

Introduction to VaR (Value-at-Risk)

Zvi Wiener*

Risk Management and Regulation in Banking

Jerusalem, 18 May 1997

* Business School, The Hebrew University of Jerusalem, Jerusalem, 91905, ISRAEL

mzwiener@pluto.msc.huji.ac.il

972-2-588-3049

The author wish to thank David Ruthenberg and Benzi Schreiber, the organizers of the conference. The author gratefully acknowledge financial support from the Krueger and Eshkol Centers, the Israel Foundations Trustees, the Israel Science Foundation, and the Alon Fellowship.

Introduction to VaR (Value-at-Risk)

Abstract

The concept of Value-at-Risk is described. We discuss how this risk characteristic can be used for supervision and for internal control. Several parametric and non-parametric methods to measure Value-at-Risk are discussed. The non-parametric approach is represented by historical simulations and Monte-Carlo methods. Variance covariance and some analytical models are used to demonstrate the parametric approach. Finally, we briefly discuss the backtesting procedure.

Introduction to VaR (Value-at-Risk)

I. An overview of risk measurement techniques

Modern financial theory is based on several important principles, two of which are no-arbitrage and risk aversion. The single major source of profit is risk. The expected return depends heavily on the level of risk of an investment. Although the idea of risk seems to be intuitively clear, it is difficult to formalize it. Several attempts to do so have been undertaken with various degree of success. There is an efficient way to quantify risk in almost every single market. However, each method is deeply associated with its specific market and can not be applied directly to other markets. Value-at-Risk (VaR) is an integrated way to deal with different markets and different risks and to combine all of the factors into a single number which is a good indicator of the overall risk level.

For example, the major risk associated with a government bond is the interest rate risk. A simple way to measure it is by duration. More subtle dependence on the term structure can be estimated by convexity or by sensitivity to a non-parallel shift. The more detailed analysis of the term structure of interest rates usually involves a parametric term structure model. Within the framework of this model the corresponding risk factors can be quantified (at least numerically if not analytically).

However, as soon as corporate bonds are included in a portfolio there is a need to deal with default (credit) risk. This is a completely different type of uncertainty which is correlated with all of the major indices of the economy. Therefore, to measure the risk we would need to analyze the counterparty. The simplest way to deal with this type of risk is by using credit rating systems (provided in many cases by credit agencies).

For more complicated financial instruments, like interest rates and currency swaps, there is a need to develop a model of default and recovery (see for example Duffie and Pan [1997]). This type of model is necessary for measuring the risk of credit derivatives, an instrument which has become very popular recently.

For an equity portfolio the most useful way to measure risk is by volatility of returns. This parameter, together with the expected level of returns and correlations between different assets, leads to standard investment models like CAPM and APT. These models use diversification as the major tool to improve the risk-return ratio and some type of correlation

(like beta) as a risk measuring tool. However, it is clear that the equity market is strongly correlated with other types of risks (like political events and litigation risk). Therefore it is not sufficient to deal with stock markets in isolation.

Derivative instruments became one of the major sources of hedging strategies, as well as speculations. The risk associated with each single option, or a portfolio of options, can be measured by the whole set of the Greeks: delta, gamma, theta, etc. Each of them represents the rate of the change in price when only one of the parameters changes (like the underlying asset, volatility, time to maturity etc.). Again, these indices are well suited to watch the very specific types of risk related to derivative instruments, but can not be applied directly to other instruments.

Foreign exchange is another area with its own risk indices. They typically include spreads and exchange rate volatilities. However, lately some other characteristics have been added, such as target zones and the term structure models of several related currencies.

We can summarize the major risk measurement tools in the following table:

Asset	Risk measure
bonds	duration, convexity, term-structure models
credit	rating, default models
stocks	volatility, correlations, beta
derivatives	delta, gamma, vega
forex	target zones, spreads

The variety of different methods still does not give a simple answer to the most basic question “What is the current risk?”, which every financial institution should ask itself. There are several reasons why this question can not be answered in a simple way. First of all, there is no simple answer. Risk is something infinitely dimensional, it depends on all of the possible events in the world, and it is impossible to represent it precisely with a single number or any set of scenarios. For example, the only genuine answer to the question “What is the maximal loss we can suffer over some time horizon?” is unfortunately, “We can lose everything!”. However, the probability of this event is very small. “Very small” must be

quantified in some way. Whenever the total risk is low this probability is quite small, however it is higher when we have a higher total risk.

Imagine a CFO of a company trying to explain the level of risk by giving the whole distribution of profits and losses under all possible scenarios. This can easily include many different parameters with complicated correlation between them. Even if he had such a graph, it would be extremely difficult to use it to compare even two different portfolios. This information must be reduced to a single number. VaR (Value-at-Risk) was introduced precisely with the purpose of providing a one dimensional approximation of the infinitely dimensional risk. Because of the simplifying nature of this index one should always take into account its limitations. However, it can be very useful for risk supervision, reporting, division of resources and overall risk management. Some interesting applications can be found in Marshall and Siegel [1997].

II. The concept of VaR

This concept came from the industry and is relatively simple. In using VaR, one must take into account that as there is no precise answer, it will be only a rough approximation with limited usage. Imagine that we are interested in a simple risk measurement of a portfolio consisting of many different types of securities. Assume first that we can price the portfolio if we knew all major market characteristics. For example take a situation in which there are only two major factors – the stock market index and the interest rate. We can plot the value of the portfolio as a function of these parameters in the following way:

(insert figure 1 here, portfolio value as a function of interest rate and stock market)

However, this figure does not tell us what the probabilities of getting different values of the basic parameters are. Moreover it does not take into account correlations between different parameters. One needs to add a joint distribution of the two basic parameters and to consider the distribution together with the graph above. The picture has already become too complicated, even with only two basic factors.

