



SOPHIS



INP Grenoble – ENSIMAG

École Nationale Supérieure d'Informatique et de Mathématiques Appliquées de
Grenoble

Rapport de projet de fin d'études

Effectué chez Sophis

Value-at-Risk for commodity portfolios

Martin Ottenwaelter

3e année – Option MIF

17 mars 2008 – 12 septembre 2008

Sophis

24-26 place de la Madeleine
75008 Paris

Responsables de stage

Gregory Cobena et Arnaud Seguin

Tuteur de l'école

Olivier Taramasco

Contents

1	Introduction: Value at Risk, a single number risk indicator	4
2	VaR definition	6
2.1	Introduction	6
2.2	VaR definition	6
2.3	Time horizon	7
2.4	Backtesting	7
2.5	Conclusion	8
3	Historical VaR	9
3.1	Introduction	9
3.2	Method	9
3.3	Experiments	9
3.4	Historical VaR in RISQUE	13
3.5	Conclusion	14
4	1st order VaR	15
4.1	Introduction	15
4.2	The 1 st order approximation formula	15
4.3	A simple example	16
4.4	Comparison of historical and 1 st order VaR	18
4.5	RISQUE's answer to the particularities of commodities	20
4.6	Variance-covariance interpolation	22
4.7	Experimental results	24
4.8	Conclusion	28

5	Principal Component Analysis: an attempt to understand the variance-covariance matrix	29
5.1	Introduction	29
5.2	The PCA methodology	29
5.3	Understanding the variance-covariance matrix	31
5.4	PCA applied to 1 st order VaR	34
5.5	Conclusion	35
6	Conclusion	36
A	Mathematical formulas	37
A.1	Daily returns	37
A.2	An unbiased estimate of the covariance, skewness and kurtosis	37
A.3	Delta, delta-cash	37
B	Risque	38
B.1	VaR computation steps	38
C	Historical VaR Server	42

1 Introduction: Value at Risk, a single number risk indicator

Managing market risk is now an integral part of the financial world. Today's most widely used tool to measure and control market risk was introduced and popularised in 1994 by J.P. MORGAN'S RISKMETRICS software and is called Value-at-Risk –henceforth VaR. VaR “is an attempt to provide a single number summarising the total risk in a portfolio of financial assets” ([Hul06]). Its simplicity in both its understanding and computation makes it such a widely used technique.

There are three methods to compute VaR:

- historical
- analytical (or parametric)
- Monte-Carlo

SOPHIS RISQUE, the portfolio and risk management solution, already features the three methods. But in this paper, we only focus on the historical and analytical methods in the particular case of commodity portfolios. We have three clear goals: (i) understand and explain the current computation methods, (ii) validate those methods and (iii) propose improvements; by keeping our focus on commodity portfolios that have certain particularities, the most important being that commodities are mostly traded with future contracts.

In the first section, we start by defining VaR and give the mathematical formula which is the foundation of any VaR method. Then, we present the backtesting method that is used in the next sections to validate the VaR models. Backtesting consists in verifying, using past records, that the portfolio losses in worst cases is well represented by the VaR. We also introduce another important aspect known as the *time horizon formula*, giving a relationship between a 1-day and a N -day VaR. In other words, it gives a relation between a very short-term risk and a longer medium-term risk.

In the second section, we start by defining the historical VaR method which consists in applying changes to the portfolio based on real past events. We then validate the method by performing a backtest on a set of sample portfolios. We also study a concrete case in which a large amount of data is artificially created for electricity contracts using the available historical data. Then, using the same historical data, we study the time horizon formula. Our experiments show that one must be careful with that formula since even pure lognormal data does not fully validate it.

Historical VaR has the advantage of being very similar to real world past events but requires many costly simulations. For this reason, the parametric VaR is an alternative. Not only it reduces the calculation time and the number of simulations, but it proposes a model of the portfolio risk. The parametric VaR is typically a 1st or 2nd order VaR, which corresponds to an analysis of the 1st or 2nd order derivatives

of the portfolio with respect to the various risk sources: spot values, volatilities, interest rates, credit spreads, etc.

In the third section, we start by computing the 1st order approximation VaR formula. After showing an example on how it is computed, both by hand and in `RISQUE`, we once again validate the method with a series of backtests. The method is then compared with the historical VaR and it is concluded that the 1st order VaR always underestimates the historical VaR, which is due to the skewed and leptokurtic historical portfolios distributions. We then jump into the particularities of commodities and propose a subtle modification to `RISQUE`'s variance-covariance matrix *mapping* method, but find that the initial method is better.

In the fourth section, we pursue the work on the variance-covariance matrix but choose to study it from a much different angle. We apply the statistical method known as Principal Component Analysis in an attempt to better understand the matrix. Although the method gives good results for interest rates, we cannot clearly make any conclusions regarding commodity portfolios with PCA.

In conclusion, we find that there is an interesting parallel between the historical VaR which is a data-hungry method and the *mapping* and PCA techniques that, on the contrary, try to trim existing data in order to keep only the meaningful characteristics. The first method requires so much data that we must sometimes find workarounds to generate the data, while the second is a model-based approach that tries to extract the important variables.

2 VaR definition

2.1 Introduction

In this section, we will first define VaR. We will then see the so-called and frequently used *time horizon formula* that gives the expression of the N -day VaR in function of the 1-day VaR. We will finally see backtesting, a method that tests the validity of a VaR computation. Backtesting will then be used in the following sections to validate both historical and analytical VaR.

2.2 VaR definition

Let V be the VaR of a portfolio, V is defined by the statement ([Hul06]):

We are $\alpha\%$ certain that we will not loose more than V dollars in the next N days.

V is function of two variables:

- N , the time horizon in days
- α , the confidence level

Let ΔP be the change in the value of the portfolio over the next N days. VaR is therefore the loss corresponding to the $(100 - \alpha)$ th percentile of the distribution of ΔP . We can either express it with the following formula:

$$P(\Delta P < \text{VaR}) = 1 - \alpha$$

or the following graph:




Figure 1: VaR as the $(100 - \alpha)$ th percentile of the portfolio's PnL

A VaR estimate is often computed on a daily basis, and for a time horizon of usually 1 or 10 days. The confidence level is usually taken equal to 95% or 99%.

2.3 Time horizon

When the daily changes in the value of the portfolio have identical normal distributions with mean zero, the relation between the 1-day VaR and the N-day VaR is exactly:

$$N\text{-day VaR} = 1\text{-day VaR} \times \sqrt{N} \quad (1)$$

Experiments to test this formula are conducted in section (3.3.2).

2.4 Backtesting

Once a method for computing VaR is chosen, its accuracy needs to be assessed. Backtesting is a common method for testing the validity of a VaR method that consists in looking at the past and counting the number of times the losses of a portfolio were greater than the VaR. In theory, if the computed VaR has a confidence level of $\alpha = 99\%$, then there should be around 1% of outliers. In order to make a serious conclusion out of the number of outliers, we need to do more than just comparing it to 1%. In the next paragraphs, we will see a statistical test that allows us to make that serious conclusion.




Figure 2: Backtesting historical VaR for a portfolio consisting of futures of Brent crude: overlay of the daily changes in the value of the portfolio with the computed VaR. There are 6 outliers out of 550 points (1.1%).

For each day i in history, we compute $\text{VaR}_i(\alpha)$, the VaR using the method we want to backtest; and $\Delta_i P = P_{i+1} - P_i$, the actual change in the value in the portfolio. Next, X_i is defined as:

$$X_i = \begin{cases} 1 & \text{if } \Delta_i P < \text{VaR}_i(\alpha) \\ 0 & \text{else} \end{cases}$$

$(X_i)_{1 \leq i \leq n}$ is a sequence of Bernoulli trials. If they are assumed independent, the random variable $X = \sum_{i=1}^n X_i$ follows a Binomial distribution with parameters (n, α) . If $\tilde{\alpha} = \sum X_i / n$ is an estimator of α , we can perform the statistical test $H_0 : \tilde{\alpha} \leq \alpha$ against $H_1 : \tilde{\alpha} > \alpha$ and see when hypothesis H_0 is rejected. [Gau07] provides an asymptotic approximation for the critical region of the test:

$$W = \left\{ \frac{x - n\alpha}{\sqrt{n\alpha(1-\alpha)}} > u_{2a} \right\}$$

where $x = \sum x_i$ is the number of outliers and a is the significance of the test.

