

Value-at-Risk (VaR)

The authors describe how to implement VaR, the risk measurement technique widely used in financial risk management.

by Simon Benninga and Zvi Wiener

In this article we discuss one of the modern risk-measuring techniques Value-at-Risk (VaR). *Mathematica* is used to demonstrate the basic methods for calculation of VaR for a hypothetical portfolio of a stock and a foreign bond.

VALUE-AT-RISK

Value-at-Risk (VaR) measures the worst expected loss under normal market conditions over a specific time interval at a given confidence level. As one of our references states: "VaR answers the question: how much can I lose with $x\%$ probability over a pre-set horizon" (J.P. Morgan, *RiskMetrics-Technical Document*). Another way of expressing this is that VaR is the lowest quantile of the potential losses that can occur within a given portfolio during a specified time period. The basic time period T and the confidence level (the quantile) q are the two major parameters that should be chosen in a way appropriate to the overall goal of risk measurement. The time horizon can differ from a few hours for an active trading desk to a year for a pension fund. When the primary goal is to satisfy external regulatory requirements, such as bank capital requirements, the quantile is typically very small (for example, 1% of worst outcomes). However for an internal risk management model used by a company to control the risk exposure the typical number is around 5% (visit the internet sites in references for more details). A general introduction to VaR can be found in Linsmeier, [Pearson 1996] and [Jorion 1997].

In the jargon of VaR, suppose that a portfolio manager has a daily VaR equal to \$1 million at 1%. This means that there is only one chance in 100 that a daily loss bigger than \$1 million occurs under normal market conditions.

A REALLY SIMPLE EXAMPLE

Suppose portfolio manager manages a portfolio which consists of a single asset. The return of the asset is normally distributed with annual mean return 10% and annual standard deviation 30%. The value of the portfolio today is \$100 million. We want to answer various simple questions about the end-of-year distribution of portfolio value:

1. What is the distribution of the end-of-year portfolio value?
2. What is the probability of a loss of more than \$20

million dollars by year end (i.e., what is the probability that the end-of-year value is less than \$80 million)?

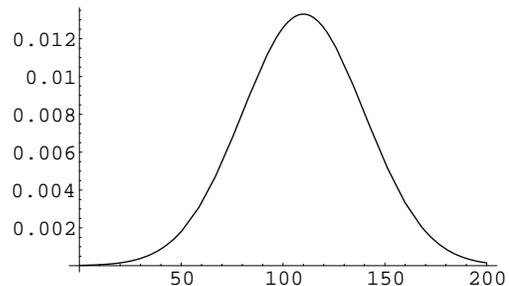
3. With 1% probability what is the maximum loss at the end of the year? This is the VaR at 1%.

We start by loading *Mathematica*'s statistical package:

```
Needs["Statistics`Master`"]
Needs["Statistics`MultiDescriptiveStatistics`"]
```

We first want to know the distribution of the end-of-year portfolio value:

```
Plot[PDF[NormalDistribution[110,30],x],{x,0,200}];
```



The probability that the end-of-year portfolio value is less than \$80 is about 15.9%.

```
CDF[NormalDistribution[110.,30],80]
0.158655
```

With a probability of 1% the end-of-year portfolio value will be less than 40.2096; this means that the VaR of the distribution is $100 - 40.2096 = 59.7904$.

```
Quantile[NormalDistribution[110.,30],0.01]
40.2096
```

We can formalize this by defining a VaR function which takes as its parameters the mean μ and standard deviation σ of the distribution as well as the VaR level x .

```
ClearAll[VaR];
VaR[mu_,sigma_,x_]:=
100-Quantile[NormalDistribution[mu,sigma],x]
VaR[110,30,0.01]
```

```
59.7904
```

LOGNORMAL DISTRIBUTIONS

As explained in our previous articles, the lognormal distribution is a more reasonable distribution for many asset prices (which can not become negative) than the normal distribution. This is not a problem: Suppose that the natural logarithm of the portfolio value is normally distributed with annual mean μ and annual standard deviation σ . Denoting the value of the portfolio by v it follows that the logarithm of the portfolio value at time T , v_T , is normally distributed:

$$\text{Log}[v_T] \sim \text{Normal} \left[\text{Log}[v] + \left(\mu - \frac{\sigma^2}{2} \right) T, \sigma T \right]$$

