Market Risk

A. Defining the Problem

What is market risk? It is the risk that a financial institution (or firm more generally) will lose money if market prices change.

The key sorts of market prices that affect the value of a FSFs portfolio, in order of importance, are:

Interest rates
Foreign Exchange rates
Equity Prices
Commodity Prices

Market risk is very important for FSFs that operate as dealers in derivatives markets. These firms are in the business of allowing less sophisticated firms to use instruments such as swaps to hedge interest rate risk. Thus, they stand ready to take either side of a swap with a counterparty. To be successful at this, these dealers need to be able to measure their exposure to market risks and reduce those exposures if they become too large.

B. Market risk management challenges

There are two key issues for market risk management.

First, how can a FSF measure its overall market risk in a portfolio? Measurement is hard because a complex portfolio will be composed of many positions and be exposed to many market prices. The firm will have positions in swaps, bonds, equities, options, forwards and futures. The market value of these positions will depend on many market prices: the yield curves in many countries, F/X rates, equity prices, etc. Moreover, the goal so to construct a single benchmark to measure the risk of the whole portfolio. Remember, portfolio theory says that the most important factor driving the risk of a portfolio is the correlation between various risks.

The second key issue is how an FSF should manage its market risk. Clearly, a precondition for rational management is a reliable measure of risk, and we will focus most of our attention on the details of how this is done in practice. However, you should consider the following **management questions**:

> Should market risk (or risks generally), be managed at the corporate level or at the business line level?
>
> How should limits be imposed on businesses or traders? Should compensation of business line managers and traders be tied to risk?

C. Measuring Market Risk with VAR

The standard tool in the market now used to measure market risk is called VAR, which stands for "Value-at-Risk." The VAR of a portfolio equals the maximum amount that may be lost on that portfolio over a given holding period and with a given probability (or confidence level).

Example. Suppose the VAR with a 1-day holding period and a 95% confidence interval is $10 million. What does this mean? It means that 95% of the time the firm will lose less than $10 million dollars on its portfolio if it freezes the current positions for exactly 1 day.

Question: What happens to VAR if you increase the holding period? What happens if you increase the confidence level?

> Key point: VAR is only defined with respect to the holding period and the confidence level. The number is TOTALLY MEANINGLESS without knowing these two things.

VAR is a statistical measure of the riskiness of a portfolio. Thus, I must digress to be sure everyone has the basic statistical concepts before going further.

A probability distribution function (pdf) assigns a likelihood (or probability) to all possible outcomes of a random variable.

Example: What is the pdf for a random variable that takes the value 0 if you flip a coin and receive heads (H) and 1 if you receive tails (T)? What is the pdf for a random variable equal to 0 if you flip a coin twice and get HH; 1 if you get HT (or TH); 2 if you get TT? These are examples of discrete pdfs where the value is a probability.

Once you have the pdf, you can compute measures of the central tendency of the random variable, and measures of its dispersion. We will focus on the expected value and variance.

For discrete pdf:

$$E(X) = \sum_s X(s)P(X(s))$$

$$\sigma^2(X) = \sum_s [X(s) - E(X)]^2 P(X(s)) = E[X - E(X)]^2$$

where X(s) is the realization of the random variable X in state s. In the coin flip example, a state would be “the coin comes up heads” and X would be the number 0. P is the pdf. It assigns a probability to each realization of X. The summation is taken over all possible states. In the single coin toss example, there are two states, heads and tails.

For continuous pdfs.

The definitions are similar, except you integrate instead of sum:

$$E(X) = \int_{-\infty}^{+\infty} XP(X)dX$$

$$\sigma^2(X) = \int_{-\infty}^{+\infty} [X - E(X)]^2 P(X)dX = E[X - E(X)]^2$$

when the random variable is continuous, the pdf P(X) is no longer a probability, since the probability that the random variable takes on any single number is infinitesimal. Instead, P assigns a likelihood to each possible realization of X. What you can do with a continuous pdf is integrate over a region to get a probability. For example, the probability that the random variable X is less than the number 14 is:

$$\Pr ob(X < 14) = \int_{-\infty}^{14} P(X)dX$$

The covariance between 2 random variables measures how much the two move together. Here is the definition:

$$Cov(X,Y) = E[(X - E(X)][Y - E(Y)]$$

Note that I did not write this in terms of the pdf. To do that, you need the joint pdf that describes the likelihood of all pairs of realizations of X and Y. The message you should be getting is that the pdf gets VERY COMPLICATED AND MESSY when you have lots of random variables, which you do in the market risk problem. The relevant random variables will be the many market prices that affect the portfolio.