In our tests, we will take $a = 95\%$, and perform the backtest on $n = 250$ days like the Basel control methodology. We will also use the R command `binom.test` that gives the exact confidence interval, although it does not directly give the critical region. For different numbers of outliers, a level of confidence $a = 95\%$ and $n = 250$ simulations, we obtain the following results:

Number of outliers	R confidence interval	$\frac{x - n\alpha}{\sqrt{n\alpha(1-\alpha)}}$	u_{2a}
1	[0.0002, 1.0000]	-0.953	1.645
2	[0.0014, 1.0000]	-0.318	1.645
3	[0.0033, 1.0000]	0.318	1.645
4	[0.0055, 1.0000]	0.953	1.645
5	[0.0079, 1.0000]	1.589	1.645
6	[0.0105, 1.0000]	2.225	1.645

Table 1: Hypothesis test for historical VaR with $a = 95\%$, $n = 250$

We see that for a number of outliers between 1 and 5, 0.01 is in the confidence interval and we are not in the critical region. We can therefore claim with a 95% confidence that if the number of outliers is less or equal to 5, then the VaR method we are testing is valid.

2.5 Conclusion

In this section we defined VaR, we saw the time horizon formula and we dealt with the backtesting method. We did not applied backtesting yet because we first need to see the VaR calculation methods, which is the purpose of the next sections.

3 Historical VaR

3.1 Introduction

In this section we will define historical VaR and detail the computation steps. We will then backtest the VaR on a set of commodity portfolios and perform some experiments with the time horizon formula. We will finally describe how RISQUE computes the historical VaR.

3.2 Method

Historical VaR is a method that uses past data to make the VaR computation. Although there are a few implicit assumptions made when doing historical VaR, they are much less strong than the model assumptions made for parametric VaR.

Suppose that we wish to compute the 99%, 1-day VaR of a portfolio. Historical VaR can be computed if we have a certain amount —500 days, for example— of historical data for every risk source of the portfolio. The algorithm goes as follow:

- for every day i of the historical data:
 - compute the daily return (A.1) between day i and day $i + 1$ of every risk source
 - apply these returns to the current portfolio P_c to obtain a new portfolio value P_i
 - compute $\Delta_i P = P_i - P_c$, the difference between this new value and the current portfolio's value
- estimate the 1% quantile of the $\Delta_i P$ s distribution, this is the historical VaR

Historical VaR is a method that consists in applying past shocks to today's portfolio. The first assumption made here is that what happened in the past is what is likely to happen tomorrow. The second assumption is that the returns are independent in time.

3.3 Experiments

In this paper, theory will need to be tested with real data. We therefore used historical data of future contracts of aluminium, brent crude, copper, WTI and PowerNext ; and conducted experiments on these data sets.

3.3.1 Counting the outliers

The classical historical VaR experimentation consists in plotting the historical VaR on top of the value of the portfolio in order to better visualise the outliers.




Figure 3: Historical VaR for a portfolio consisting of futures of Brent crude: 6 outliers out of 550 points (1.1%)

Experiments are then conducted with every data set. The available historical data is sometimes longer than 250 days. Because the test is more accurate when there is more data, we always perform the backtest with the most possible data.

Commodity	Number of backtests	Outliers (historical)	$\tilde{\alpha}$	Maximum number of outliers to pass
Aluminium	242	2	0.83%	4
Copper	239	2	0.84%	4
Brent Crude	550	6	1.09%	9
WTI	250	1	0.40%	5
PowerNext Base	678	11	1.62%	11

Table 2: Historical VaR results

Every line of the table corresponds to one backtest with a particular portfolio. We see that the historical VaR is validated by the hypothesis test for every commodity.

3.3.2 Verifying the time horizon formula

Another experiment that can be done is the one that attempts to verify the time horizon formula (1) when computing the historical VaR. The first experiment is done on simulated lognormal returns. In that particular case, we know that, in theory, the formula must be confirmed. Daily lognormal returns are simulated using a spot price that follows a geometric brownian motion:

$$S_t = e^{\left(\mu - \frac{\sigma^2}{2}\right)t + \sigma W_t}$$

The returns are therefore equal to:


$$\begin{aligned} r_{1,t} &= \log\left(\frac{S_t}{S_{t-1}}\right) \\ &= \left(\mu - \frac{\sigma^2}{2}\right) + \sigma N(0,1) \end{aligned}$$

where $\mu = 0$, $\sigma = 0.2/\sqrt{252}$ and $N(0,1)$ is a standard normal random number. Then, the 10-days returns are computed by adding consecutive daily returns:


$$\begin{aligned} r_{10,t} &= \log\left(\frac{S_t}{S_{t-10}}\right) \\ &= r_{1,t} + r_{1,t-1} + \dots + r_{1,t-9} \end{aligned}$$

Before computing the VaR, we can compare the density of the daily returns with the density of the 10-days returns divided by $\sqrt{10}$:

```
1 compare_density <- function()
2 {
3   d1 <- density(returns)
4   d10 <- density(returns10)/sqrt(10)
5
6   plot(d1$x, d1$y, type="l")
7   points(d2$x, d2$y, type="l", col="red", lwd=2)
8 }
```



(a) 500 simulations



(b) 5000 simulations

Figure 4: Comparison of the density of simulated 1-day returns (thin) and 10-day returns divided by $\sqrt{10}$ (red)

We observe that, even with simulated lognormal returns, it takes a certain amount of simulations before the formula is significantly validated. When computing VaR, financial institutions rarely use more than 500 days of past data, so we should be careful when using the time horizon formula because not only it is an approximation when the prices don't exactly follow a lognormal distribution, but even with pure lognormal returns, the result is only exact with a high amount of historical data. The same experiment with real historical data gives the following result:




Figure 5: Comparison of the density of simulated 1-day returns (thin), its estimated normal distribution (dashed) and 10-day returns divided by $\sqrt{10}$ (red)

We fit on top of the returns (black) a normal distribution (dashed). We observe that both curves are nearly identical for brent, which means that the returns are nearly lognormal ; and they are quite different for copper, which means that the returns are not lognormally distributed. Then we overlay the 10-day returns divided by $\sqrt{10}$ distribution (red) and observe once again that the curve is nearly identical to the 1-day returns for brent, but much different for copper. The brent result is surprising because we concluded with the previous simulation that even pure lognormal returns don't give such a good result ; but the bad result for copper was predictable because the normal distribution fitting was already bad.

Having these remarks in mind, we can now compare the actual VaR values using the two methods:

Commodity	Error
Simulated lognormal	8.7%
Aluminium	13.3%
Copper	16.2%
Brent Crude	29.6%
WTI	30.0%
PowerNext Base	7.5%

Table 3: Time horizon formula experiments: error made when converting 10-day VaR to 1-day VaR. The converted VaR is compared to the actual 1-day VaR.

The error is pretty high for all the commodities and even for the simulated normal returns. The VaR time horizon formula must therefore be used carefully.

3.3.3 Generating more historical data for a 30 days VaR

In this subsection, we will show how historical data can be generated in order to produce more data necessary to compute a 30-days VaR for an electricity portfolio. Due to lack of time, the results are not presented here, but will be during the oral presentation.

3.4 Historical VaR in RISQUE

Historical VaR is not computed directly within RISQUE, but using a set of software and technologies called “Historical VaR Server”. Historical VaR can be long to compute since all the portfolios need to be recalculated a large number of times, so the tasks need to be parallelised on multiple computers. This is the role of the Historical VaR Server whose different components are the following:

- the *scheduler* and the *manager*, which distribute the tasks
- the *activation service*, that manages the resources
- the *var calculator*, that does the actual computations by communicating with RISQUE

Finally, the results are viewed in a reporting application. Screenshots of the Historical VaR Server are in appendix (C).

Because we want to compare the historical and 1st order VaR methods, we need to treat 1st order VaR first. Therefore, the experiments with historical VaR are conducted in the next section.