The term $\frac{\sigma^2}{2} T$ appears due to Ito's Lemma (see [Hull 1997]). In our case, this means that $v = 100$, $\mu = 10\%$, $\sigma = 30\%$. Thus the end-of-year log of the portfolio value is distributed `NormalDistribution[Log[100] + (0.1-0.3^2/2), 0.3] = NormalDistribution[4.666017, 0.3]`. This means that the probability that the end-of-year value of the portfolio is less than 80 is given by:

```
CDF[NormalDistribution[Log[100] +
(0.1-0.3^2/2), 0.3], Log[80]]
0.176926
```

Similarly the VaR function has to be redefined:

```
ClearAll[lognormalVaR];
lognormalVaR[mu_, sigma_, x_] :=
100-Exp[Quantile[NormalDistribution[
Log[100] + (mu-sigma^2/2), sigma], x]];
lognormalVaR[0.10, 0.30, 0.01]
47.4237
```

Thus a portfolio whose initial value is \$100 million and whose annual returns are lognormally distributed with parameters $\mu = 10\%$ and $\sigma = 30\%$, has an annual VaR equal to \$47.42 million at 1%.

Most VaR calculations are not concerned with annual value at risk. The main regulatory and management concern is with loss of portfolio value over a much shorter time period (typically several days or perhaps weeks). It is clear that the distribution formula $\text{Log}[v_T] \sim \text{Normal}[\text{Log}[v] + (\mu - \frac{\sigma^2}{2})T, \sigma T]$ can be used to calculate the VaR over any horizon. Recall that T is measured in annual terms; if there are 250 business days in a year, then the daily VaR corresponds to $T = 1/250$ (for many fixed income instruments one should use 1/360, 1/365, or 1/365.25 depending on the market convention):

```
ClearAll[lognormalVaR];
lognormalVaR[mu_, sigma_, x_, T_] := 100-Exp[
Quantile[NormalDistribution[Log[100] + (mu-
sigma^2/2)*T, sigma*T], x]];
lognormalVaR[0.10, 0.30, 0.01, 1/250]
lognormalVaR[0.10, 0.30, 0.01, 5/250]
lognormalVaR[0.10, 0.30, 0.01, 21/250]
0.256831
1.27758
5.25717
```

The daily VaR of the portfolio at 1% is \$256,831. The probability that the firm will lose more than this amount on its portfolio over the course of a single day is less than 1%. Similarly, the weekly and the monthly VaRs at 1% are \$1.28 and \$5.26 million.

A THREE-ASSET PROBLEM: THE IMPORTANCE OF THE VARIANCE-COVARIANCE MATRIX

As can be seen from the above examples, VaR is not—in principle, at least—a very complicated concept. In the implementation of VaR, however, there are two big practical problems (both problems are discussed in much greater detail in the material available on the J.P. Morgan Web site):

1. The first problem is the estimation of the parameters of asset return distributions. In “real world” applications of VaR, it is necessary to estimate means, variances, and correlations of returns. This is a not-inconsiderable problem! In this section we illustrate the importance of the correlations between asset returns. In the following section we give a highly-simplified example of the estimation of return distributions from market data. For example you can imagine that a long position in Deutschmarks and a short position in Dutch guildens is less risky than one leg only, because of a high probability that profits of one position will be mainly offset by losses of another.
2. The second problem is the actual calculation of position sizes. A large financial institution may have thousands of loans outstanding. The data base of these loans may not classify them by their riskiness, nor even by their term to maturity. Or—to give a second example—a bank may have offsetting positions in foreign currencies at different branches in different locations. A long position in Deutschmarks in New York may be offset by a short position in Deutschmarks in Geneva; the bank's risk—which we intend to measure by VaR—is based on the net position.

We start with the problem of correlations between asset returns. We continue the previous example, but assume that there are three risky assets. As before the parameters of the distributions of the asset returns are known: all the means: μ_1, μ_2, μ_3 , as well as the variance covariance matrix of the returns:

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

The matrix S is of course symmetric; σ_{ij} is the covariance of the returns of assets i and j (if $i = j$, σ_{ii} is the variance of asset i 's return).