With the definitions of expected value, variance and covariance, it is easy to prove the following, which you should know from portfolio theory:

$$E(AX + BY) = AE(X) + BE(Y)$$

$$\sigma^2(AX + BY) = A^2\sigma^2(X) + B^2\sigma^2(Y) + 2ABCov(X,Y)$$

where A and B are constants and X and Y are random variables. Another way to express the covariance of the random variables is in terms of their correlation coefficient.

$$Cov(X,Y) = \sigma(X) \times \sigma(Y) \times \rho(X,Y)$$

So,

$$\sigma^2(AX + BY) = A^2\sigma^2(X) + B^2\sigma^2(Y) + 2AB\sigma(X)\sigma(Y)\rho(X,Y)$$

A mathematical definition of VAR

The VAR can be defined with respect to the pdf for the change in the value of the portfolio over the holding period. Let x be a random variable equal to the change in the value of the portfolio over the holding period, and P(x) be its pdf. Then the VAR with a confidence level of α solves the following equation:

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

Here is a graphical representation of this definition (more bad art...):

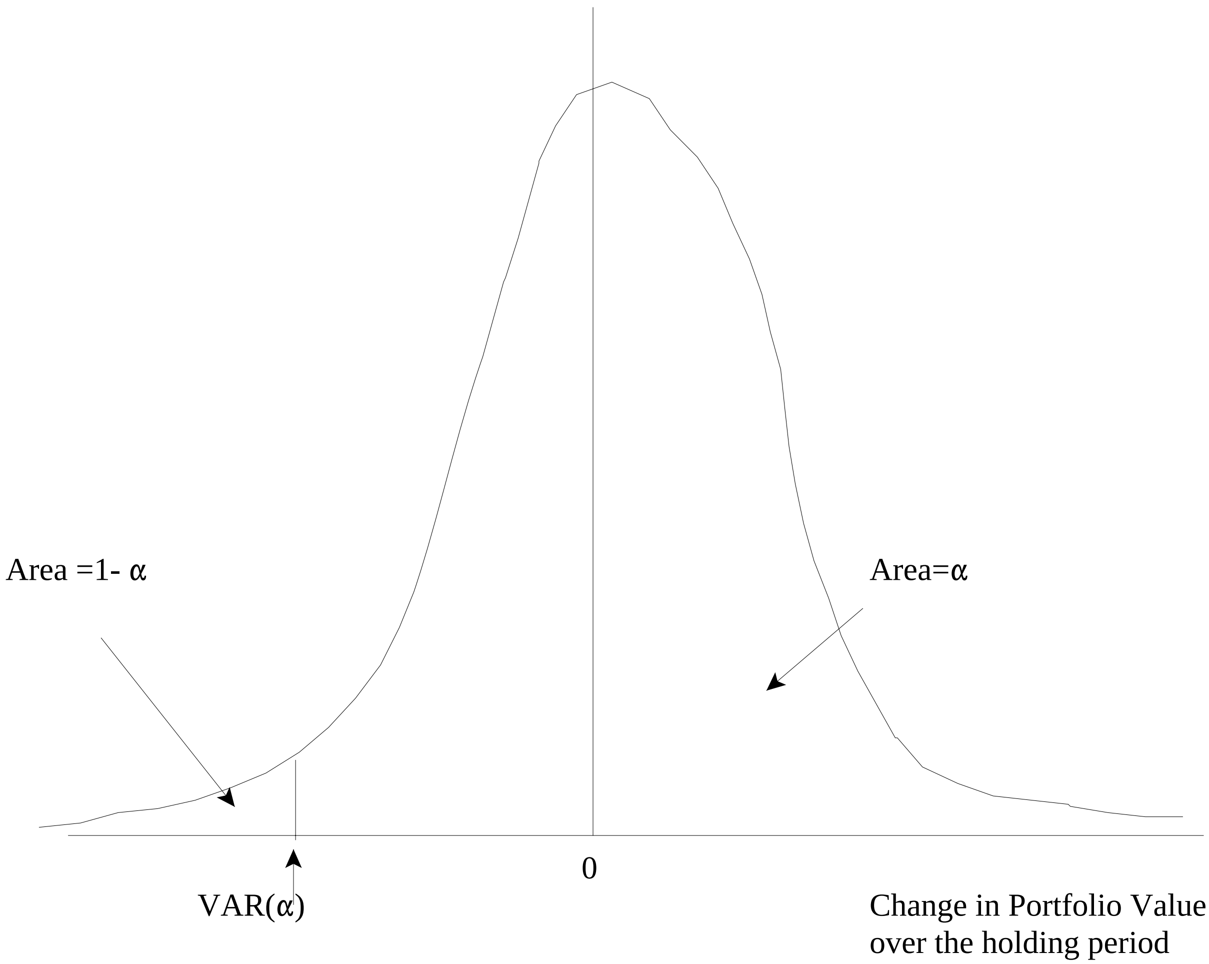

A few words about the normal distribution

There is an incredibly useful pdf called the "normal distribution." Normal distributions are cool because they occur naturally. For example, if you take the average of a bunch of independent draws from any random variable and think of this average as a new random variable, it will have a pdf that is approximately normal. The reason is complicated, so I won't go into it. Moreover, the normal pdf has 2 incredibly useful properties:

> Fact 1 about normal pdf: To construct the normal pdf, all you need to know about the random variable is its expected value and its variance. This means that if a portfolio has returns that are normally distributed, a natural measure of risk is the variance (or standard deviation).
>
> Fact 2 about normal pdf: Linear combinations (AX+BY) of normally distributed variables are also normally distributed.
>
> Fact 2 means that if we have a bunch of components of a portfolio with normally distributed returns, then the portfolio itself has a normal distribution. Also, we know how to compute its expected value and variance from the formulas above, as long as we can get estimates of the expected values, variances and covariances.

D. How to Compute VAR

Step 1

> The first step is to define what market prices affect the value of your positions. This is the easy part. If there are *n* prices, we will name then $S^1,...,S^n$. These are the risk factors.

Step 2