3.5 Conclusion

In this section we defined historical VaR and detailed the computation steps. We then saw that the historical VaR on commodity portfolios was validated by the backtest. Experiments with the time horizon formula showed that its accuracy can be limited: even with simulated lognormal returns, the approximation is quite severe. We finally described how `RISQUE` computes the historical VaR, but left the actual computation for the next section, in which it will be compared with the 1st order VaR.

4 1st order VaR

4.1 Introduction

In this section, we will see the 1st order approximation formula for analytical VaR and immediately apply it to a simple example. We will then see more complex examples and compare the results with the historical VaR, as defined in the previous section. Finally, we will try to enhance the *mapping* method used by RISQUE as an answer to the particularities of commodity portfolios.

4.2 The 1st order approximation formula

Let P be the value of a portfolio with n risk sources which are typically spot prices, volatilities, rates or CDS spreads, for example. Let δ_i be the delta-cash sensitivity (A.3) of the portfolio to the i th asset ($1 \leq i \leq n$). If r_i is the return (as defined in A.1) on asset i in 1 day, then the dollar change of our investment sensitive to asset i is $\delta_i r_i$. It follows that a linear approximation of the variation of the value of the portfolio is

$$\Delta P = \sum_{i=1}^n \delta_i r_i \quad (2)$$

Recall that the VaR for an α confidence level and a N days horizon is defined by

$$\int_{-\infty}^{\text{VaR}} \delta P(x) dx = 1 - \alpha$$

If we assume the r_i follow a centered normal distribution, then δP is normally distributed and

$$\text{VaR} = \sigma_P \cdot \mathcal{N}^{-1}(1 - \alpha) \cdot \sqrt{N/252}$$

where

- \mathcal{N}^{-1} is the inverse standard normal distribution function,
- $\sigma_P = \sqrt{\Delta^\top \Sigma \Delta}$ is the portfolio's volatility where Σ is the joint distribution r_1, \dots, r_n 's variance-covariance matrix and $\Delta = (\delta_i) \in \mathcal{M}_{n,1}$.

In conclusion,

$$\text{VaR} = \sqrt{\Delta^\top \Sigma \Delta} \cdot \mathcal{N}^{-1}(1 - \alpha) \cdot \sqrt{N/252} \quad (3)$$

Remarks:

- we did two approximations here: a first order approximation on ΔP and the assumption that the returns are normally distributed.
- formula (3) is a direct proof of the time horizon formula (1).

4.3 A simple example

4.3.1 Computation by hand

Consider a portfolio consisting of \$10 million in shares of Apple and \$7 million in shares of Microsoft, i.e. $\Delta = (10 \ 7)^\top$. We suppose that the variance-covariance matrix of the pair (Apple, Microsoft) is

$$\Sigma = \begin{pmatrix} 0.4 & -0.2 \\ -0.2 & 0.1 \end{pmatrix}$$

The 99% 1-day VaR is therefore equal to

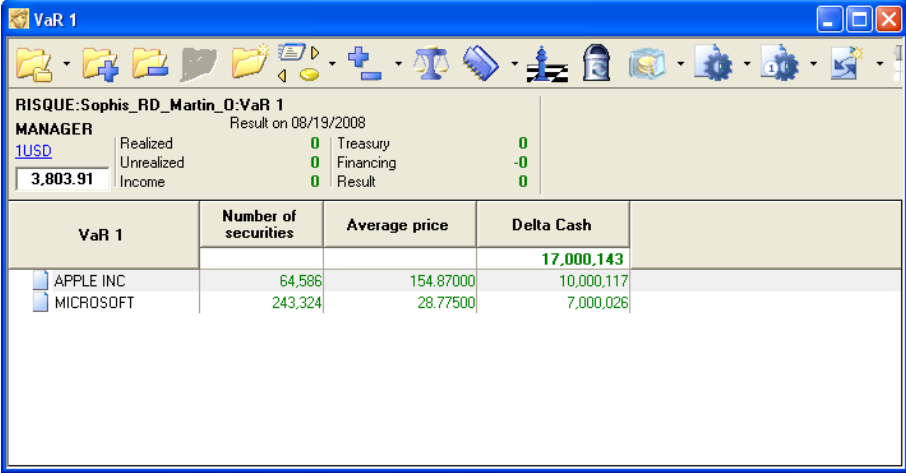
$$\begin{aligned} \text{VaR} &= \sqrt{10^2 \times 0.4^2 + 7^2 \times 0.1^2 + 2 \times 10 \times 7 \times (-0.2) \times 0.4 \times 0.1} \\ &\quad \times \mathcal{N}^{-1}(0.01) \times \sqrt{1/365} \\ &= -\$0.477381 \text{ million} \end{aligned}$$

In conclusion, we are 99% certain that we will not loose more than \$477,381 by tomorrow.

4.3.2 Computation in RISQUE

We now want to confirm this result by making the same computation in RISQUE:

1. we create a new portfolio and book the two deals



The screenshot shows the 'VaR 1' dialog box in the RISQUE software. The title bar is 'VaR 1'. The main window has a menu bar and a toolbar. Below the toolbar, there's a section for 'RISQUE:Sophis_RD_Martin_0:VaR 1' with a 'MANAGER' field set to '1USD'. To the right, it says 'Result on 08/19/2008'. Below this, there's a table with 4 columns: 'Realized', 'Unrealized', 'Income', and 'Result'. The values are: Realized: 0, Unrealized: 0, Income: 0, Result: 0. To the right of this table, there's another table with 2 columns: 'Treasury' and 'Financing'. The values are: Treasury: 0, Financing: -0. Below these tables, there's a table with 4 columns: 'VaR 1', 'Number of securities', 'Average price', and 'Delta Cash'. The values are: VaR 1: 3,803.91, Number of securities: 17,000.143, Average price: 10,000.117, Delta Cash: 7,000.026. Below this table, there's a table with 4 columns: 'APPLE INC', 'MICROSOFT', 'Number of securities', 'Average price', and 'Delta Cash'. The values are: APPLE INC: 64,586, 154.87000, 10,000.117; MICROSOFT: 243,324, 28.77500, 7,000.026.

RISQUE:Sophis_RD_Martin_0:VaR 1	
MANAGER	
Result on 08/19/2008	
Realized	0
Unrealized	0
Income	0
Treasury	0
Financing	-0
Result	0

VaR 1	Number of securities	Average price	Delta Cash
3,803.91	17,000.143	10,000.117	7,000.026

APPLE INC	MICROSOFT	Number of securities	Average price	Delta Cash
64,586	243,324	154.87000	28.77500	10,000.117

2. we call the *Parametric VaR* dialog and input the parameters. The variance-covariance matrix has been previously saved in a .csv file.

Parametric VaR

Time horizon: 1d

Confidence level: 99.00 (from 80% to 99.99%)

Correlation matrix: C:\Documents and Settings\motenwaelter\Desktop\ ...

Calculation type: Delta

Extraction criterion:

Consolidation: XXX

Entry point: RISQUE:Sophis_RD_Martin_0:VaR 1 107419

Breakdown for Rho: XXX

Scenario configuration:

Evaluate Correl. matrix Delta vector Covar matrix Gamma matrix Cancel

3. the result is displayed in a new window

VaR results for VaR 1

+

+

+

RISQUE	VaR	Delta VaR	Vega VaR	Rho
Sophis_RD_Martin_0	477,497.24	477,497.24		
VaR 1				

< >

The VaR is equal to \$477,497, which is nearly equal to the previous computation. The difference comes from the fact that we didn't book exactly \$10 and \$7 millions.

4.4 Comparison of historical and 1st order VaR

Now that both historical VaR and 1st order VaR methods have been seen, experiments can be conducted parallelly to compare them. In the following section, we use the fact that 1st order VaR can be easily deduced from historical VaR. If the variance-covariance matrix is computed with the exact same data used for historical VaR, then the 1st order VaR is equal to the 1% quantile of the normal distribution that fits the historical distribution. Therefore, we plot, on top of the historical distribution, the normal distribution that fits and can visually compare both computations.

The first series of tests has been conducted on simple portfolios consisting in future contracts. The tests were conducted with the R software and using data previously extracted from a client's database.