Suppose that the total portfolio value today is \$100 million, with \$30 million invested in asset 1, \$25 million

in asset 2, and \$45 million in asset 3. Then the return distribution of the portfolio is given by:

$$\begin{aligned} \text{mean return} &= x \cdot \mu = x_1 \mu_1 + x_2 \mu_2 + x_3 \mu_3 \\ \text{variance of return} &= x \cdot S \cdot x^T, \end{aligned}$$

where $x = \{x_1, x_2, x_3\} = \{0.3, 0.25, 0.45\}$ is the vector of proportions invested in each of the three assets. Assuming that the returns are normally distributed (meaning that prices are lognormally distributed), we may calculate the VaR by:

```
Clear[x,s,means,portfolioMean,portfolioSigma];
X=Table[x[i],{i,3}];
S=Table[sigma[i,j],{i,3},{j,3}];
means=Table[mu[i],{i,3}];
portfolioMean[X_,means_,initial_]:=
  X.means*initial;
portfolioSigma[X_,s_,initial_]:=
  Sqrt[X.s.X]*initial;
```

To implement this:

```
Clear[Var];
X={0.3,0.25,0.45};
means={0.1,0.12,0.13};
initial=100;
Varlevel=0.01;
S={{0.1,0.04,0.03},{0.04,0.2,-0.04},
  {0.03,-0.04,0.6}};
Print["The portfolio mean = ",
  portfolioMean[X,1+means,initial] ]
Print["The portfolio sigma = ",
  portfolioSigma[X,S,initial]]

VaR[X_,s_,initial_,level_]:=
  initial-Quantile[
    NormalDistribution[
      portfolioMean[X,means,initial],
      portfolioSigma[X,s,initial]],Varlevel]
Print["The portfolio VaR at the ",
  Varlevel*100, "% level is ",
  VaR[X,S,initial,Varlevel] ]
```

```
The portfolio mean = 111.85
The portfolio sigma = 38.4838
The portfolio VaR at the 1. % level is 177.677
```

USING EMPIRICAL DATA TO CALCULATE THE VAR

In this section we use market data to do some VaR calculations. Our data consists of 40 business days of data for a market index, foreign interest rates and foreign exchange rates. We use this data to do three kinds of VaR calculations: Historic simulations, variance-covariance calculations, and Monte Carlo simulations.

Data Description and Preliminary Calculations

We consider the VaR of a portfolio manager who has invested in only two assets: A domestic stock index and a foreign bond. For simplicity we assume that the value of this portfolio is wholly determined by only several parameters:

- The current price of the stock index
- The foreign market interest rate (the bond is assumed to have a zero coupon, so that only the interest rate

until the bond's maturity determines its price in foreign currency).

- The time until the bond's maturity; along with the foreign market interest rate, this will determine the bond's foreign currency price.
- The exchange rate between the foreign and the domestic currency.

We use historic price and return data to do our VaR calculations. It is convenient to translate all dates to a numerical format. We set January 1, 1997 to be day number 0 and then any day is given by a number of days since the initial date (we assume that all days are business days for simplicity).

```
Needs["Miscellaneous`Calendar`"];
Needs["Statistics`Master`"];
Needs["Statistics`MultiDescriptiveStatistics`"];
Day0 = {1997, 1, 1};
dayN[ day_]:=DaysBetween[ Day0, day];
```

Thus January 13, 1997 is the day number 12 and so on.

```
dayN[{1997, 1,13}]
12
```

Consider February 10, 1997. Suppose that on February 10, 1997, the stock index value is 293., the foreign interest rate is 5.3%, and the exchange rate is 3.4; February 10, 1997 is the 40th business day since January 1, 1997. Thus we write the data for this day as a list: {40, 293., 5.3, 3.4}. Obviously any real market data must contain much more information, however will use this simple example to illustrate the VaR approach.