> Once we have a set of risk factors, we need to have a valuation model. This defines the value of the portfolio as a function of the risk factors. Once we know the risk factors and the valuation model, we can define $V(S^1,...,S^n)$ as the value of the portfolio as a function of *n* market prices (or risk factors).

Example. If the portfolio were a zero coupon bond with face value of $100 and maturity of t, how would we would write V?

Step 3:

Now that we have $V(S^1,\ldots,S^n)$, we want to construct the pdf for the change in V over the holding period. Once we have this, we can construct the VAR with a given level of confidence.

This is the hard part! There are three approaches that are typically used. Two of these fall under the category of simulation. The third uses an approximation to the true pdf for V based on the normal distribution. The simulation methods are conceptually easy but computationally hard; the normal method is harder to understand analytically, but much easier to use in practice.

Sophisticated banks that do lots of trading use simulation; everyone else uses the normal approximation.

Simulation Models

Let us write the change in the portfolio as follows:

$$\Delta V = V(S^1+\Delta S^1,\ldots,S^n+\Delta S^n) - V(S^1,\ldots,S^n)$$

To make things concrete, suppose we have a portfolio with a single 1-year zero bond paying $100 million and interest rates today are 3.4%. Then,

$$\Delta V = \frac{100}{1.034 + \Delta r} - \frac{100}{1.034}$$

In historical simulation, we go out and measure actual changes in the risk factor(s) over the holding period during some estimation period. We then simply compute ΔV over the estimation period. This defines an empirical pdf. To compute VAR with a given confidence level (say 95%), we just use that pdf.

Example . Suppose we use a 30-day historical estimation period to compute the VAR on the zero-coupon bond portfolio with a 1-day holding period. Here are interest rates on 1-year T-bills during last August. Use historical simulation to construct the VAR with 95% confidence. (Use a spreadsheet.)

Date	r	Δr	ΔV
07/31/2001	3.53	N/A	N/A
08/01/2001	3.56	0.03	??
08/02/2001	3.57	0.01	??
08/03/2001	3.57	0	??
08/06/2001	3.56	-0.01	??
08/07/2001	3.56	...	
08/08/2001	3.46		
08/09/2001	3.48		
08/10/2001	3.45		
08/13/2001	3.43		
08/14/2001	3.46		
08/15/2001	3.47		
08/16/2001	3.43		
08/17/2001	3.39		
08/20/2001	3.44		
08/21/2001	3.41		
08/22/2001	3.44		
08/23/2001	3.46		
08/24/2001	3.48		
08/27/2001	3.51		
08/28/2001	3.46		
08/29/2001	3.44		
08/30/2001	3.38	...	
08/31/2001	3.41	0.03	??