Another way is to represent the risky nature of the investment is to plot a graph which gives the probabilities of changes in the market value of the whole portfolio. This can be done either as a density function as in Figure 2 or as a cumulative probability as in Figure 3.

(insert figures 2,3 here, density function, cumulative density)

This graph deals directly with the variable we are interested in – the probability of heavy losses. Since we are interested in the most unfavorable outcomes, we would like to measure the level of losses which will not be exceeded with some fixed probability. The critical level typically used for this fixed probability is either 5% or 1%. This quantile is the corresponding Value-at-Risk as shown in figures 2 and 3. Thus the idea behind VaR is to measure the heaviest losses which will not occur by the end of the basic time interval with some fixed probability.

To draw such a graph of probabilities the following steps should be taken:

1. Data must be gathered regarding the current position.
2. The basic time period must be set.
3. A model must be built, predicting the distribution of the prices of the securities in the portfolio as a function of basic market parameters.

After these three steps are completed, one can draw the value-probability graph, at least in theory, since in practice this might require too much time and other resources.

Historically, VaR was introduced by financial companies in the late 80's. Since then this approach has become popular among practitioners, the academic community and – most significantly – among central banks. A more detailed description of the historical development of VaR can be found in Duffie and Pan [1997], Linsmeier and Pearson [1996], and <http://pw2.netcom.com/~bschacht/varbiblio.html>. The RiskMetrics™ initiative of J.P.Morgan played a very positive role, and made VaR a widely used risk management tool.

We want to emphasize that VaR itself is not “the true way to measure risk” in the absolute sense. There are many other risk measurements. The major advantage of VaR is that it has become a widely accepted standard! Once the industry and the regulators agreed to use it as the primary risk management tool, it plays an important unifying role, and enables a

a good basis for comparison of risk between portfolios, institutions and financial intermediaries. Two necessary conditions for the success of VaR are relative simplicity of implementation and stability of results. In what follows we will discuss some major methods of implementation of VaR.

III. Choice of parameters and current position

Any risk management project should begin with meticulous collection of data. One can not expect a model to produce reasonable results when it has very inaccurate data. The data collection consists of three major steps.

First, we must know our current portfolio. Even this simple question is not as innocent as it seems. In the modern global financial system, trading is ongoing throughout the entire day. This raises the question: when to collect data. The most popular method is to collect the data at the end of the working day on the major market (the market with major risk exposures for the financial institution). However, this approach can not be used by intraday speculators or market makers who typically have large positions throughout the day but close or hedge most of them by the end of the day.

The second type of data one needs to collect is the current market data. Just to demonstrate that this is not a simple task, let us recall that equity prices are disclosed almost continuously, interest rates may have significant jumps around certain dates (for example meetings of central banks), inflation rates are published with a delay of several weeks and credit ratings of companies and countries are also typically published with a significant delay. This raises the question of data timing. The solution to this question mainly depends on the type of portfolio and the major goal of VaR measurement, which we will discuss below.

The third type of data is historical observations. The amount of historical information is huge. Moreover the historical information is needed to estimate the major market characteristics for the next basic period. In effect, we are basing our predictions on a very problematic assumption, i.e. that the future behavior of the markets can be predicted by its past behavior. Although this approach is problematic, there is no other way to study the market. Although there is no need to assume stationarity of markets, the models we discuss are typically based on historical estimates.

There is an alternative way to measure basic market characteristics without relying solely on historical data. In financial markets which are well developed, one can use the implied parameters (like implied volatilities, see Jackwerth and Rubinstein [1997]). However, this approach is problematic to the extent that by measuring VaR we are typically interested in estimating the maximal realistic losses. By using the implied parameters we rely heavily on market expectations and can suffer from underestimation of “fat tails” by the whole market.

The next question we have to deal with is – what level of confidence and basic time period to use. The answer depends primarily on the final goal. When the primary goal is to satisfy imposed restrictions, such as central bank requirements, there is not much choice. Most central banks have adopted the Basle requirements, which allow to choose between the standardized approach and an internal model (also called the alternative approach). The standardized approach is simple for implementation, but it does not account for correlations between different markets. In most situations it leads to a significant overestimation of risk. The alternative approach allows the use of correlations between markets. This method is more time consuming but much more flexible and typically leads to more realistic risk estimates. The alternative approach for banks requires a 99% confidence level (1% of worst outcomes) and a 10 day time horizon. A simple calculation shows that this leads to a very conservative estimate. An exception (a loss exceeding the predicted level) can be expected only once in every 100-10 days, i.e. very rarely (once every few years). However, when the primary goal is to build an efficient internal risk management model it is strongly recommended to use parameters which lead to an observable and verifiable result. Typically firms use a 95% confidence level (5% of worst outcomes) and a short time horizon, such as 1 day. In this case, losses greater than the predicted level can be expected every 20-1 days. That is, roughly once a month we can suffer a loss which is heavier than the VaR level¹.

Another important issue is the different type of risks taken by hedgers and speculators. Here is one example. Consider two American firms both having occasionally the same position: £1M in cash and a future contract to exchange £1M to \$1.6M 3 months from now. The first institution is a pure speculator and occasionally has an almost offsetting position. In one day's time there will be a new round of trade and there are almost no

¹ We learned this explanation from D. Pyle on the “Risk Management and Regulation in Banking” conference, Jerusalem 1997.

chances that this parity will remain. The second institution is a pure hedger and tries to prevent losses in the current position by a simple hedge. This strategy is precommitted and will be kept until maturity with daily readjustments. When working with a 10 day time horizon, how can we compare the riskiness of these firms? Intuitively we would like to say that the hedger is less risky. However, if we use the standard frozen approach we assume that the current position is unchanged during the whole period, and we get the same results. This clearly is not an accurate interpretation of the available information about the two very different companies.