The whole data set with which we will be working is the following (in a real-world situation, we would obviously use much more data about many more assets):

```
dataVAR = {
  {1, 282., 5.28, 3.5},{2, 283., 5.26, 3.47},
  {3, 285., 5.23, 3.46},{4, 280., 5.24, 3.45},
  {5, 282., 5.25, 3.45},{6, 281., 5.24, 3.46},
  {7, 282., 5.24, 3.45},{8, 286., 5.25, 3.43},
  {9, 285., 5.25, 3.47},{10, 286., 5.26, 3.443},
  {11, 288., 5.27, 3.42},{12, 289., 5.28, 3.42},
  {13, 290., 5.28, 3.41},{14, 289., 5.28, 3.42},
  {15, 291., 5.29, 3.46},{16, 293., 5.31, 3.41},
  {17, 294., 5.32, 3.40},{18, 290., 5.34, 3.49},
  {19, 287., 5.35, 3.47},{20, 288., 5.34, 3.48},
  {21, 289., 5.35, 3.46},{22, 281., 5.36, 3.44},
  {23, 283., 5.23, 3.45},{24, 285., 5.24, 3.42},
  {25, 288., 5.25, 3.41},{26, 289., 5.26, 3.41},
  {27, 287., 5.26, 3.43},{28, 285., 5.28, 3.42},
  {29, 290., 5.27, 3.44},{30, 291., 5.27, 3.42},
  {31, 289., 5.27, 3.37},{32, 288., 5.29, 3.39},
  {33, 290., 5.28, 3.41},{34, 293., 5.31, 3.44},
  {35, 292., 5.32, 3.41},{36, 293., 5.28, 3.42},
  {37, 293., 5.30, 3.42},{38, 293., 5.31, 3.44},
  {39, 292., 5.32, 3.41},{40, 293., 5.30, 3.4};
```

Suppose that the financial institution holds two shares of the stock index portfolio and a short position in a zero-coupon foreign bond having face value 100 rubles with maturity May 8, 2000. We represent this portfolio as a list:

```
portfolio = { {STOCK, {}, 2},
             {FBOND, {dayN[{2000, 5, 8}]}, -1}}
```

Note that the empty parameters list in the stock part means that there are no parameters and we always keep the same portfolio of stocks. The list of parameters of the bond part consists of one number - day of maturity. In general both lists can have many parameters describing a specific portfolio of stocks or bonds.

Risk Mapping

“Risk mapping” is a nice wording for a pricing function. In a general case we might decompose the pricing function for each type of financial instrument into simple risk factors (perhaps using the option pricing techniques explained in our previous articles). However, in our simple example the current value of each instrument is a simple one-dimensional function of current market data. For example, the price of the stock is just its market value and the price of the bond is its discounted future payoff translated to a local currency according to the current exchange rate.

Next, we define the pricing functions for the stock index and the foreign bond. Suppose we are given a list of market data, `market={dayN[{1997, 3, 6}], 283., 5.26, 3.47}`. The stock market price is the second item in this list, whereas the bond price in the local currency is given by $100 * 3.47 * \exp[-5.26\% * (\text{time to maturity in years})]$. We use *Mathematica*'s calendar functions to define two dollar pricing functions:

```
Clear[stock,fbond]
stock[ param_, market_]:= market[[2]];
fbond[ param_,market_]:=
  market[[4]]*100*
  Exp[-market[[3]]/100*
    (param[[1]]-market[[1])/365.25];
junk={dayN[{1997,8,9}],355.,5.77,6.};
stock[{},junk]
fbond[{1223},junk]
355.
512.08
```

Next we define

```
pricingFun[instr_]:=
  Switch[instr, STOCK,stock, FBOND,fbond]
valueP[ portf_, mrkt_]:=
  Module[{valueList, i},
    valueList =
      Table[Apply[ pricingFun[ portf[[i,1]] ],
            {portf[[i,2]], mrkt}]*portf[[i,3]],
          {i,Length[ portf]};
    Apply[ Plus,valueList]
  ];
```

Check that it works

```
mrkt9Feb97 = {40, 293., 5.30, 3.4};
currentPrice = valueP[ portfolio, mrkt9Feb97]
299.63
```

If everything is correct we should get 299.63 since the current value of our stock position is $586 = 2 * 293$, and

the bond position has a value of $-286.37 = -100 * 3.4 * e^{-0.053 * 1183/365.25}$.