The other simulation method is called Monte Carlo Simulation. In this approach, we *assume a distribution* for the change in the risk factors. For example, we might assume that changes in interest rates are normally distributed. We then use data over the estimation period to compute the mean and variance for changes in interest rates. This gives our estimated pdf for the risk factor. We can now sample changes in interest rates from the estimated pdf to construct a pdf for the change in the value of the portfolio.

The advantage of Monte Carlo, over Historical, simulation is that you can consider more scenarios. For example, if you do Historical simulation with a 30-day estimation period, you only get 30 realizations of the change in your portfolio. That is not many. Monte Carlo lets you explore the "tails" of the distribution.

Warning about Monte Carlo: You must make an assumption about the distribution of the risk factors to do Monte Carlo. If this is wrong, you will have a misleading model.

Delta-Normal Method

The delta normal method makes 2 assumptions. First, we assume that the portfolio is a linear function of the risk factors and, second, that the percentage change in the risk factors are normally distributed. With these two assumptions, it is very easy to construct VAR. All you need are the variances and covariances of the underlying risk factors over the holding period.

Here is the idea. Say there is one risk factor S. S_0 is the initial value, and S_1 is its value at the end of the holding period.

$$V(S_1)-V(S_0)=\sim\Delta(S_1-S_0) = \Delta\ S_0\ [\ (S_1-S_0)/S_0\ ]$$

For small changes in S, this approximation becomes exact. That is what is known as a "Taylor Series Approximation." So, for small changes:

$$dV = \Delta S(dS/S)$$

where $\Delta=\delta V/\delta S$. Since dV is a linear function of dS/S (which has a normal distribution by assumption), dV also has a normal distribution with a standard deviation that is equal to ΔS times the standard deviation of the risk factor (dS/S). Let σ equal the standard deviation of the risk factor. We can estimate σ using data on percentage changes in the risk factor over the holding period. Under these conditions, VAR is a fixed multiple of the standard deviation of dV:

$$VAR(99\%) = 2.33\Delta S\sigma$$

$$VAR(95\%) = 1.65\Delta S\sigma$$

A word on the portfolio's expected return

Under the preceding approach, we implicitly assumed that the expected change in the value of the portfolio over the holding period is zero. This is commonly assumed because the expected return on a portfolio over a very short holding period such as 1 day is very small relative to the standard deviation of the change in the portfolio value. Over longer holding periods, however, the expected return becomes more important. This is true because the standard deviation increases with the square root of time, whereas the expected return increases with time (i.e. the expected return over 10 days is 10 times as big as the expected return over 1 day). To adjust for a non-negligible expected return, you would simply subtract it from the VAR formula above.

Delta Normal with more than one risk factor

Delta Normal also works with more than one risk factor. Let's do it with 2. Again, we use the Taylor Series to justify the following linear approximation:

$$dV = \Delta^1 S^1 (dS^1/S^1) + \Delta^2 S^2 (dS^2/S^2)$$

Again, dV is a linear combination of 2 normally distributed variables, so it too is normally distributed. Thus, the VAR is proportional to the standard deviation of dV. The delta normal approaches allows you to decompose risk into different components. In this example:

$$VAR^1(95\%) = 1.65\Delta^1 S^1 \sigma^1$$
$$VAR^2(95\%) = 1.65\Delta^2 S^2 \sigma^2$$

To compute the VAR for the whole portfolio, you aggregate these VARs for each risk factor using the rule for aggregating variances:

$$PortfolioVAR = \sqrt{(VAR^1)^2 + (VAR^2)^2 + 2\rho VAR^1 VAR^2}$$

where ρ is the correlation coefficient between the two risk factors.

Example. Here is a simple example with 2 risk factors. Suppose you have a portfolio with 100 shares of IBM, and each share is worth $115, and you hold 300 shares of Microsoft, with each share worth $25. You have decided to use delta normal with 2 risk factors: the daily return on IBM and the daily return on Microsoft. The standard deviation of daily returns on IBM is 2.5%, and the standard deviation of the daily return on Microsoft is 4%. The correlation between the two is 0.5.

