrix over a period of 3 months. In this case, the returns $\mathbf{R}(t)$'s have normal distribution with unknown variance and the standard VaR estimate can be applied.

Time series analysis of the returns suggests that the Student's t-distribution is a better fit than the normal distribution. The standard Student distribution is given by the density

$$f(x) = \frac{1}{\sqrt{\nu\pi}} \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})} \left(1 + \frac{x^2}{\nu}\right)^{-\frac{\nu+1}{2}},$$

where ν is the number of degrees of freedom (shape parameter). For our purposes the non-standardized Student's t-distribution with the density

$$f(x|\mu,\sigma,\nu) = \frac{1}{\sigma\sqrt{\nu\pi}} \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})} \left(1 + \frac{\left(\frac{x-\mu}{\sigma}\right)^2}{\nu}\right)^{-\frac{\nu+1}{2}},$$

where μ is the mean (location), σ is the scale parameter, ν is the the number of degrees of freedom (shape parameter).

Note that for $\nu \to \infty$, the the non-standardized Student's distribution $F_{\mu,\sigma,\nu}$ approaches the normal distribution $\mathcal{N}(\mu, \sigma^2)$. Therefore, the application of Student's t-distribution has an advantage over the normal law when the distribution of the returns has lighter tails, and can recover the normal law, when necessary.

Again, we assume that $\mu = 0$, the corresponding VaR value is the product of the scale parameter σ and the appropriate quantile of the distribution.

The VaR models described above have been evaluated (backtested) for a number of different fictitious portfolios in the following fashion: for each trading day the corresponding VaR value has been computed, and the number of trading days when the loss exceeded the predicted VaR has been computed. Ideally, the fraction of such days should be close to the chosen confidence level α.

Here are the results for several portfolios and confidence levels α. We have tested two types of portfolios. For the first type, positions in Arabica and Robusta are long, and for the second type, one is long in Robusta, and short in Arabica. This choice corresponds to test the performance of NEDCOFFEE's estimator, since it is constructed differently for long/long and long/short portfolios. Secondly, we tested the length of the past period used in estimation of the parameters $M = 60$ days and $M = 90$ days. Finally, we performed backtesting for confidence levels $\alpha = 0.15, 1.5,$ and 5 percent.

Tables below give the performance of the estimators, most accurate in bold.

N=[3000,2000,100, 2000, 600, 200], M=60

Confidence Level	NEDCOFFEE	GAUSS	STUDENT
0.15	0.33	0.66	**0.25**
1.50	0.66	2.40	**1.74**
5.00	2.23	**5.21**	5.87

N=[3000,2000,100, 2000, 600, 200], M=90

Confidence Level	NEDCOFFEE	GAUSS	STUDENT
0.15	0.34	0.51	**0.17**
1.50	0.76	2.46	**1.61**
5.00	2.12	**5.09**	5.85

N=[3000,2000,100, -2000, 600, 200], M=60

Confidence Level	NEDCOFFEE	GAUSS	STUDENT
0.15	0.83	0.58	**0.41**
1.50	2.15	1.65	**1.49**
5.00	5.77	4.58	**5.00**

N=[3000,2000,100, -2000, 600, 200], M=90

Confidence Level	NEDCOFFEE	GAUSS	STUDENT
0.15	0.76	0.59	**0.42**
1.50	2.29	**1.53**	1.44
5.00	5.77	4.58	**5.00**

6 Conclusions and outlook

The current VaR estimator for low α severely overestimates the true VaR for long/long portfolios, and underestimates the VaR for mixed portfolios. Suggested extensions demonstrate better performance. In particular, Student's t-distribution offers significant improvement. We also have found that Nedcoffee should consider various levels of confidence, e.g., 1 or 5 percent. The current level of 0.15 percent seems too small to provide accurate risk assessment.

In this report we primarily discussed construction of VaR estimators based on univariate time series $S(t)$ or $\Delta S(t)$. Multivariate modelling of the returns $\boldsymbol{R}(t)$ might provide a better insight into the dynamics of underlying assets. Moreover, multivariate modelling opens a possibility for portfolios optimisation. Another important direction for future work is incorporation of options in the analysis similar to the one performed above.

References

P. Artzner, F. Delbaen, J.M. Eber, D. Heath (1999). Coherent measures of risk. *Mathematical Finance*, 9:203–228, 1999.

P. Carr, H. Geman, D. B. Madan and M. Yor (2002). The fine structure of asset returns: an empirical investigation. *Journal of Business*, 75:305–332.

P. Jorion (2007). *Value at risk: the new benchmark for managing financial risk*. McGraw-Hill, 3rd edition.

L. Ortiz-Garcia and C. W. Oosterlee (2013). Efficient VaR and expected shortfall computations for non-linear portfolios within the delta-gamma approach. Preprint.

A. Silvennoinen and T. Teräsvirta (2009). Multivariate GARCH Models. In T. G. Andersen, R. A. Davis, J. P. Kreiß and Th. V. Mikosch (Eds.), *Handbook of Financial Time Series*, Springer-Verlag, Berlin-Heidelberg, 2009, pages 201–229.