For a VaR calculation we have to fix two parameters: the time horizon and the confidence level (1 - quantile). We choose a one day time horizon and an 80% confidence level. This time horizon is a typical one for financial institutions, however the confidence level 80% is very low, but it will allow us to use a short data file. For a higher precision much more data is necessary.

Historical Simulation

The first method we are presenting here is the historical simulation. We pretend that the change in the market conditions from today to tomorrow are the same as the changes that took place some time in the past. There is an important question on what type of current changes are the same as historic changes. The main difficulty is to distinguish between multiplicative and additive types of market variables. We provide here a simplistic scheme assuming that all changes are additive, for a detailed explanation of this problem see [Grundy, Wiener 1996].

This procedure pretends that the change in market parameters from today to tomorrow will be the same as it was some time ago. The routine below returns a list of the historical prices. Note that this procedure is completely different from just pricing the current portfolio according to old data, in which case the global change in the market level is lost. Parameters of the historical simulation function are: **portf** - the current portfolio, **mrkt** - the current market data, **histData** - the historical data.

```
histSimul[ portf_, mrkt_, histData_]:=
  Module[{histPrice={}, i},
    For[ i=1, i<Length[histData], i++,
      AppendTo[ histPrice,
        valueP[portf,
          mrkt+histData[[i+1]]-histData[[i]]]
      ];
    histPrice
  ];
```

To use this simulation for VaR measurement we use the following function, which requires in addition to the parameters described above the quantile.

```
HistApproach[ portf_, mrkt_, hData_, quant_]:=
  Module[{currPrice, histPr, changes},
    currPrice = valueP[ portf, mrkt];
    histPr = histSimul[ portf, mrkt, hData];
    changes = histPr - currPrice;
    {Quantile[ changes, quant], changes, histPr}
  ];
```

The outcome of this function is a list of three objects. The first one is the VaR, the second one is a list of pretended market changes in the value of our portfolio. The third is a list of resulting portfolio values.

```
HistApproach[portfolio, mrkt9Feb97,
             dataVAR, 0.2]
{-3.0144, {4.30181, 4.52345, -9.10677, 4.05111,
-2.97692, 2.80083, 9.73534, -5.41111, 4.32482,
5.98798, 2.05111, 2.80083, -2.88394, 0.682656,
8.35297, 2.89323, -15.4328, -4.26466, 1.02308,
3.73534, -14.2647, 1.90504, 6.57746, 6.89323,
2.05111, -5.72633, -3.0144, 8.18042, 3.64322,
0.170379, -3.53993, 2.18042, 3.71166, 0.577459,
0.743977, 0.143757, -1.63311, 0.577459,
2.61594}, {303.932, 304.153, 290.523, 303.681,
296.653, 302.431, 309.365, 294.219, 303.955,
305.618, 301.681, 302.431, 296.746, 300.313,
307.983, 302.523, 284.197, 295.365, 300.653,
303.365, 285.365, 301.535, 306.207, 306.523,
301.681, 293.904, 296.615, 307.81, 303.273,
299.8, 296.09, 301.81, 303.342, 300.207,
300.374, 299.774, 297.997, 300.207, 302.246}}
```

Thus the simulated daily VaR at 20% level is a loss of \$3.0144.

Variance Covariance Approach

This method is based on the assumption that the short term changes in the market parameters and in the value of the portfolio are normal. This method also reflects the fact that the market parameters are not independent, however it is restricted to the first degree of dependence - correlation.

First based on historical data we build a vector of average daily changes in each parameter, and a historical variance-covariance matrix of these changes. Second we calculate a linear (so-called delta) approximation of our portfolio for small changes in the market (first term of the Taylor series). Third, we measure the variance of the portfolio assuming normally distributed changes in the market with this variance covariance matrix. Fourth, we measure the lowest quantile of P&L (profit and loss) for the required level of confidence.

To calculate the additive covariance matrix we use

```
AddChangesMx = Drop[dataVAR,1] - Drop[dataVAR,-1];
AddMeansVec = Mean[ AddChangesMx ];
AddCovMx = CovarianceMatrix[ AddChangesMx ];
sqrtAddCovar = MatrixPower[ AddCovMx, 1/2 ];
```

This completes the first step. We also recommend to verify that the covariance matrix is positive definite (and well-defined), since otherwise you may experience some problems with stability.

