CURRENCY BASKET AS ASSET OR BASE CURRENCY IN VALUE-AT-RISK COMPUTATION

MARC HENRARD

ABSTRACT. This note describes the problem arising from using a currency basket in the computation of value-at-risk. This applies mainly when the basket is used as *base currency*. A solution based on the modification of the historical time series is proposed. The solution is easy to implement and doesn't have important draw-back.

1. Introduction

Value-at-risk (VaR) is a quantile approach to risk. The VaR of a portfolio is the forecast of a given (high) quantile of the distribution of the returns of the portfolio over a given period. In other words it is a monetary estimation of the level beyond which the lost will go only with a given (small) probability.

There are several classical approaches to VaR. One, called the parametric approach, is to suppose that the returns of the risk factors follow a given probability distribution (the normal distribution is often chosen). The parameters of the distribution are estimated from recent market history. A second approach, called historical simulation approach, creates a profit series by applying historical changes in risk factors to today's positions. The quantile is directly estimated from the series.

A currency basket is a unit composed of a fixed number of units of several component currencies. A typical example is the SDR[1], used by several international institutions. The basket is defined as a fixed number of units and not a fixed weight. With the evolution of the exchange rates the weights of the different currencies in the basket change.

Due to this problem, the historical time series of the basket exchange rates cannot be used in the way that they are used for a true currency. This is in particular true for the VaR computation presented in this note but more generally for studies that use historical time series. The computation of the VaR is based on the current value (weight) in the different currencies and the historical time series on a number of units. These two views cannot be used simultaneously when any of the VaR methods is used (parametric or historical).

If we look more specifically at the normal parametric method for VaR, one of its founding hypotheses is that the (log) returns of currencies are normal. But it is impossible to have at the same time the returns of series which are normal and the returns of a linear combination of those series which are also normal (the ratio of a linear combination is not a linear combination of the ratios). The use of both the components and the basket as risk factors is incoherent with one of the hypotheses of the model.

Date: First version: August 2002. This version: January 2003.

Key words and phrases. Currency basket, SDR, value-at-risk, historical series.

Adapted from Chapter 6 of *Value-at-risk: the delta normal approach* [1].

[1]SDR: special drawing right, currently composed of USD 0.577, EUR 0.426, GBP 0.0984 and JPY 21.

This problem appears not only for currency baskets but also for other baskets. In particular stock indices that are based on a number of stocks present the same problem. In this last case the problem is even exacerbated by the change of number of units or the change of constituents of the index.

One can accept this as part of the method and use historical series to compute the variance-covariance matrix. But then if one owns the exact number of units in each of the underlying currencies, the VaR will not be zero. In other words, the VaR will not be zero even if the risk is null. Numerical examples of the size of the residual VaR are given at the end of Section 3.

Here we propose to twist the data slightly to solve the problem and obtain a "clean" risk number. The goal is to obtain zero return if one owns the exact number of units in each of the underlying currencies. A way to do this is explained in Section 4. This solution, that we have implemented on a large scale, is for us the fastest, the cleanest and the more coherent with the way we use VaR.

This is not the only way to compute the VaR with a basket as base currency. Another is, in the delta-normal VaR with log-return framework, to modified the computations in the following way. First the basket position is split in its component currencies. Then a fake position equal to the total value of the portfolio split in the basket components is removed from the portfolio. The VaR is computed in any currency and finally the result obtained is converted in the basket. But this technique is, to our opinion, heavier to implement in an automatic way and presents some problems on reading the results as the value of the portfolio has been changed (it is always 0). The explanation on why this other technique is working and its draw-backs are presented at the end of Section 4.

2. Notation and VaR set-up

The basket is composed of n currencies (four for the SDR) with number of units α_i ($i = 1, \ldots, n$). Let $x_{i,t}$ ($T_0 \leq t \leq T$) be the exchange rates of the different currencies with respect to an (arbitrarily selected) base currency. The exchange rate of the basket with respect to that base currency is

$$X_t = \sum_{i=1}^{n} \alpha_i x_{i,t}.$$

A currency neutral position is to hold α_i units of the underlying currencies and to be short of one basket. So the vector

$$(\alpha_1 x_{1,T}, \ldots, \alpha_n x_{n,T}, -X_T)$$

representing this position converted in the base currency, should have no risk.