Note that the frozen assumption is invalid for both companies. However, we know that the speculator will change positions according to his strategy and the hedger will always readjust his position according to changes in exchange rates and \$, £ interest rates. The possible solution can be based on Dynamic VaR – a modification of the standard VaR which takes into account some precommitted hedging strategies. This approach should be used when dealing with a pure speculator who uses stop-loss orders to prevent big losses. Here again, one should take into account the precommitted strategy of position change.

IV. Methods for VaR measurement

We will discuss several widely used methods to calculate VaR. Typically this is not a very precise type of calculation. There is a need to rely on many different parameters, and each one has a small error. As a result the raw data already has a lot of noise. It does not make sense to perform all of the calculations very precisely once we know that the initial data is not precise. Typically for bank supervision and for internal control it is enough only to estimate VaR from above, but the problem is more subtle when VaR is used for distribution of resources between trading desks. In this situation one should perform a deeper analysis to decide how much risk each separate division should be allowed to have. The common policy of central banks is to set capital requirements as a function of the VaR as specified by the Basle committee. In many cases an upper bound is sufficient when the bank capital is above the required level (a common situation) and there is no requirement for additional capital.

We discuss later the two basic types of methods: parametric and non-parametric. Parametric methods will include the variance-covariance approach and some analytical methods. The non-parametric model includes historical simulation and the Monte-Carlo approach. All VaR measurement approaches use a similar scheme:

- a. Selection of basic parameters (time horizon, confidence level, time of measurement)
- b. Selection of relevant market factors
- c. Risk mapping
- d. VaR calculation

For step (a) we define the relevant parameters according to our goals and resources. The next two steps, (b) and (c), assume some kind of model, either just a set of relevant factors or a completely specified pricing model. In any case the relatively small set of relevant parameters should be defined, and some method for portfolio valuation based on this set should be established. Step (d) includes the calculation itself. This step can be very time consuming, especially when Monte-Carlo methods are used. There are numerous techniques for speeding the calculations.

The following are the different types of techniques to calculate VaR.

A. Historical simulations.

This is probably the simplest non-parametric method. There is no assumption of a complex structure of the markets. Instead we observe the historical behavior of our current portfolio over the last few years. We begin by measuring the daily percentage changes in the market parameters. Then we apply these changes to our current portfolio and measure the corresponding profits and losses. The most useful version of this approach is when the risk mapping procedure defines the price of the whole portfolio as a deterministic function of the market parameters $P(p)$. Here P is the pricing function and p is the vector of all relevant market parameters. Then today's (day t) price is $P(p_t)$. The market parameters at some day j were p_j and on day $j+1$ the parameters were p_{j+1} . Then we can model the possible changes in today's parameters in the following ways. We can use the relative change, where each market parameter is multiplied by the ratio of the same parameter at day $j+1$ and day j . Another approach is when we add to today's value the difference between the values at day $j+1$ and day j for each parameter. The multiplicative method is applicable when the volatility increases with the level of the parameter. This method is useful for stock indexes, exchange rates, etc. The additive approach assumes that the volatility is level independent. For example, for the additive approach we would take as a possible price tomorrow

$P(p_t + (p_{j+1} - p_j))$. More complex combinations of both methods can be used as well. For example when modeling exchange rates in a band or interest rates.

Using a moving window, we calculate the profits and losses for each ten day period. After ordering all the resulting data, we set the level of VaR at the 5% quantile of worst outcomes (assuming that we are working with a 95% confidence interval). An alternative method is to calculate the profits and losses for each one day period, proceed as above, and multiply the result by $\sqrt{10}$. Both methods are applicable as per the Basle regulations.

This approach is relatively simple and it does not require simulations or the development of an analytical model. Moreover it can easily incorporate non-linear instruments such as options.

A typical problem with this approach is that there is not enough data. The further we go into the past for data, the less relevant this information is to today's market. This is not a simple trade-off. On the one hand, we would like to have more data in order to observe the rare events, especially the heavy losses. On the other hand, we do not want to build our current risk estimates on very old market data. Let's assume that we have agreed to take the last five years of data for our VaR estimate. If there was a big loss on a particular day, then exactly five years later the big market jump will not appear in the set of data we use. This will lead to a jump in our VaR estimate from one day to the next. This demonstrates that the results are not stable when using the historical simulations approach.

One important situation in which the historical simulation approach can not be used is for technical trading strategies developed on the basis of historical data. Technical trading strategies are generally conceived on the basis of historical data, and produce the best results on this data (i.e. most of big losses are a posteriori excluded). In such a case, one can certainly not use the data which was already used for calibration as the data set for the VaR estimate.

B. Variance covariance.

This is a parametric method, based on the assumption that the returns are normally distributed. Historical data is used to measure the major parameters: means, standard deviations, correlations. The overall distribution of the market parameters is constructed from this data. Using the risk mapping technique, the distribution of the profits and losses

over the time horizon (typically one day) can be found. When the market value of the portfolio is a linear function of the underlying parameters, the distribution of the profits is normal as well. Therefore, the 5% quantile corresponding to VaR can be calculated at $1.65 \cdot \sigma$ below the mean ($2.33 \cdot \sigma$ will give the 1% level). One significant advantage of this scheme is that for many market parameters all of the relevant data is well known. The J.P.Morgan's RiskMetrics^{TM2} is probably the best source for this type of data in many markets. It is free and easily accessible at <http://www.jpmorgan.com>.

The strong side of this approach is that it is flexible, simple and widely used. It also enables the addition of specific scenarios and enables the analysis of the sensitivity of the results with respect to the parameters. However, it relies heavily on the important assumption that all of the major market parameters are normally distributed. Therefore when a significant portion of the portfolio is not linear (with options for example) this method can not be used directly.