Figure 6: Overlay of historical and 1st order VaR for a portfolio of brent futures

Figure (4.4) is very representative of all the commodity portfolios that we have tested: the historical distribution is slightly skewed and pretty much leptokurtic. Therefore, the historical VaR is higher than the normal VaR. It shows that it is important to take into account the kurtosis and skewness of commodity portfolios so that the VaR is not underestimated.

The second series of tests has been conducted on real life portfolios. They can contain more complex products like exotic options or hybrid products.




Figure 7: Overlay of historical and 1st order VaR for a portfolio of soft commodities

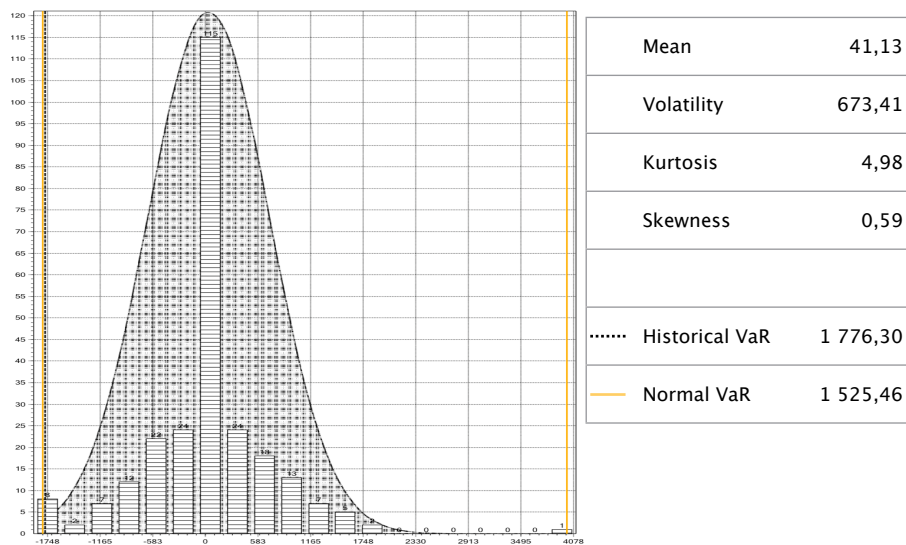


Figure 8: Overlay of historical and 1st order VaR for a portfolio of commodity hybrid products




Figure 9: Overlay of historical and 1st order VaR for a portfolio whose underlyings are metals

We see that the results on complex portfolios are pretty much the same that we obtained on simple portfolios. The distributions are similarly leptokurtic and the 1st order VaR underestimates the historical VaR.

4.5 RISQUE's answer to the particularities of commodities

Computing VaR for a portfolio whose risk sources are shares is a straightforward process:

- compute Σ , the variance-covariance matrix of the returns on the spot price of the shares. For example, historical data can be used to compute an estimator of the variance-covariance – see (A),
- compute the VaR with formula (3).

Unfortunately, the process is not as simple for commodity portfolios because commodities don't have a spot price, but are instead traded with futures contracts. Most commodities have at least a future contract expiring on every month of the year, some every day and electricity has base and peak futures! There can be a lot of different futures for the same commodity and, consequently, a huge variance-covariance matrix needs to be computed. This matrix is not easy to compute for various reasons:

1. historical data is not always available for every future – for example, the client can't store that much data;

2. some historical data is not reliable enough to compute the variance – for example, because there are so many futures, some have a too low trading volume for the price to be meaningful;
3. it is better to look at the historical data of futures with the same time to maturity rather than at the plain historical data. This is done by a technique called *future rolling*;
4. in some cases, finding the future with the same time to maturity is not sufficient: the futures should also be with a similar seasonality.

Contract Name	Expiry Date	Settlement Price (USD)
Jul08	2008-06-13	132.11
Aug08	2008-07-16	131.87
Sep08	2008-08-14	132.13
Q3 08	2008-06-13	131.86
Cal 09	2008-12-16	131.77

Table 4: A few out of the many ICE Brent futures available on 2008-05-23

Points 3 and 4 are perhaps what is the most complex about commodities: the *historic correlations* module of RISQUE doesn't support rolling. Instead, users can input a variance-covariance matrix that was computed by another method that does rolling.

RISQUE addresses problems 1 and 2 with variance interpolation: the variances that can't be computed are not computed but are instead interpolated. This method is also named Δ -mapping because you can do it the other way around:

- a set “standard” maturities is chosen, for which the variance-covariance matrix is known. These maturities are chosen because their prices are believed to guide the nearby futures' prices.
- a new Δ' is computed as follows: for each element Δ_i of Δ , Δ_i is splitted on to its adjacents standard maturities' Δ'_k and Δ'_{k+1} proportionally to the time interval between the maturities.
- VaR is computed with Δ' and the standard maturities' variance-covariance matrix.

Let's apply the method on a simple portfolio consisting of 3 futures expiring respectively in 30, 60 and 120 days. We suppose we want to map $\Delta = (2 \ 12 \ 3)^\top$ to the standard maturity set (30 days, 120 days). We are looking for a new $\Delta' = (\Delta'_{30} \ \Delta'_{120})^\top$. The algorithm goes as follow:

- $\Delta' = (0 \ 0)^\top$;

- 2 maps entirely to Δ'_{30} thus $\Delta' = (2 \ 0)^\top$;
- 12 is proportionnaly splitted on to Δ'_{30} and Δ'_{120} . Because it corresponds to a 60 days maturity, $12 \times (120 - 60)/(120 - 30) = 8$ is mapped to Δ'_{30} and $12 \times (60 - 30)/(120 - 30) = 4$ is mapped to Δ'_{120} . At the end of this step, $\Delta' = (10 \ 4)^\top$;
- 3 maps entirely to Δ'_{30} , thus $\Delta' = (10 \ 7)^\top$.

If we suppose that the variance-covariance matrix of the standard maturity futures is equal to:

$$\Sigma = \begin{pmatrix} 0.4 & -0.2 \\ -0.2 & 0.1 \end{pmatrix}$$

the VaR can be computed with formula (3), $\text{VaR} = \$477,000$.

4.6 Variance-covariance interpolation

As said previously, Δ -mapping can be done the other way around and be interpreted as variance-covariance interpolation. This subsection describes this method.

In this section, the variance of X will be written $\text{var}(X) = \sigma_x^2$ and the covariance between X and Y will be written:

$$\text{cov}(X, Y) = \sigma_{xy} = \sigma_x \sigma_y \rho_{xy}$$

Thus, we also have $\text{var}(X) = \sigma_{xx}$.

4.6.1 The method

Variance-covariance interpolation is a two steps method:

- select a small subset of the risk sources for which the variance-covariance matrix is known
- reconstruct the initial variance-covariance matrix by interpolating between the values of the subset's variance-covariance matrix

4.6.2 An example

Let (F_1, F_2, F_3) be the risk sources of a portfolio and σ_{ij} the covariance of the historical returns of F_i and F_j . We have

$$\Sigma = \begin{pmatrix} \sigma_{11} & \sigma_{12} & \sigma_{13} \\ \sigma_{21} & \sigma_{22} & \sigma_{23} \\ \sigma_{31} & \sigma_{32} & \sigma_{33} \end{pmatrix}$$

We choose (F_1, F_3) as the subset, the values $\sigma_{11}, \sigma_{13}, \sigma_{31}, \sigma_{33}$ are known. We need to fill the gaps in Σ' the new variance-covariance matrix:

$$\Sigma' = \begin{pmatrix} \sigma_{11} & \sigma_{13} \\ \sigma_{31} & \sigma_{33} \end{pmatrix}$$

If we assume that the time between two expiry dates of two adjacents futures is 1 month, bilinear interpolation gives:

$$\Sigma' = \begin{pmatrix} \sigma_{11} & \frac{1}{2}(\sigma_{11} + \sigma_{13}) & \sigma_{13} \\ \frac{1}{2}(\sigma_{11} + \sigma_{31}) & \frac{1}{4}(\sigma_{11} + \sigma_{13} + \sigma_{31} + \sigma_{33}) & \frac{1}{2}(\sigma_{33} + \sigma_{33}) \\ \sigma_{31} & \frac{1}{2}(\sigma_{31} + \sigma_{33}) & \sigma_{33} \end{pmatrix}$$