The mean vector and covariance matrix of the market data are

```
AddMeansVec
{1, 0.282051, 0.000512821, -0.0025641}
```

AddCovMx // MatrixForm

$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 6.10256 & -0.0122537 & -0.00328408 \\ 0 & -0.0122537 & 0.000662888 & -0.0000302294 \\ 0 & -0.00328408 & -0.0000302294 & 0.0007132 \end{pmatrix}$$

The first line and column correspond to the time change, which is deterministic and thus uncorrelated with the rest of the data.

Since our portfolio is not linear in the market parameters we should measure its sensitivity to a small changes in each parameter. There are different ways to do this. We demonstrate below a way when we price the portfolio under the current market data, then we increase one parameter by 1% of its value, and price the portfolio again. Then take the finite difference and estimate the derivative.

Denote the price of a portfolio by P, this is a deterministic function of the market data P(x). To measure the sensitivity we consider $P(x + \Delta_i x) - P(x)$ where $\Delta_i x = (0, 0, \dots, 0, 0.01x_i, 0, \dots, 0)$. To estimate the partial derivative of P with respect to the i-th argument we calculate

$$\frac{\partial P(x)}{\partial x_i} \approx \frac{P(x + \Delta_i x) - P(x)}{0.01x_i}$$

```
sensitivity[ portf_, mrkt_]:=
Module[{currentPrice, unitVec,
eps=0.01, deriv, i},
currentPrice = valueP[ portf,mrkt];
unitVec = Table[ 1, {4}];
deriv = Table[(valueP[ portf,
mrkt*ReplacePart[ unitVec, 1+eps, i]]-
currentPrice)/(mrkt[[i]]*eps), {i, 4}
];
```

Applying this function to our portfolio we get:

```
sensitivity[ portfolio, mrkt9Feb97 ]
{-0.0415553, 2., 9.26722, -84.2265}
```

We will also need the matrix of second derivatives of the pricing function. Using finite differences it can be calculated as:

```
ClearAll[secondDeriv];
secondDeriv[portf_, mrkt_]:=
Module[{zerVec= Table[ 0, {4}],
unitVec=Table[ 1, {4}], eps=0.01, i,j},
Table[
(valueP[ portf,
mrkt*(unitVec+ReplacePart[ zerVec, eps, i]+
ReplacePart[ zerVec, eps, j]])-
valueP[ portf,
mrkt*(unitVec+ReplacePart[ zerVec, eps, i]+
ReplacePart[ zerVec, -eps, j]])-
valueP[ portf,
mrkt*(unitVec+ReplacePart[zerVec, -eps, i]+
ReplacePart[ zerVec, eps, j]])+
valueP[ portf,
mrkt*(unitVec+ReplacePart[zerVec, -eps, i]+
ReplacePart[ zerVec, -eps, j]])]/4/
(mrkt[[i]]*eps)/(mrkt[[j]]*eps),{i,4},{j,4}
];
```

```
secondDeriv[ portfolio, mrkt9Feb97] //MatrixForm
```

$$\begin{pmatrix} -6.02975 \times 10^{-6} & 0. & -0.00649451 & -0.0122218 \\ 0. & 0. & 0. & 0. \\ -0.00649451 & 0. & -0.300412 & 2.72799 \\ -0.0122218 & 0. & 2.72799 & 0. \end{pmatrix}$$

Zeros indicate that the part of our portfolio that depends on the stock market is not influenced directly by the interest and exchange rates. It still is influenced by correlations between different market factors, but indirectly only.

Price of a portfolio is a function of the market data, say $P(x)$, where x is the vector of the market data. The current parameters of the market are known x_0 , however tomorrow the market will move to a new vector x_1 . The important simplifying assumption of the variance covariance approach is that the changes of the parameter vector are assumed to be normally distributed. Then we can write

$$P(x_1) = P(x_0) + P'(x_0)(x_1 - x_0) + \frac{1}{2}P''(x_0)(x_1 - x_0)^2 + o(x_1 - x_0)'$$