Another field of limited usage is foreign exchange, especially with European currencies. Whenever a target zone is used one can rely on this approach only when the exchange rate is in the middle of the band. However this method becomes very problematic when the rate is close to the boundaries, because in this situation the distribution of changes is far from normal. Figure 4 demonstrates that whenever the exchange rate is close to the middle of the target zone (point A), the next day's distribution of the exchange rate is practically normal. However, when the current exchange rate is at point B, the distribution has a completely different shape. Moreover, as some models of exchange rates show (see Ingersoll [1996]), as soon as the time horizon increases, the future distribution of the exchange rate becomes du-modal, even when the rate is currently in the middle of the band.

(insert figure 4 here)

An additional problem with this approach is that historical distributions of market returns are far from being normal. This problem is well known and is related to fat tails

² After July 17, 1997 the data set is available through Reuters WEB page: <http://www.riskmetrics.reuters.com>

(kurtosis). One can find several techniques of how to deal with these tails in the paper by Duffie and Pan [1997].

All of the problems above must be taken into account when using this approach. Even having solved them, one must remember that the direct usage of the variance covariance method for wide portfolios is restricted to simple linear instruments.

The major market information used in this approach is in the variance covariance matrix (and means). All of these parameters are typically extracted from historical data. The question of which time interval to use for historical correlations immediately arises. The proposed alternative Basle approach requires an estimation of the variances from at least the last year. A naive formula of covariances based on historical observations has two major drawbacks (both are described in details in Alexander, Leigh [1997]).

First, an important property of the variance covariance matrix is that it must be positive definite. However, because of the very high degree of correlation, even a small error in the data can easily lead to loss of this property. The standard procedure applied in this case is to decompose the matrix either to eigenvalues or to singular values and to replace small negative values (arising from numerical errors in data) by zeroes.

A second difficulty is related to the fixed rolling time window. The problem here is exactly the same as the one we have mentioned when discussing the historical simulations approach. Any big market move creates an abrupt change in all major parameters when the window moves.

The solution can be found with an appropriate weighting scheme. More weight is given to the recent past and less weight to the older events. One specific realization of this method is described in Alexander, Leigh [1997]. Many similar versions of the weighting scheme can be used depending on the situation.

C. Monte Carlo simulations.

This is another non parametric method. It is probably one of the most popular methods among sophisticated users. It does not assume a specific form of the distributions. The first step is to identify the important market factors. Next, one should build a joint distribution of these factors based on one of the following: historical data; data implicitly implied by observed prices; data based on specific economic scenarios. Finally, the

simulation is performed, typically with a large number of scenarios. The profit and losses at the end of the period are measured for each scenario. As in the other methods, these numbers should be ordered and the 5% quantile of the worst results is the VaR estimate.

This method has several important advantages. First, it does not assume a specific model and can be easily adjusted to economic forecasts. The results can be improved by taking a larger number of simulated scenarios. Options and other nonlinear instruments can be easily included in a portfolio. In addition, one can track path-dependence because the whole market process is simulated rather than the final result alone.

One important disadvantage is very slow convergence. Any Monte Carlo type simulation converges to the true value as $\frac{1}{\sqrt{N}}$, where N is the total number of simulated trajectories. This means that in order to increase the precision by a factor of 10 one must perform 100 times more simulations. This problem is the most serious disadvantage of this method. However in many cases there are well developed techniques of variance reduction. They are typically based on known properties of the portfolio, such as correlations between some markets and securities, or known analytical approximations to options and fixed income instruments.

Another class of useful techniques for speeding the standard Monte Carlo approach is portfolio compression, by which one can represent a large portfolio of similar instruments as a single artificial instrument with risk characteristics similar to the original portfolio. One important requirement for this bunching is that all of the instruments that we wish to unify have very similar risk characteristics.

When using the Monte Carlo simulations one should generate a large number of trajectories for all of the major parameters and then price every single instrument along these trajectories. The pricing of all of the instruments is very time consuming, however it can be reduced significantly when one divides all of the instruments into similar groups. This method of variance reduction is based on the knowledge that the instruments that we want to group have similar risk characteristics.

For example, consider a large portfolio of similar bonds. Each one has the current value of $\sum_{j=1}^{T_i} C_{ij} e^{-r(t_{ij})t_{ji}}$, assuming a bond with an index i and coupons C_{ij} paid at time t_{ij} . One can replace the whole portfolio with one artificial bond with the coupon rate C, and time to

maturity T , such that its price is equal to the price of the whole portfolio and its duration is the same as the portfolio's duration. Thus the risk characteristics of this bond are similar to the properties of the whole portfolio. Now, instead of mapping the risk for each individual instrument we can consider the influence of market moves on the two parameters C and T . This can save a significant amount of computational time while keeping high accuracy. With this method there is no need to price every single bond along every path when running many simulations. Instead, one can use the artificial bond as a good approximation to the whole portfolio.

An additional problem with Monte Carlo simulations is that one needs to know the joint distribution of many market parameters. When there are more than 3-4 important parameters it is not easy to clean all the data and to build this multidimensional distribution. This is especially true in cases in which the variables are strongly correlated. Imagine a set of bonds with different maturities and different credit ratings. Together they form a very complex random structure, with different variables which are interconnected and which can not be measured easily.

An advantage of the Monte Carlo method is that it allows the use of the preliminary results of all of the methods mentioned above. The historical simulations can give a first approximation to the distribution functions. The variance covariance shows which connections between variables are important and which can be neglected.

In addition, one can easily perform stress testing on the Monte Carlo simulation or perform a more detailed analysis of a specific set of scenarios, including dynamic strategies, such as prepayments or partial recoveries.

D. Analytical methods.

This set of methods is based on a specific parameterization of the market behavior. Typically these methods can not lead to a precise evaluation of VaR, however the upper and lower bounds can be found. These bounds are very useful in many situations in which there is no immediate need to raise additional capital. The analytical methods for VaR evaluation can be used in the several fields. For example, analytical methods can be used to measure the VaR of a portfolio of options in a Black-Scholes setting. In addition, an analytical estimate of VaR is applicable to some fixed income instruments with an underlying term structure

model. Exchange rates in a target zone can provide another framework in which an analytical solution for VaR can be found. In most of the cases above the analytical approach is only when we have a dynamic model of the market leading to some bounds on the probability distribution of market parameters.