4.6.3 Interpolation coefficients

Method in RISQUE Let us formalise the previous example: the interpolation coefficients are computed using the expiry dates of the futures. Let F_x be the future for which the covariance needs to be interpolated and t_x its expiry date. There are 3 cases:

- if there exists F_a, F_b such as $t_a < t_x < t_b$, then

$$\left\{ \begin{array}{l} \sigma_{ax} = \frac{t_b - t_x}{t_b - t_a} \sigma_{aa} + \frac{t_x - t_a}{t_b - t_a} \sigma_{ab} \\ \sigma_{bx} = \frac{t_b - t_x}{t_b - t_a} \sigma_{ab} + \frac{t_x - t_a}{t_b - t_a} \sigma_{bb} \\ \sigma_{xx} = \frac{t_b - t_x}{t_b - t_a} \sigma_{ax} + \frac{t_x - t_a}{t_b - t_a} \sigma_{bx} \\ \quad = \left(\frac{t_b - t_x}{t_b - t_a} \right)^2 \sigma_{aa} + 2 \left(\frac{t_b - t_x}{t_b - t_a} \right) \left(\frac{t_x - t_a}{t_b - t_a} \right) \sigma_{ab} + \left(\frac{t_x - t_a}{t_b - t_a} \right)^2 \sigma_{bb} \end{array} \right.$$

- otherwise there exists F_a such as $t_x < t_a$, then $\sigma_{ax} = \sigma_{xx} = \sigma_{aa}$,
- or there exists F_b such as $t_b < t_x$, then $\sigma_{bx} = \sigma_{xx} = \sigma_{bb}$.

Giving more importance to the variances In this paragraph, we will follow [Ale01] that suggests that we should rather interpolate the variances instead of the covariance. We consider the case where there exists F_a, F_b such as $t_a < t_x < t_b$. We are looking for a new interpolation coefficient λ such that $F_x = \lambda F_a + (1 - \lambda) F_b$. In this case,

$$\sigma_{xx} = \lambda^2 \sigma_{aa} + (1 - \lambda)^2 \sigma_{bb} + 2\lambda(1 - \lambda) \sigma_{ab}$$

Because we want to give importance to the variances rather than to the covariances, we suppose that the variance σ_{xx} is a linear combination of the variances σ_{aa} and σ_{bb} :

$$\sigma_{xx} = \frac{t_b - t_x}{t_b - t_a} \sigma_{aa} + \frac{t_x - t_a}{t_b - t_a} \sigma_{bb}$$

λ is found by picking the positive solution of the quadratic equation


$$\lambda^2 \sigma_{aa} + (1 - \lambda)^2 \sigma_{bb} + 2\lambda(1 - \lambda) \sigma_{ab} = \frac{t_b - t_x}{t_b - t_a} \sigma_{aa} + \frac{t_x - t_a}{t_b - t_a} \sigma_{bb}$$

Once lambda is found, the previous method can be applied with the new interpolation coefficient.


4.7 Experimental results

4.7.1 Graphical interpretation of interpolation


The first experiment consists in plotting in three dimensions the brent's variance-covariance matrix and its two interpolated variants.



(a) Original variance-covariance matrix



(b) Interpolated variance-covariance matrix



(c) Interpolated variance-covariance matrix by giving more importance to the variance

Figure 10: Variance-covariance interpolation for Brent Crude.

We clearly see the effect of the interpolation between the two first graphs. The difference between the two last graphs is more subtle, but we can still see that the difference is on the diagonal. It is normal because the difference between the two mapping methods resides in the way the variance is interpolated.

We also see that the quality of the result depends heavily on the choice of the *reference futures*, the ones on which the mapping is done. This particular brent variance-covariance matrix has such a particular shape that a poor choice in the reference futures could result in an important error. This choice is therefore left to the user that knows which maturities are known to guide the market. For example, for aluminium, the most traded future are the 3, 15 and 27 months.

4.7.2 Comparison of VaR computations

Now, the interesting comparison is the one between the standard VaR computation and the two mapping computations. The comparisons are performed on brent, copper and PowerNext base electricity. VaR is computed using randomly generated Δ vectors and the error displayed is an average on a high number of experiments.

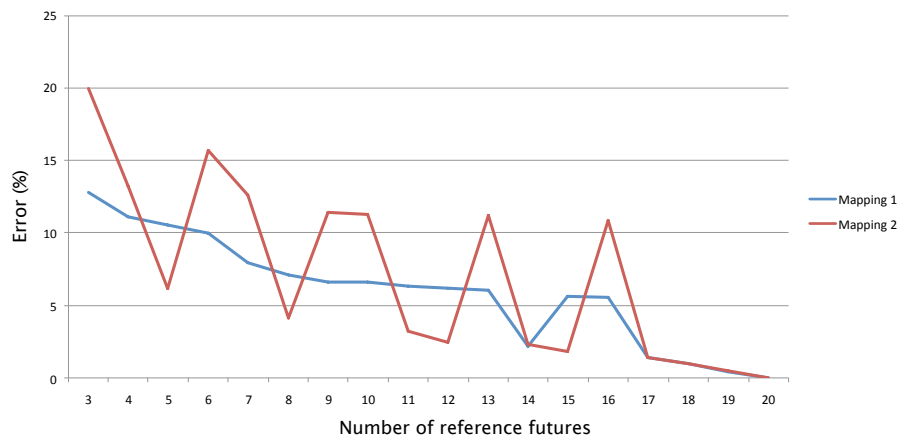


Figure 11: Error made when computing a brent portfolio VaR with the two mapping methods.

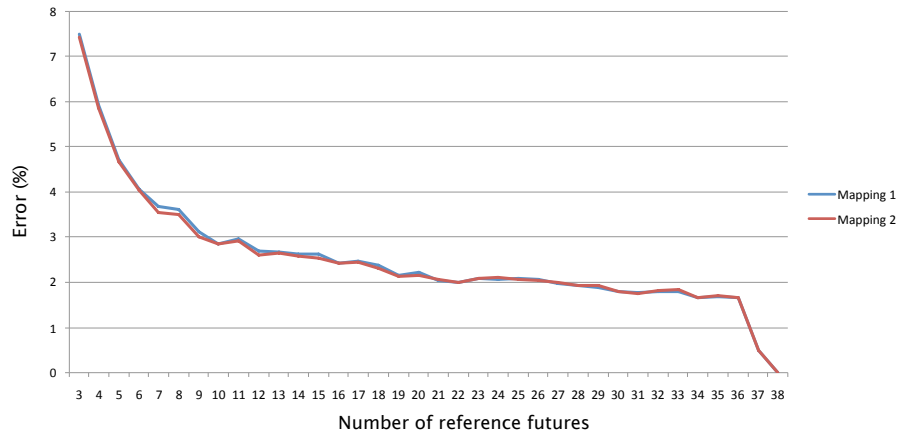


Figure 12: Error made when computing a copper portfolio VaR with the two mapping methods.

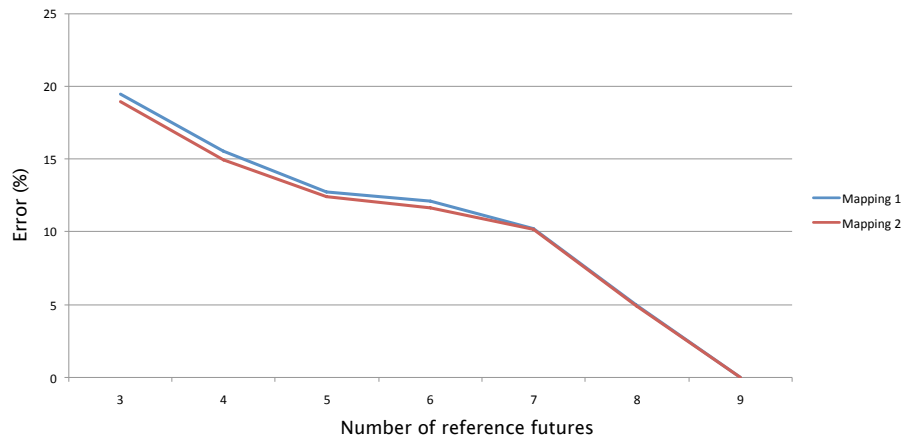


Figure 13: Error made when computing an electricity portfolio VaR with the two mapping methods.