The relative return of the currencies and the basket with respect to the base currency are

$$c_{i,t} = \frac{x_{i,t} - x_{i,t-1}}{x_{i,t-1}} \qquad C_t = \frac{X_t - X_{t-1}}{X_{t-1}}.$$

The log returns are given by

$$d_{i,t} = \ln(x_{i,t}/x_{i,t-1}) \qquad D_t = \ln(X_t/X_{t-1}).$$

Let R denote the matrix of the returns (weighted or unweighted) used to compute the historical scenarios or the variance-covariance matrix. In our notations

$$R = \operatorname{diag}(\beta_T, \ldots, \beta_{T_0+1}) \begin{pmatrix} c_{i,T} & \cdots & c_{n,T} & C_T \\ \vdots & \ddots & \vdots & \vdots \\ c_{i,T_0+1} & \cdots & c_{n,T_0+1} & C_{T_0+1} \end{pmatrix}$$

where β is a weighting scheme (for historical scenarios, all the β are equal to 1, see Mina and Yi Xiao [3, Section 2.1] or Hull [2, p. 356, 14.8]).

The Delta-normal VaR of a position p is given, up to a multiplicative factor depending of the probability level of the VaR, by

$$\mathrm{VaR}(p)^2 = p^T(R^T R)p = (Rp)^T Rp.$$

3. Why it does not work

To obtain a VaR equal to 0 in the Delta-Normal framework it is necessary that $Rp = 0$ for the position $p = (\alpha_1 x_{1,T}, \ldots, \alpha_n x_{n,T}, -X_T)$. The same condition applies to have a historical VaR equal to 0 at all probability levels. In the relative return framework[2], this is the case only if for all t

$$\sum_{i=1}^{n} \alpha_i x_{i,T} c_{i,t} - (\sum_{i=1}^{n} \alpha_i x_{i,T}) \frac{\sum \alpha_i c_{i,t} x_{i,t-1}}{\sum \alpha_i x_{i,t-1}} = 0.$$

This would be the case if $x_{i,t-1} = x_{i,T}$, but this is not the case in general.

As the daily returns of the position are not all zero, the implication is that the VaR computed using the historical simulation method would also be non-null.

Note that all the above arguments are valid for an arbitrary base currency. In particular they are valid when the base currency is the currency basket itself. In the following example the SDR is used as base currency.

We have computed the VaR of a portfolio made of the component currencies with the exact number of units to make up SDR 1 billion (USD 577 million, EUR 426 million, GBP 98.4 million and JPY 21 billion). This was done using five years of data ending on 14 June 2002 (for delta-normal VaR, the parameters are computed using EWMA and a decay factor of 0.97). The residual VaR for a 97.5% probability level and a one-day horizon is

Base currency: SDR

Currency	Delta-normal VaR (in SDR)	Historical VaR (in SDR)
EUR	1,740,279	2,739,039
GBP	403,523	845,471
JPY	1,124,301	1,541,932
USD	1,991,825	2,647,140
Total	96,989	1,061,558

Note that due to the exponential weighting, the delta-normal VaR is in practice computed with less than one year of data.

The covariance matrix was given by

Currency	Volatilities	GBP	EUR	JPY	USD
GBP	0.1835	1	-0.0103	-0.3674	0.0155
EUR	0.2847	-0.0103	1	0.0361	-0.8604
JPY	0.4393	-0.3674	0.0361	1	-0.5062
USD	0.2282	0.0155	-0.8604	-0.5067	1

4. How to make it work

Instead of using historical time series of the basket we can rewrite the history using the *current weight* (not the number of units) for the different currencies.

Let w_i be the current weights of the different components:

$$\alpha_i x_{i,T} = w_i X_T.$$

To obtain a clean risk figure, we need a time series that gives the same results with a position in the basket or its decomposition into the component currencies.

[2] In the log return framework, the formulas are a little bit longer but the results are similar.