Let's consider an example of a simple stochastic model of market prices. Assume that s is a market index and its dynamic is approximated by:

$$ds = v(s,t)dt + z(s,t)dB_t.$$

Here v is a drift and z – the diffusion parameter equal to the volatility times the current price $z(s,t) = \sigma(s,t) s$. B_t indicates the standard Brownian motion representing the market noise. The standard Black-Scholes option pricing model is a specific case of this approach. Among different methods we will describe briefly the change of variables method which leads to a tight lower and upper bound for a portfolio of derivatives. A detailed description of the method can be found in Grundy, Wiener [1996].

Assuming the above the major problem is in modeling the behavior of such a diffusion process with non constant volatility. One can use the following monotonic change of variables:

$$F(s,t) = \int_{A(t)}^{s\alpha(t)} \frac{a(t)}{\sigma\left(\frac{x}{\alpha(t)}, t\right)x} dx$$

Here $A(t)$, $a(t)$ and $\alpha(t)$ are free parameters which can be used to simplify the new equation. They have the following meanings: $A(t)$ sets the zero level of the new variable, $a(t)$ gives the new diffusion parameter (in most cases the simplest choice for $a(t)$ is to be a constant), $\alpha(t)$ represents the predictable component of price changes (like inflation). The new variable F gives the generalization of returns for the original variable s , as one can see from the example below.

To demonstrate the method we consider the simple Black-Scholes type of dynamics. In this case there is no real need to apply the change of variables, however we use it to demonstrate the method. Let the price of the underlying asset follow the geometric Brownian motion with a drift:

$$ds = \mu(s,t)s dt + \sigma s dB_t$$

Applying the change of variables suggested above with $A(t) = K$ (typically set at the strike price), $a(t) = 1$, $\alpha(t) = 1$, we get:

$$F(s) = \int_K^x \frac{dx}{\sigma x} = \frac{1}{\sigma} \ln \frac{s}{K}$$

The new variable F is the return of the underlying asset measured relatively to the price to strike ratio. Whenever the diffusion parameter $z(s,t)$ has a more complex form, the new variable F is a generalization of returns.

Substituting the new variable F and applying the Ito's lemma we immediately obtain:

$$dF = \left(\frac{\mu(s,t)}{\sigma} - \frac{\sigma}{2} \right) dt + dB_t$$

which in the case of a risk-neutral process ($\mu = r$) has the familiar form:

$$dF = \left(\frac{r - 0.5\sigma^2}{\sigma} \right) dt + dB_t$$

This process has a known drift and a random component arising from the standard Brownian motion. Thus its density function can be calculated easily (a Normal distribution with variance t and mean $t(r - 0.5\sigma^2)/\sigma$). Therefore, it is easy to find any quantile and the corresponding VaR, first in the new variable F and then in the original variable s .

This approach is based on a monotonic change of variables. It leads to bounds on almost all of the quantitative characteristics of the stochastic process. Numerous examples of bounds on probabilities, deltas, and VaRs can be found in Grundy, Wiener [1996]. Here we give without a proof only one result related to VaR (Lemma 4).

Suppose that the true process for the underlying asset starts at level s at time t and follows the diffusion $ds = \mu(s,t)sdt + \sigma(s,t)sdB_t$, with the following bounds on its major parameters: $\underline{\mu} \leq \mu(s,t) \leq \bar{\mu}$ and $\underline{\sigma} \leq \sigma(s,t) \leq \bar{\sigma}$. Let the inflation to be i and denote by

$$M = \max \left[\frac{\bar{\mu} - i}{\underline{\sigma}}, \frac{\bar{\mu} - i}{\bar{\sigma}} \right] \text{ and } m = \max \left[\frac{\underline{\mu} - i}{\underline{\sigma}}, \frac{\underline{\mu} - i}{\bar{\sigma}} \right].$$

Then using the standard inverse of the normal density ($\Phi(k)$ – the 1- $k\%$ quantile) we obtain:

$$\text{if } \frac{\partial z}{\partial s} \geq 0 \text{ for all } s \text{ and then}$$

$$\text{VaR}(k) \geq \begin{cases} 1 - \exp\left((i + M\underline{\sigma})(T-t) - \Phi(k)\underline{\sigma}\sqrt{T-t}\right) & \text{if } M\sqrt{T-t} \leq \Phi(k) \\ 1 - \exp\left((i + M\bar{\sigma})(T-t) - \Phi(k)\bar{\sigma}\sqrt{T-t}\right) & \text{if } M\sqrt{T-t} > \Phi(k) \end{cases}.$$

if $\frac{\partial \sigma}{\partial s} \geq 0$ for all s and then

$$\text{VaR}(k) \geq \begin{cases} 1 - \exp\left((i + (M - 0.5\underline{\sigma})\underline{\sigma})(T-t) - \Phi(k)\underline{\sigma}\sqrt{T-t}\right) & \text{if } (M - 0.5\underline{\sigma})\sqrt{T-t} \leq \Phi(k) \\ 1 - \exp\left((i + (M - 0.5\bar{\sigma})\bar{\sigma})(T-t) - \Phi(k)\bar{\sigma}\sqrt{T-t}\right) & \text{if } (M - 0.5\bar{\sigma})\sqrt{T-t} > \Phi(k) \end{cases}.$$

if $\frac{\partial \sigma}{\partial s} \leq 0$ for all s and then

$$\text{VaR}(k) \leq \begin{cases} 1 - \exp\left((i + (m - 0.5\bar{\sigma})\bar{\sigma})(T-t) - \Phi(k)\bar{\sigma}\sqrt{T-t}\right) & \text{if } (m - 0.5\bar{\sigma})\sqrt{T-t} \leq \Phi(k) \\ 1 - \exp\left((i + (m - 0.5\underline{\sigma})\underline{\sigma})(T-t) - \Phi(k)\underline{\sigma}\sqrt{T-t}\right) & \text{if } (m - 0.5\underline{\sigma})\sqrt{T-t} > \Phi(k) \end{cases}.$$

Various numerical results based on simulations and using this lemma are presented in Table II in Grundy, Wiener [1996].