We keep the second order term due to Ito's lemma. Assume that the market parameter x follows a simple arithmetic Brownian motion. $dx = \mu dt + \sigma dB$, then by Ito's lemma we can keep only terms up to the order dt by choosing

$$P(x_1) \approx P(x_0) + P'(x_0)(\mu \Delta t + \sigma \Delta B) + \frac{1}{2}P''(x_0)\sigma^2 \Delta t + o(\Delta t)$$

We assume that all components of the vector $\Delta x = x_1 - x_0$ are normally distributed, and the covariance matrix is known, $\sigma^2 = S(x,t)$. Then the right hand side is a linear combination of normally distributed random variables. As soon as all means are very small (note that daily expected changes are close to zero), the change in value can be approximated by a normally distributed random variable with mean and variance defined below:

$$\begin{aligned} \mu_p &= E(P(x_1) - P(x_0)) = \\ & \left(\mu \cdot P'(x_0) + \frac{1}{2} \text{tr}(P''(x_0) \cdot S) \right) \Delta t \\ \sigma_p^2 &= \text{var}(P(x_1) - P(x_0)) = P'(x_0) \cdot S \cdot (P'(x_0))^T \Delta t \end{aligned}$$

Here the superscript T denotes transposition and we use the standard vector product of lists. Then the lower q-quantile can be approximated by a q-quantile of the normal distribution $N(\mu_p, \sigma_p)$

$$\text{VaR}(q) = \mu_p + \sigma_p \cdot \text{Quantile}(N(0, 1), q)$$

The sign plus here corresponds to the lowest quantile, since, the $\text{Quantile}(N(0,1), q)$ is negative for $q < 0.5$. Note that $E(\Delta x)$ and S are characteristics of the market and not of the portfolio. Thus they can be used for measuring risks of different portfolios. The *Mathematica* implementation of this idea looks like:

```
Clear[VarCovarAdd]
VarCovarAdd[ portf_, mrkt_, varCovarMx_,
             drift_, quant_]:=
Module[{sens, secDeriv,i},
  sens = sensitivity[ portf, mrkt];
  secDeriv=secondDeriv[ portf,mrkt];
  drift.sens+
  Sum[(varCovarMx.secDeriv)[[i,i]], {i,4}]/2+
  Sqrt[sens.varCovarMx.sens]*
  Quantile[ NormalDistribution[0,1], quant]
];
VarCovarAdd[portfolio, mrkt9Feb97,
             AddCovMx, AddMeansVec, 0.2]
-3.87143
```

Note that **Trace** means in *Mathematica* something completely different from the mathematical trace of a matrix, thus we use **Sum** instead.

Monte Carlo Approach

This method is based on the assumption that we have some information about the joint distribution of market changes. Then using this distribution we can draw randomly a large number of scenarios and price the portfolio for each scenario. A rich set of scenarios will give a good approximation for the distribution of final value of the portfolio. The lowest q-quantile of this distribution can be used as an approximation to VaR. Moreover this method allows a dynamic improvement. One can run a small set of simulations, get a preliminary result and then improve it by running additional simulations if necessary.

In the example below we choose the simplest form of the joint distribution - all market parameters are distributed jointly normal with the same mean and covariance as we have measured above. In a general case one can provide any reasonable distribution and use the same method.

```
nor[mu_, sig_]:=
  Random[ NormalDistribution[mu,sig] ];
MCvarAdd[portf_, mrkt_, n_:10]:=
Module[
  {tbl, simulatedParam, values},
  tbl = Table[ sqrtAddCovar.Table[ nor[0,1],
  {Length[AddCovMx]}] + AddMeansVec, {n} ];
  simulatedParam = Table[mrkt,{n}] + tbl;
  values =
  Map[valueP[ portf,#]&, simulatedParam ];
  {Mean[values], StandardDeviation[values],
  StandardErrorOfSampleMean[values],
  values, simulatedParam }
];
```

To generate a vector of correlated normally distributed random variables we generate first a vector of independent random variables (normally distributed) and then multiply it by the square root of the covariance matrix. It is left as an exercise to verify that this leads to the required correlation. The `tbl` table is the resulting table of market changes.