So for any position $(p_1,\ldots,p_n,p^*)$ with $(p_1,\ldots,p_n)$ the positions in the component currencies[3] and p^* the position in the basket, we should find that in term of risk

$$(p_1,\ldots,p_n,p^*) \equiv (p_1+\alpha_1 p^*,\ldots,p_n+\alpha_n p^*,0).$$

For all t the daily returns of the two positions should be equal. For the relative return this implies that

$$\sum_{i=1}^n p_i x_{i,T} c_{i,t} + p^* X_T C_t = \sum_{i=1}^n (p_i+\alpha_i p^*) x_{i,T} c_{i,t}.$$

Denoting by K_t the modified time series, this is equivalent to

$$\frac{K_t - K_{t-1}}{K_{t-1}} = \sum w_i c_{i,t}.$$

To obtain a clean risk figure, we have to take

$$K_T = X_T \quad \text{and} \quad K_{t-1} = \frac{K_t}{1+\sum w_i c_{i,t}}.$$

For the log return the time series is

$$L_T = X_T \quad \text{and} \quad L_{t-1} = \frac{L_t}{\exp(\sum w_i d_{i,t})}.$$

These series can be used for normal VaR but also for historical VaR. Note that these series have to be recreated every day. They are not cumulative. This means the full history has to be rewritten daily.

The series to be used depends of the way covariance matrix and historical return are computed.

For the computation of the normal VaR, it depends on the type of return used to compute the covariance matrix.

Parametric VaR

Covariance matrix	Asset return	Series
Relative return	normal	K
Log return	normal	L

For the historical VaR, the series to be used depends on the type of return computed and on the way the historical information is used to compute the profit (see for example Mina and Yi Xiao [3, Section 3.1] and Hull [2, Section 14.8]).

Historical VaR

Historical return	Profit computation	Series
Log return	$\text{PnL} = P_0(e^r-1)$	K
Relative return	$\text{PnL} = P_0 r$	K
Log return	$\text{PnL} = P_0 r$	L

A completely coherent choice for the covariance matrix computation, asset return hypothesis in parametric VaR and historical VaR is to choose the relative return, a normal asset return hypothesis and a profit formula of the form $\text{PnL} = P_0 r$ for the historical simulation.

Another set of hypotheses is, however, often chosen. It consists of the log return for the covariance matrix computation, the normal return hypothesis, and a profit formula of the form $\text{PnL} = P_0(e^r-1)$ for the historical profit series. The incoherence here is that the assets are supposed to be log-normal to compute the parameters and normal to compute the parametric VaR. The historical profit computation is coherent with the returns computation. To compute the VaR with SDR as base currency with those hypotheses, one has to use the "L" series for parametric VaR

[3] Some of the units α_i can be zero, so the description we give takes also into consideration the non-component currencies and the non-currency risks

and the "K" series for historical VaR. Even if the choice seems incoherent it is the one done by several implementations.

To the first order the normal and the log-normal hypotheses are equivalent. So if the normal hypothesis is associated to a *delta* approach (first-order approximation of the positions), there is no incoherence.

We have computed the VaR with the example of the previous section. The log-normal rescaling is used for the parametric VaR and the relative return for the historical VaR as described in the previous paragraph.

Base currency: SDR

Currency	Delta-normal VaR (in SDR)	Historical VaR (in SDR)
EUR	1,720,794	2,425,955
GBP	403,301	830,502
JPY	1,120,337	1,653,650
USD	2,022,451	2,689,322
Total	683	602

The small residual VaR comes from rounding errors.

Remark: Suppose we use the log return. If a portfolio does not include the basket and its total value is zero, the change of time series has no impact on its VaR even if the basket is used as base currency. To show this, let $p = (p_1, \ldots, p_n)$ be a portfolio with $\sum p_i = 0$. The daily historical returns of the position are

$$\begin{aligned}
\sum_{i=1}^{n} p_i d_{i,t} &= \sum_{i=1}^{n} p_i \ln\left(\frac{x_{i,t}}{x_{i,t-1}}\right) \\
&= \ln\left(\prod_{i=1}^{n}\left(\frac{x_{i,t}}{x_{i,t-1}}\right)^{p_i}\right) \\
&= \ln\left(\left(\prod_{i=1}^{n-1}\left(\frac{x_{i,t}}{x_{i,t-1}}\right)^{\sum_{j=1}^{i} p_j}\left(\frac{x_{i+1,t}}{x_{i+1,t-1}}\right)^{-\sum_{j=1}^{i} p_j}\right)\left(\frac{x_{n,t}}{x_{n,t-1}}\right)^{\sum_{j=1}^{n} p_j}\right) \\
&= \ln\left(\prod_{i=1}^{n-1}\left(\frac{x_{i,t}/x_{i+1,t}}{x_{i,t-1}/x_{i+1,t-1}}\right)^{\sum_{j=1}^{i} p_j}\right).
\end{aligned}$$

This value depends only of the ratios $x_{i,t}/x_{i+1,t}$ and so does not depend on the base currency and its time series. A trading book that starts with a total value of zero and that borrows to invest is not affected at all by the change of the base currency or by the "adjustment" of the basket time series.