This change of variables can be used for more than one variable. Let us demonstrate this in the case of two variables (the general approach is similar):

$$\begin{cases} ds = v(s, q, t)dt + z(s, q, t)dB_t \\ dq = v(s, q, t)dt + w(s, q, t)dB_t \end{cases}$$

Then one can perform the change of variables simplifying the volatility to a constant form. In other words we are looking for a new pair of variables $S(s, q, t)$ and $Q(s, q, t)$ such that the new dynamic has a simple form. Using the Ito's lemma one can write the new variables as:

$$\begin{cases} dS = S_1 ds + S_2 dq + S_3 dt + 0.5S_{11} (ds)^2 + S_{12} (dq)^2 + 0.5S_{22} (dq)^2 \\ dQ = Q_1 ds + Q_2 dq + Q_3 dt + 0.5Q_{11} (ds)^2 + Q_{12} (dq)^2 + 0.5Q_{22} (dq)^2 \end{cases}$$

Thus in order to get a simple form for the random part of the new dynamic we have to check

$$\begin{cases} dS = \dots dt + \left(S_1(s, q, t)z(s, q, t) + S_2(s, q, t)w(s, q, t)\right)dB_t \\ dQ = \dots dt + \left(Q_1(s, q, t)z(s, q, t) + Q_2(s, q, t)w(s, q, t)\right)dB_t \end{cases}$$

We have to solve separately the two (independent) partial differential equations of the first order (of Pfaff type) where the $a(t)$ and $b(t)$ are the new deterministic diffusion parameters:

$$\begin{aligned}
S_1(s, q, t)z(s, q, t) + S_2(s, q, t)w(s, q, t) &= a(t) \\
Q_1(s, q, t)z(s, q, t) + Q_2(s, q, t)w(s, q, t) &= b(t)
\end{aligned}$$

The solution is based on a general solution of a system of ordinary differential equations:

$$\begin{cases}
\frac{ds}{dt} = z(s, q, t) \\
\frac{dq}{dt} = w(s, q, t) \\
\frac{dS}{dt} = a(t)
\end{cases}$$

There are many methods of how to solve such a system and we will not investigate it here further.

Another example of the analytical approach to VaR is through the delta or delta-gamma approach. This method is useful when there is an analytic formula (like Black-Scholes or the CEV model for options). In this case in order to measure a risk associated with a complex derivative (an option in this case) one can use a linear (delta approach) or a quadratic (delta-gamma approach) to measure the price distribution density of the derivative. For a detailed discussion and some implementations see Schaefer [1997] and Duffie and Pan [1997].

V. Backtesting

The Basle standard requires backtesting – a procedure in which one check (a posteriori) how often the actual losses have exceeded the level predicted by VaR. As soon as a 99% confidence interval and the 10 day time horizon are used, there should not be too many cases in when the actual losses are greater than the predicted ones.

There are three zones. If during the last year (approximately 250 business days) there are four or less exceptions (losses that exceed the VaR level), the model is said to be in a green zone and it is acceptable. If there are five to nine exceptions it is in the yellow zone and certain actions (such as increase of the safety multiplier from 3 to 3.65) are recommended³. When there are ten or more exceptions the whole model should be revised.

³ An interesting explanation of this multiplier and the assumptions about the probability distribution can be found in Stahl [1997].

This mechanism prevents banks from setting the VaR too low. However when current capital reserves are high the bank can safely set an upper bound of VaR instead of trying to calculate a precise value.

References

Alexander, Carol and C. Leigh, On the Covariance Matrices Used In VaR Models. *Journal of Derivatives* 4 (Spring), p. 50-62, 1997.

Amendment to the Capital Accord to Incorporate Market Risks, Bank for International Settlements, 1996.

An Internal Model-Based Approach to Market Risk Capital Requirements, Basle Committee on Banking Supervision, 1995.

Duffie, D. and J. Pan, An Overview of Value at Risk. *Journal of Derivatives* 4 (Spring), 7-49, 1997.

Grundy Bruce D. and Zvi Wiener. The Analysis of VAR, Deltas and State Prices: A New Approach. Working paper of The Rodney L. White Center for Financial Research, The Wharton School, 11, 1996.

Ingersoll Johnathan E. Jr., Valuing Foreign Exchange Derivatives with a Bounded Exchange Rate Process. *Review of Derivatives Research*, vol.1 , p. 159-181, 1996.

Jackwerth Jens Carsten, Rubinstein Mark, Recovering Probability Distributions from Option Prices. *The Journal of Finance*, vol. 51, no. 5, December 1996, p. 1611-1631.

Linsmeier, Thomas and Neil Pearson, Risk Measurement: An Introduction to Value at Risk. Working paper of the University of Illinois, 1996.

Marshall, Chris and Michael Siegel, Value at Risk: Implementing A Risk Management Standard. *Journal of Derivatives*, vol. 4 (Spring), 91-110, 1997.

Schaefer Stephen, Management of Non-Linear Market Risk. Working paper, London Business School, 1997.

Stahl Gerhard, Three Cheers. *Risk*, vol. 10, no. 5, May 1997, p. 67-69.