The first remark that we can make is that the more *reference futures* there are, the less important is the error, which is the expected behaviour. The second remark is that the second mapping method does not give a significant increase in the accuracy of the method for copper and electricity, and gives a very instable result for brent, which is probably due to the specific shape of brent's variance-covariance matrix. We therefore conclude that the second mapping technique does not improve the accuracy of mapping and is unstable.

4.8 Conclusion

In this section, we applied the 1st order VaR to commodity portfolios and compared the results with historical VaR. We observed that our portfolios had skewed and leptokurtic distributions, resulting in the 1st order VaR being inferior to the historical VaR. Risk managers must therefore be aware that 1st order VaR underestimates the actual VaR for such portfolios. Finally, we saw the *mapping* method that deals with the particularities of commodity portfolios. We proposed an enhancement to this method but did not find it conclusive. The next section continues on with the particularities of commodities by attempting to understand better the variance-covariance matrix.

5 Principal Component Analysis: an attempt to understand the variance-covariance matrix

5.1 Introduction

We have seen in the previous section that the complexity of computing 1st order VaR resides in the computation of the variance-covariance matrix. In this section, we attack the variance-covariance matrix with a different angle, in an attempt to understand it better. To do so, we applied a method known as Principal Component Analysis (PCA).

5.2 The PCA methodology

5.2.1 Introduction

“PCA is a rotation of axes in multi-dimensional space that allows one to find linear combinations –the principal components– of the original variables that summarise as much of the information as possible. The eigenvalue of each component is an estimate of the amount of total variance explained by that particular component. Inspecting the eigenvalues allows one to pick the minimum number of components summarising enough of the total variance, while inspecting the eigenvectors leads to financial interpretations of the principal components.” ([PV02])

In the current context, the goal is similar to what we did with mapping: reduce the amount of initial data by only selecting data that has a meaning. In the case of mapping, there was a simple and natural choice of meaningful data: the most traded maturities. In this case, we are looking for more complex combinations.

5.2.2 Data

Let $X = (x_{ij}) \in \mathcal{M}_{m,n}$ be a matrix of data. In the case of financial data, x_{ij} is the i th historical return on the j th asset.

Before performing PCA, data must be normalised. Let \bar{X}_j and $\sigma(X_j)$ respectively be the empirical average and empirical volatility of the j th column. If \tilde{X} is the normalised matrix, then:

$$\tilde{x}_{ij} = \frac{x_{ij} - \bar{X}_j}{\sigma(X_j)} \quad (4)$$

5.2.3 Variance-covariance matrix

Let $\Sigma \in \mathcal{M}_{n,n}$ be the variance-covariance matrix of X , it is also the correlations matrix of \tilde{X} :

$$\Sigma = \frac{1}{m} \tilde{X}^\top \tilde{X}$$

PCA consists in diagonalising Σ and sort the eigenvalues and eigenvectors decreasingly:

$$\Sigma = W D W^\top$$

where $D = \text{diag}(\lambda_1, \dots, \lambda_n)$ is the diagonal eigenvalue matrix with $\lambda_1 > \dots > \lambda_n$, and $W = (W_j) \in \mathcal{M}_{n,n}$ the eigenvectors matrix where W_j is λ_j 's eigenvector.

5.2.4 Principal components

The k th principal component is defined as:

$$P_k = w_{1k} \tilde{X}_1 + w_{2k} \tilde{X}_2 + \dots + w_{nk} \tilde{X}_n$$

and the principal components matrix $P \in \mathcal{M}_{m,n}$ is therefore equal to:

$$P = \tilde{X} W$$

Conversely, $\tilde{X} = P W^\top$, thus

$$\tilde{X}_j = w_{j1} P_1 + \dots + w_{jn} P_n \quad (5)$$

5.2.5 Dimension reduction

Our goal is to reduce the number of terms in the sum in (5). To that purpose, we compute P 's variance-covariance matrix:

$$\begin{aligned} \Gamma &= \frac{1}{m} P^\top P \\ &= \frac{1}{m} W^\top \tilde{X}^\top \tilde{X} W \\ &= \frac{1}{m} W^\top (m W D W^\top) W \\ &= D \end{aligned}$$

We deduce that the j th principal component is equal to λ_j . Since the eigenvalues are sorted ascendingly, $\alpha = \sum_{l=1}^k \lambda_l / n$ is the proportion of variance contained in the first k principal components.

A minimum value of α is chosen depending on the initial data set, so that k can be deduced. k is the number of principal components that is kept. The principal components representation of \tilde{X} is therefore equal to:

$$\tilde{X}_j \approx w_{j1} P_1 + \dots + w_{jk} P_k \quad (6)$$

Once the w_{ij} are computed, the size of the problem has moved from n to k .

5.2.6 Recipe

PCA can be summarised with the following *recipe*:

- normalise the data
- diagonalise Σ , the variance-covariance matrix
- for $1 \leq k \leq n$, plot the sum of the k first eigenvalues. Each sum is a percentage that tells us how much the first k principal components are responsible for the variance of the initial data.
- plot the first few eigenvectors to study their shape
- compute Σ_k , an estimation of Σ using only k principal components

5.3 Understanding the variance-covariance matrix

5.3.1 Theory

PCA is known to work well with interest rates. An example is given by [\[Fry97\]](#): PCA is applied to 1,543 daily observations of US Treasury rates between 1989 and 1995. There are 10 series that have a maturity between 3 months and 30 years. The sums of the first eigenvalues show that the first three principal components are responsible for more than 95% of the variance of the data:




Figure 14: Sums of the k first eigenvalues

If 95% is considered to be a reasonable threshold, then we can decide to only keep the first three eigenvectors. The next step is to plot them in order to study their shape and their meaning:





Figure 15: The first three eigenvectors


The first component “corresponds to a roughly parallel shift in the yield curve” ([Hul06]), the second to “a twist or steepening of the yield curve” and the third to “a bowing of the yield curve”. Interest rates seem to follow this pattern, but what about commodities?

5.3.2 Practice

We now apply the PCA methodology to our data sets. Each set contains historical prices for future contracts. The data is first extracted from the database. It is then cleaned because there are always inconsistencies. Finally, we apply the *future rolling* method as described in (3).




(a) Sums of the first k eigenvalues




(b) First three eigenvectors

Figure 16: PCA results for copper




(a) Sums of the first k eigenvalues




(b) First three eigenvectors

Figure 17: PCA results for Brent



(a) Sums of the first k eigenvalues



(b) First three eigenvectors

Figure 18: PCA results for electricity

PCA works remarkably well with copper: the first three principal components account for more than 98% of the variance and their shape follows the “shift, twist, bow” rule of interest rates. But for Brent and electricity, the results are not as good: the first three principal components don’t account for more than 95% of the variance and it is more difficult to make any conclusions regarding the shape of all the eigenvectors. More experiments have been conducted on other commodities and no conclusions can be made on the shape of the eigenvectors.

5.4 PCA applied to 1st order VaR

In the case of 1st order VaR, PCA is applied to the variance-covariance matrix Σ . We want a formula for $\Sigma^* \in \mathcal{M}_{n,n}$, an estimation of the variance-covariance matrix, that uses the principal components obtained by PCA. (4) gives:

$$X_j = \bar{X}_j \mathbf{1} + \sigma(X_j) \tilde{X}_j$$

By injecting (6), we obtain

$$X_j \approx \bar{X}_j \mathbf{1} + \sigma(X_j) (w_{j1} P_1 + \dots + w_{jk} P_k)$$

Therefore, $\Sigma^* = X^\top X$. Because P is orthogonal, we find that:

$$\Sigma_k^* = W^k D W^{k\top}$$

where

$$\begin{cases} W^k = (w_{ij} \sigma(X_i)) \in \mathcal{M}_{n,k} \\ D = \text{diag}(\sigma(P_1), \dots, \sigma(P_k)) \end{cases}$$

Recall that the formula for 1st order VaR is known (3):

$$\text{VaR} = \sqrt{\Delta^\top \Sigma \Delta} \cdot \mathcal{N}^{-1}(1 - \alpha) \cdot \sqrt{N/252}$$

An approximation of the VaR using PCA is given by:

$$\text{VaR}_k = \sqrt{\Delta^\top \Sigma_k^* \Delta} \cdot \mathcal{N}^{-1}(1 - \alpha) \cdot \sqrt{N/252}$$

Note that the new formula for VaR is a function of k , the number of principal components kept, which is a variable that controls the extent of the approximation. The following graph shows the error made when computing VaR with Σ_k^* instead of Σ , in function of k , for a Brent portfolio.