The following function incorporates all the necessary steps for Monte Carlo approach to VaR. The parameters of this function are: **portf** portfolio, **mrkt** current market

data, **quant** quantile, **n** number of generated scenarios. This function returns the VaR estimate and the list of all changes in the value of the portfolio that were observed in this simulation.

```
MCApproachAdd[ portf_, mrkt_, quant_, n_:10 ] :=
Module[{currentValue, changes},
  currentValue = valueP[ portf, mrkt];
  changes = MCvarAdd[portf, mrkt, n][[4]] -
    currentValue;
  {Quantile[ changes, quant], changes}
];
```

To get consistent results this function can be combined with the random seed generator. For 1000 simulations (about 1 minute of computational time) we get:

```
SeedRandom[1];
MCApproachAdd[ portfolio, mrkt9Feb97, 0.2,
1000][[1]]
-3.71068
```

Running this simulation again we get a similar result:

```
SeedRandom[7];
MCApproachAdd[ portfolio, mrkt9Feb97,
0.2, 1000][[1]]
-3.76359
```

Discussion

The historical simulation method is useful when the amount of data is not very large and we do not have enough information about the profit and loss distribution. It is usually very time consuming, but its main advantage is that it catches all recent market crashes. This feature is very important for risk measurement.

The variance covariance method is the fastest. However it relies heavily on several assumptions about the distribution of market data and linear approximation of the portfolio. It is probably the best method for quick estimates of VaR. However one should be very careful when using this method for a non-linear portfolio, especially in the case of high convexity in options or bonds.

The Monte Carlo simulation method is very slow, but it is probably the most powerful method. It is flexible enough to incorporate private information together with historical observations. There are many methods of speeding calculations, so-called variance reduction techniques.

The results of all three methods are similar and our goal was to demonstrate a very basic approach to risk measurement techniques using *Mathematica*.

REFERENCES

BEDER, TANYA (1996), VAR: Seductive But Dangerous, *Financial Analyst Journal*, September-October, pp. 12–24.

GRUNDY, BRUCE D. and ZVI WIENER. 1996. The Analysis of VAR, Deltas and State Prices: A New Approach. mimeo The Wharton School, University of Pennsylvania.

HULL, JOHN C. *Options Futures, and Other Derivatives*, Third Edition. Prentice-Hall, 1997.

JORION, PHILIPPE, *Value at Risk, the New Benchmark for Controlling Market Risk*, McGraw-Hill, 1997.

LINSMEIER, THOMAS and NEIL PEARSON. Risk Measurement: An Introduction to Value at Risk. mimeo, University of Illinois, 1996.

Several Web sites offer information on value-at-risk:

<http://pw2.netcom.com/~bschacht/varbiblio.html>—an excellent source of VaR related materials.

<http://www.riskmetrics.reuters.com/>—Reuters materials on riskmetrics.

<http://www.jpmorgan.com/RiskManagement/RiskMetrics/RiskMetrics.html>

ABOUT THE AUTHORS

The authors acknowledge grants from Wolfram Research, the Krueger and Eshkol Centers at the Hebrew University, and the Israeli Academy of Science. Wiener's research has benefited from a grant from the Israel Foundations Trustees, the Alon Fellowship and the Eshkol grant.

Simon Benninga is professor of finance at Tel-Aviv University (Israel) and the Wharton School of the University of Pennsylvania. He is the author of *Financial Modeling* (MIT Press, 1997) and of *Corporate Finance: A Valuation Approach* (with Oded Sarig, McGraw-Hill, 1997); he is also the editor of the *European Finance Review*.

Simon Benninga
Faculty of Management
Tel-Aviv University, Tel-Aviv, Israel
benninga@post.tau.ac.il
<http://finance.wharton.upenn.edu/~benninga>

Zvi Wiener is assistant professor of finance at the business school of the Hebrew University of Jerusalem. His finance research concentrates on the pricing of derivative securities, Value-at-Risk, computational finance and stochastic dominance. He wrote this article while visiting at the Olin School of Business at Washington University in St. Louis.

Zvi Wiener
Finance Department, Business School
Hebrew University, Jerusalem, Israel
mswiener@mscc.huji.ac.il
<http://pluto.mssc.huji.ac.il/~mswiener/zvi.html>

ELECTRONIC SUBSCRIPTIONS

Included in the distribution for each electronic subscription is the file `varisk.nb`, containing *Mathematica* code for the material described in this article.