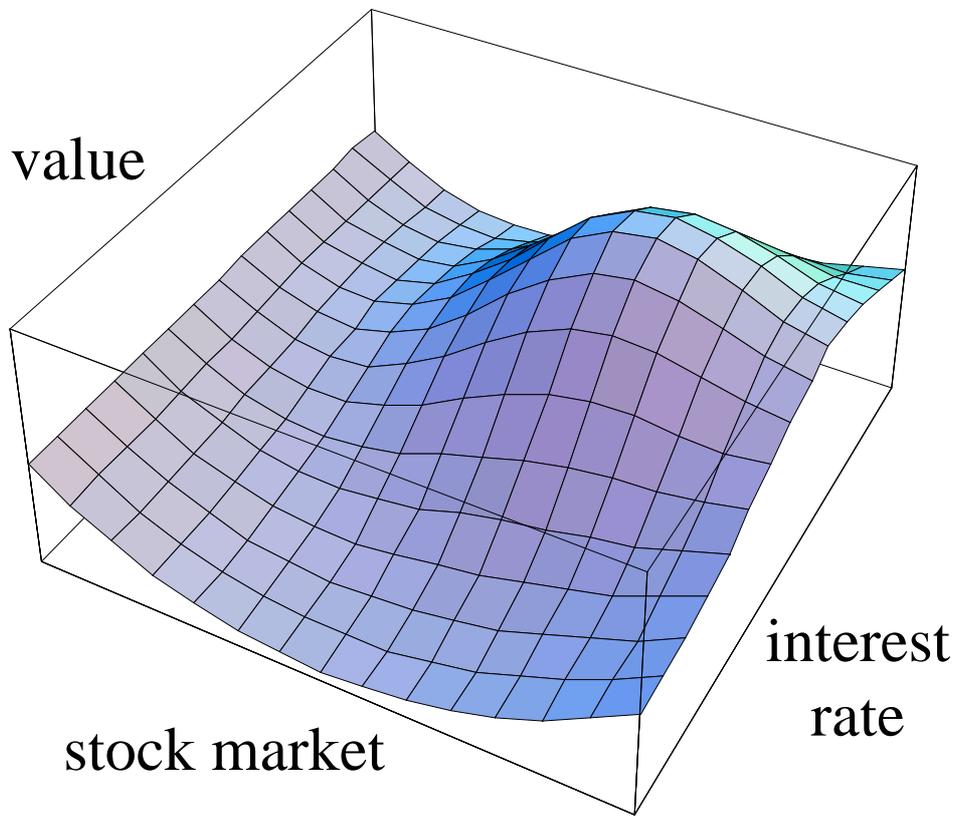


Figure 1. Value of the portfolio as a function of market parameters (which can be correlated).

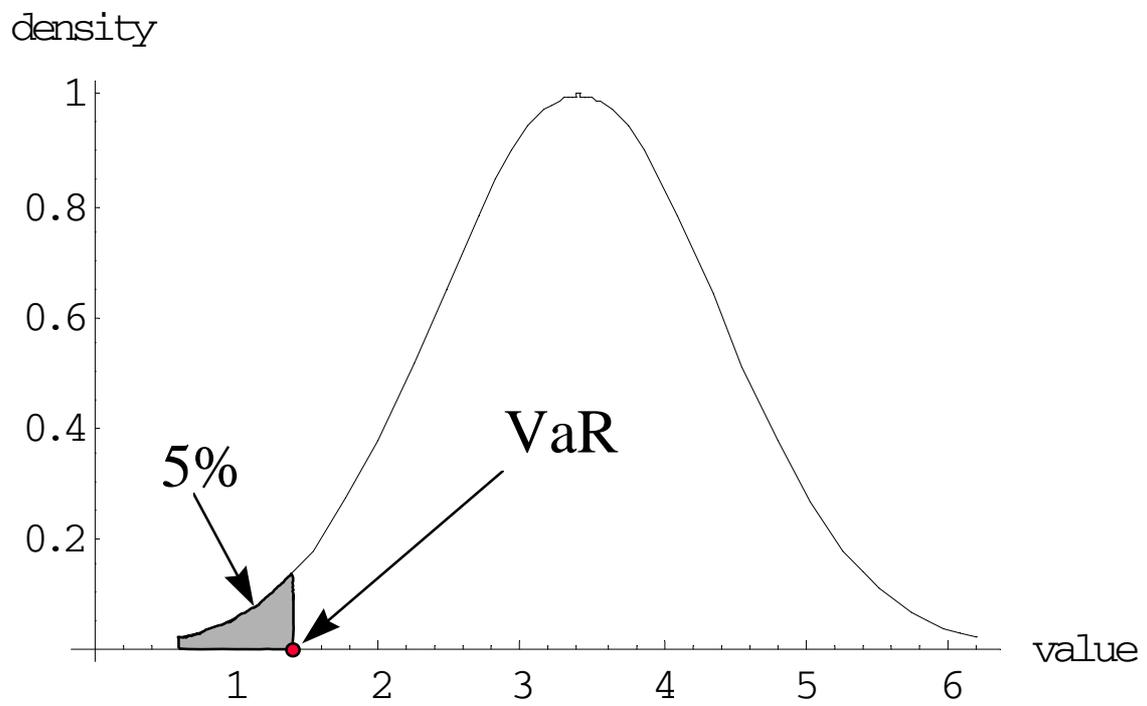


Figure 2. 5% VaR is defined on the probability density function.