Figure 19: The error made when computing 1st order VaR by applying PCA to a Brent portfolio.

5.5 Conclusion

In this section we saw the Principal Component Analysis method. When PCA works well –with interest rates, for example– we find that the first three principal components (the shift, twist and bow) are more representative of the variance-covariance matrix than the actual historical data. In an attempt to make similar conclusions with commodities, we applied PCA to our data sets. Unfortunately, we did not find as significant regular patterns in the principal components for commodities.

6 Conclusion

In this paper, we studied Value-at-Risk for commodity portfolios. We started by defining VaR. Then we saw the historical and 1st order methods. We concluded with a backtest that these methods are valid. We also compared each of the methods and found out that since our portfolios have skewed and leptokurtic distributions, the 1st order VaR will always underestimate the historical VaR. We compared two *mapping* techniques and concluded that RISQUE's method is already satisfying. Finally, we studied Principal Component Analysis in an attempt to better understand the variance-covariance matrix. But the experiments were not as conclusive for commodities as they are for interest rates.

Appendix

A Mathematical formulas

A.1 Daily returns

Let S_i be the price of a financial asset at the end of day i . The i th daily return is defined by:

$$r_i = \frac{S_i - S_{i-1}}{S_{i-1}}$$

In finance, log-returns are more commonly used:

$$r_i = \log \left(\frac{S_i}{S_{i-1}} \right)$$

A.2 An unbiased estimate of the covariance, skewness and kurtosis

Let $x = (x_i)_{1 \leq i \leq n}$ and $y = (y_i)_{1 \leq i \leq n}$ be two series of historical data. An unbiased estimate of the covariance is:

$$\sigma_{xy} = \frac{1}{n-1} \sum_{k=1}^n (x_k - \bar{x}) (y_k - \bar{y})$$

The k th sample moment of x is defined by:

$$m_k = \sum_{i=1}^n (x_i - \bar{x})^k$$

An unbiased estimate of the skewness of x is:

$$\text{sk}_x = \frac{n^2}{(n-1)(n-2)} m_3$$

and an unbiased estimate of the kurtosis of x is:

$$\text{ku}_x = \frac{n^2}{(n-1)(n-2)(n-3)} \left[(n+1)m_4 - 3(n-1)m_2^2 \right]$$

A.3 Delta, delta-cash

Let P be a portfolio. The *delta* sensitivity to the i th risk source whose price is S_i is defined as:

$$\Delta_i = \frac{\partial P}{\partial S_i}$$

and the similar *delta-cash* sensitivity is equal to:

$$\delta_i = S_i \Delta_i$$

B RISQUE

B.1 VaR computation steps

The screenshot shows the 'Sophis_RD_Martin_O' application window. The top section displays a summary of the portfolio, including the current value of 3,000,000 USD and various risk metrics (Delta, Gamma, Rho, etc.). Below this, a table lists the portfolio holdings:

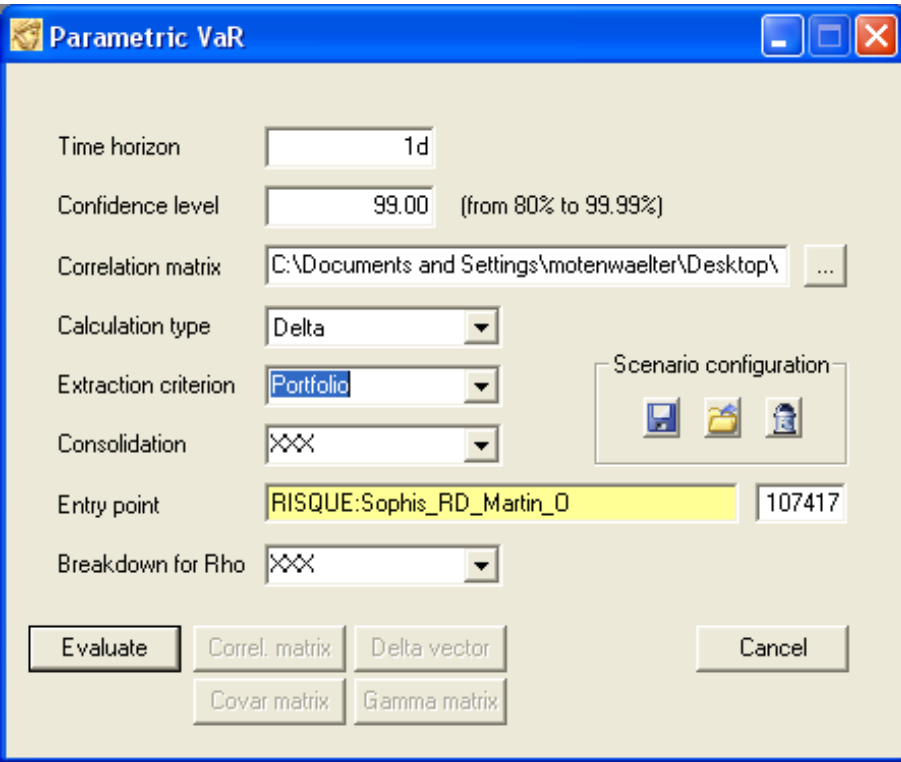
Sophis_RD_Martin_O	Number of securities	Theoretical	Maturity	Delta	Delta Cash	Ident	Inst
BRENT CRUDE - Oct 08	1,000	104,550,000	09/15/2008	2.973	311,606,255	107,417	
BRENT CRUDE - Dec 08	1,000	103,945,000	11/13/2008	1.00000	103,945,000	2,559,136	
BRENT CRUDE - Mar 09	1,000	103,111,252	02/12/2009	1.00000	103,111,252	2,559,137	

Figure 20: A brent portfolio with 3 futures

The screenshot shows the 'Historic correlations for Sophis_RD_Martin_O' dialog box. It contains several sections for configuring the historical correlation analysis:

- History:** Backward day count (30), Latest date included (02/01/2008).
- Output:** Reference (SOPHISREF), Save to file (C:\Documents and Settings\motenw).
- Scenario configuration:** Load, Save, Delete buttons.
- Extraction:** Consolidation (dropdown), Breakdown List (dropdown).
- Table:** A table with columns Name, Reference, Type, and Rate. It lists three Brent crude oil futures contracts: BRENT CRUDE - Oct 08 (LCO V8, S & V), BRENT CRUDE - Dec 08 (LCO Z8, S & V), and BRENT CRUDE - Mar 09 (LCO H9, S & V).
- Buttons:** Update, Remove, Compute, Output, Cancel.

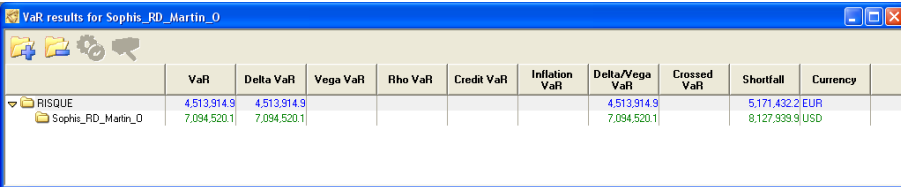
Figure 21: The historical correlations dialog



The dialog box is titled "Parametric VaR". It contains the following fields and controls:

- Time horizon:** 1d
- Confidence level:** 99.00 (from 80% to 99.99%)
- Correlation matrix:** C:\Documents and Settings\motenwaelter\Desktop\ ...
- Calculation type:** Delta
- Extraction criterion:** Portfolio
- Consolidation:** XXX
- Entry point:** RISQUE:Sophis_RD_Martin_0 107417
- Breakdown for Rho:** XXX
- Scenario configuration:** A sub-dialog box with three icons (a folder, a document, and a server).
- Buttons:** Evaluate, Correl. matrix, Delta vector, Covar matrix, Gamma matrix, and Cancel.