Compute the 1-day VAR with 99% confidence levels for each risk factor and for the whole position. What would happen if the Microsoft position were short rather than long?

E. VAR Pitfalls

1. Confidence Level. VAR only measures the worst loss with a given level of confidence. If the 1-day VAR with 95% confidence is $10 million, that means that you lose more than $10 million 5% of the time. VAR tells you nothing about how big the losses could be on those days! The losses could be very big, as we will see when we study the hedge fund Long-Term Capital Management (LTCM).

One solution to this problem is to augment VAR with "stress testing" or "scenario analysis."

> In stress testing, the risk manager attempts to construct VAR for higher and higher confidence levels. An additional useful approach is to compute the expected losses during the days when the portfolio does worse than VAR.
>
> In scenario analysis, the risk manager creates a bad scenario and then tries to construct losses under that scenario. The idea is to identify things that make it hard to sleep at night and see how bad the losses would be. For example, a large insurance company contemplates the effect of a major earthquake in Los Angeles hitting simultaneously with one in Tokyo. Would the firm survive? Of course, the risk manager may not contemplate scenarios that actually occur, such as both WTC towers falling (reports claimed that the WTC was insured against terrorist destroying 1 but not both towers).

Another scenario that always ought to be considered is one in which the correlations of positions goes up. Portfolio theory teaches that adding independent risks to a portfolio reduces total variation. This is great, but what if there is a market shock (like last week) that makes all positions move together?

2. Holding Period. Remember the VAR is only defined with respect to a holding period. The longer the holding period, the greater the risk. The key factor is liquidity; if the portfolio is highly liquid, a short holding period is OK because the firm could dump its risky positions quickly. We will again see the importance of this issue when we study the LTCM fiasco.

One simple, back-of-the-envelope way to consider the effects of the holding period is to use the "square root of t" rule. In the delta-normal model, VAR is proportional to the standard deviation of the change in the value of the portfolio over the holding period. If changes in the value of the portfolio are independent, then the variance of changes increases with t and the standard deviation increases with $\sqrt{t}$.

Example: Compute the VAR for the portfolio above over a 1-week, 1-month and 1-year holding period.

3. Estimation Period. What is the right estimation period? In the example above, I had you use a very short period of 30 days. The advantage of using the recent past is that it is more likely to

be reflective of the actual statistical process generating the price data. However, with fewer observations, it is harder to get a good estimate of the statistics you want. Also, you may leave out important days of dramatic volatility in the markets, such as October 19, 1987.

If you use a long estimation period, you solve some of these problems but introduce a new one: what if the distributions generating price changes are not constant over time?

4. Non-normality. In the delta-normal method, the assumption of normality in the price change data is key. However, many financial time series have "fat tail" distributions relative to the normal. That means that the probability of very large outliers is much higher than what the normal distribution says. A dramatic example is the 20% drop in US equities on October 9, 1987. That is a decline of at least 10 standard deviations, which is all but impossible under a normal distribution.

Another problem is the use of the Taylor Series Linearization. This is only a good thing to do for small changes. Thus, if you have a portfolio that is a very non-linear function of the underlying risk factors, and you are considering a holding period of more than 1 day, the delta-normal model may behave very badly.

An example of a non-linear position is an option. The problem occurs because a non-linear function of a normal random variable does NOT have a normal distribution.

5. Validation. How do you know the VAR model is working? If you compute VAR with 1-day holding period and 95% confidence, you would expect about 12.5 days per year to lose more than VAR. What if you lose more than VAR 20 times during a year. Does this mean that VAR is bad? The answer is that it is very hard to be sure. The problem gets worse with VAR measured at 99% confidence level.

What makes this issue even more vexing is that the models are constantly being updated and improved, so you are trying to validate a moving target.

F. The Bottom Line on VAR

As you can see, there are many pitfalls with VAR. However, it is very useful to have a way to assess the risk of an entire portfolio, or portions of a portfolio. VAR can be used to determine how much capital a FSF should set aside. For example, if the FSF determines that it wants to be sure to survive the worst annual losses 99% of the time, then it wants to have capital equal to VAR with 99% confidence over a 1 year holding period.

So, risk managers should (and most do) use VAR; but they should use it with care. Don't get complacent if VAR is low, as the managers at LTCM did. Supplement VAR with "stress testing" and scenario analysis.