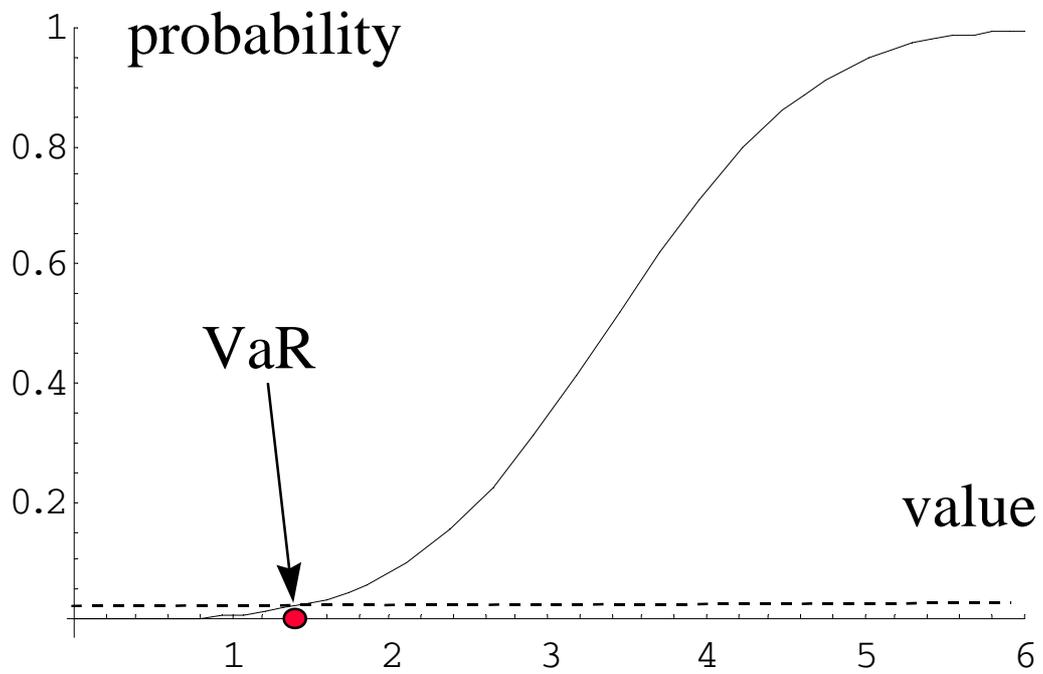


Figure 3. 5% VaR is defined on the cumulative probability function as a 5% quantile.

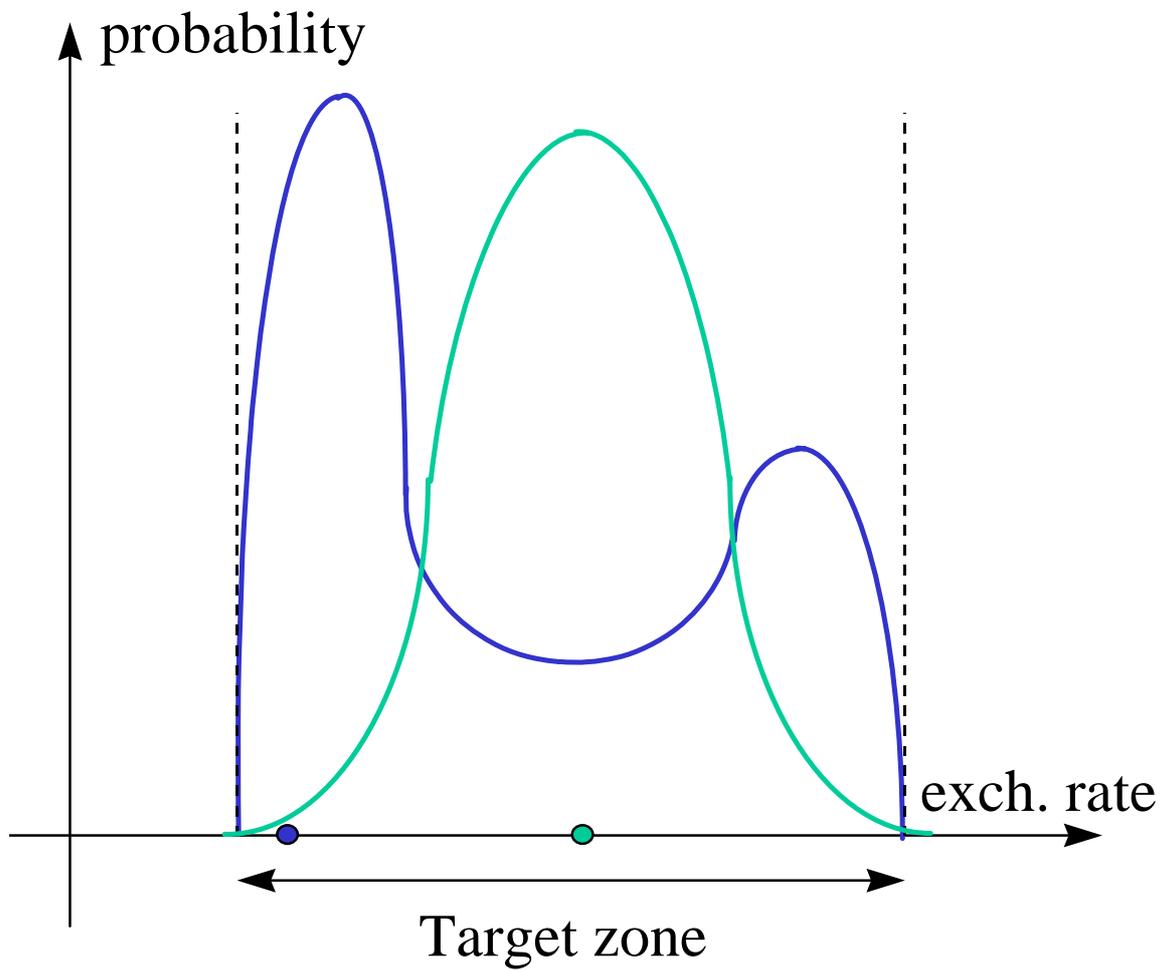


Figure 4. Probability density of the exchange rate under a credible target zone regime. The normality assumption is reasonable when the current exchange rate is close to the middle of the zone. When the exchange rate is close to the boundary the probability distribution is very far from being normal.