Figure 22: The VaR computation dialog



The table shows the results of the VaR computation for the entry point "RISQUE:Sophis_RD_Martin_0".

	VaR	Delta VaR	Vega VaR	Rho VaR	Credit VaR	Inflation VaR	Delta/Vega VaR	Crossed VaR	Shortfall	Currency
RISQUE	4,513,914.9	4,513,914.9					4,513,914.9		5,171,432.2	EUR
Sophis_RD_Martin_0	7,094,520.1	7,094,520.1					7,094,520.1		8,127,939.9	USD

Figure 23: The 1st order VaR result

Standard Option "Mart Option Brent Oct08" (OTHER)

UNDERLYING Calculation Date: 07/22/2008 Reference: LCO V8 Name: BRENT CRUDE - Oct 08	PAYMENT Delivery: Future Payment curr.: USD Quotation: USD Fixing Type: Last Fixing Forex: Calc. Agent:	OPTION Reference: MobOct08 Name: Mart Option Brent Oct08 Pay Off: $\text{Max}(S-K, 0)$ Forward Start Date: Premium Pay Date: Expiration date: 09/10/2008 Deferred payoff date: <input type="checkbox"/> Performance Pay Off Description: Call Exercise Style: European Option Unit: Barrel US Strike Unit: Barrel US Strike: USD 110.00000 Basis: 0 Notional: 0.00 Number of points: 100 B & S: No optima: <input checked="" type="checkbox"/> Validation	RESULTS Value in USD Value in % Delta Gamma Vega Theta Epsilon Equ Dividend Yield Forward Price Break Even Rho / USD
--	--	---	--

SIMULATION

Spot: 104.55000
 Repo (%): 0.0000
 Rate: 2.8273
 Volatility: 30.41
 Overvolatility: 0.00

Market Data Underlying Swaption Param Implied Strike

Figure 24: Option creation dialog

Sophis_RD_Martin_0

RISQUE: Sophis_RD_Martin_0
 MANAGER
 BRENT CRU
 106.07

Result on 07/22/2008

	Realized	Unrealized	Income	Treasury	Financing	Result	Delta	Gamma	Vega	Theta	Epsilon
	0	-1	0	0	0	-1	3,000.345	3,412	0	142	-43
	0	0	0	0	0	0	0	0	0	0	-3.4

	Number of securities	Theoretical	Maturity	Delta	Delta Cash	Gamma	Gamma cash
Sophis_RD_Martin_0				2,974	311,642,333	31m	3,412
BRENT CRUDE - Oct 08	1,000	104.55000	09/15/2008	1.00000	104,550,000		
BRENT CRUDE - Dec 08	1,000	103.94500	11/13/2008	1.00000	103,945,000		
BRENT CRUDE - Mar 09	1,000	103.11125	02/12/2009	1.00000	103,111,252		
Mart Option Brent Oct08	1,000	2.56042	09/10/2008	0.34508	36,078	0.03121	3,412

Figure 25: The same brent portfolio with the just created option

The screenshot shows a window titled "Commodity Delta Parametric VaR: ...". Inside, there is a table with two columns. The first column contains a folder icon and the text "Sophis_RD_Martin_0". The second column is labeled "Delta" and contains the value "7,095,369".

	Delta
Sophis_RD_Martin_0	7,095,369

Figure 26: 1st order VaR

The screenshot shows a window titled "Commodity Delta/Gamma Parametri...". Inside, there is a table with two columns. The first column contains a folder icon and the text "Sophis_RD_Martin_0". The second column is labeled "Delta/Gamma" and contains the value "7,095,740".

	Delta/Gamma
Sophis_RD_Martin_0	7,095,740

Figure 27: 2nd order VaR

C Historical VaR Server

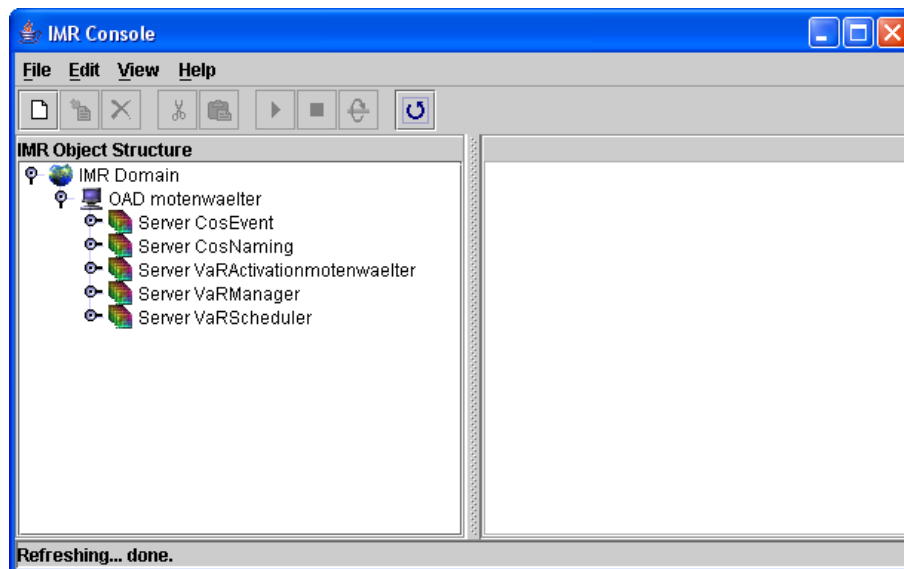


Figure 28: The IMR console

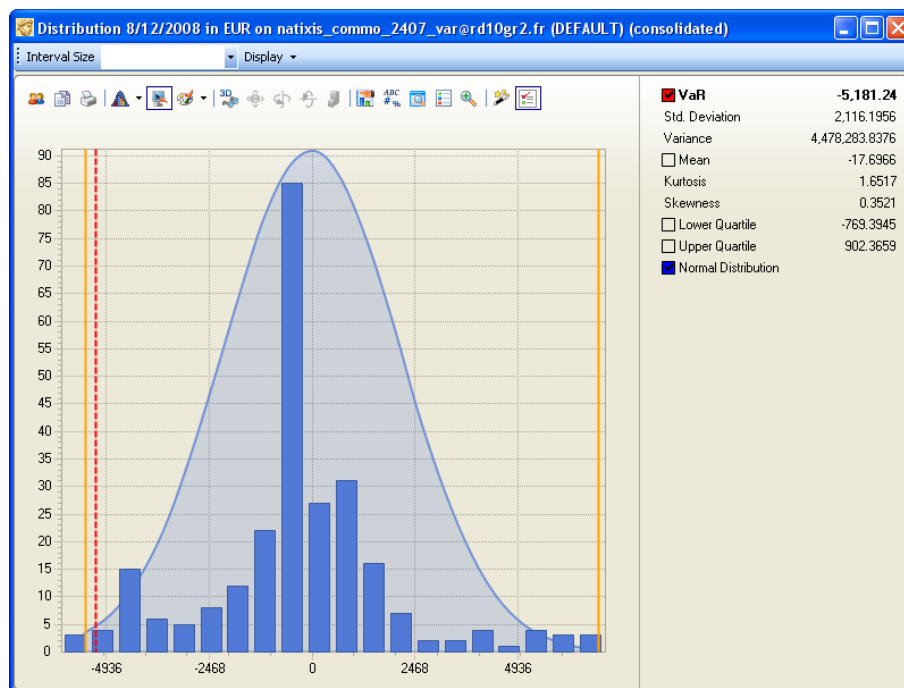


Figure 29: The reporting application

References

- [Ale01] Carol Alexander. *Market Models: A Guide to Financial Data Analysis*. Wiley, 2001.
- [Fry97] Jon Frye. Principals of risk: Finding var through factor-based interest rate scenarios. *Risk Publications*, 1997.
- [Gau07] Olivier Gaudoin. *Principe et méthodes statistiques*, 2007.
- [Hul06] John C. Hull. *Options, Futures and Other Derivatives*. Pearson, 2006.
- [PV02] Christophe Pérignon and Christophe Villa. Components proponents. *Risk*, 2002.