If we use the relative return, a change of base currency is not coherent (a series and the series of the inverse of its elements cannot be normal simultaneously). The computations in the different base currencies and series will be slightly different (the same currency loss will be translated in a different way in the base currency).

In any case, whatever the value of the portfolio is, if there is no basket in the portfolio and the the basket is not used as base currency, the VaR of the portfolio is not affected by the changes.

Another is, in the delta-normal VaR with log-return framework, to modified the computations in the following way. First the basket position is split in its component currencies.

Remark: As said in the introduction, our proposal is not the only possible. We prove here that in the delta-normal set-up with the variance-covariance matrix computed by $R^T R$, the VaR with a basket as base currency can be computed on a modified portfolio with another base currency and then converted. First the basket position is split in its component currencies. Then a fake position equal to the total value of the portfolio split in the basket components is removed from the portfolio. The VaR is computed in any currency and finally the result obtain is converted in the basket.

Let P be the value of the portfolio p ($P = \sum p_i$), $\bar{p} = p - P(w_1, \ldots, w_n)$ and Σ be the matrix in the new base currency. Using the notations of [1], we have

$$\begin{aligned}
\bar{p}\Sigma\bar{p} &= (p - Pw)^T(R^TR - X^TR - R^TX + v_{k,k}1)(p - Pw) \\
&= p^TR^TRp + P^2(w^TR^TRw) - (p - Pw)^T(X^TR + R^TX - v_{k,k}1)(p - Pw) \\
&= p^TR^TRp
\end{aligned}$$

where we use that $w^TR^TRw = 0$ as w is the components weights and the last term is 0 as $\sum(p_i - Pw_i) = 0$. So we have that the VaR (converted in the basket) of the modified portfolio in the new base currency is equal to the VaR of the original portfolio with the basket as base currency.

Note that to obtain the result, we used the change of base currency and explicitly the fact that the matrix is computed with the log-return time series.

This approach gives the same theoretical VaR, but it requires the creation of a fake adjustment position for each position for which one wants to compute the VaR. In particular if the VaR is computed for parts of the portfolio, then the creation of the adjustment for all those parts is necessary. Moreover the value of the adjusted portfolio is by construction always zero. This is a problem, specially if one computed the VaR on a relative (percentage) basis.

5. Conclusions

The computation of parametric or historical VaR in the case where a currency basket is an asset or the base currency by using the historical time series of the basket creates incoherences. This is coming from the definition of currency basket that uses a fixed number of units where the computation of the VaR is base on *current weights*. The use of historical time series will create residual risk even in the absence of real risk. This problem also appears for stock indices.

By modifying the time series of the currency basket using the current weights of the different currencies, the problem can be overcome. Depending of the exact way the data is used (return, log return), the modification is different but uses the same idea: a portfolio with the currency basket or with its decomposition in the component currencies of the basket should have the same risk. The solution proposed here is the only one for which the two portfolios have the same risk.

Using this technique we present an example based on a portfolio of SDR 1 billion. After the modification the residual risk has completely disappeared.

Acknowledgments

The views expressed here are those of the author and not necessarily those of the Bank for International Settlements. The author wishes to thank his colleagues, and in particular W. Gehlen, J.-M. Mahler, and J.-P. Matt, for fruitful discussions and valuable comments on previous versions of this document.

References

[1] M. Henrard. Value-at-risk: the delta-normal approach. In preparation.

[2] J. C. Hull. *Options, futures, and other derivatives.* Prentice Hall, fourth edition, 2000.

[3] J. Mina and J. Yi Xiao. *Return to RiskMetrics: The Evolution of a Standard.* RiskMetrics, 2001.

Current address: Derivatives Group, Banking Department, Bank for International Settlements, CH-4002 Basel (Switzerland)

E-mail address: marc.henrard@bis.org