Measuring market risk: the heart of the matter

Susan Thomas
IGIDR, Bombay
http://www.igidr.ac.in/∼susant

January 28, 2011

The problem

What is risk?

- Risk is realising a valuation that is different from what is expected.
- The direction of the deviation is not as important as it's magnitude: a realised value very different from that expected is **worse** than a realised value that is similar to what was expected.
- For example, getting an average $+5\%$ on a long position is **more risk** than getting an average -5 bps on the same position.

Traditional view on measuring risk

- In finance, risk is sometimes used interchangeably with volatility, or variance.
- We assumed that financial returns had a character that was similar to numbers from a normal distribution:
  - Values between $+\infty$ to $-\infty$.
  - Typically with equal probability of an up-tick or a down-tick.
- This lead to using the statistical measurement of deviation for a normal distribution, which is the variance σ^2, or standard deviation σ.
- This was the state of the art for a while.
- This changed from the '80s when several crisis in the US banking sector forced a fresh look at risk measurement.

Why revisit the issue of risk measurement?

- People running a business have always “guesstimated” the level of risk of a project that they were entering into. Today seems to have an emphasis on numerical values for how much risk. Why?
- There are several reasons:
  1. Old problems of good governance:
     - Shareholders and the Board of Directors assessment of risk as opposed to that of the management.
     - Skewed incentives for high risk-taking by the management. There are high rewards upon success but capped losses upon failure.
  2. New reasons:
     - External compulsions: today’s regulators demand it.
     - Incentive for high “alpha”: rewards must be calculated net of risk. We need to have a fine measure of risk.

Regulatory compulsions for financial firms

Basel norms

- In the late eighties, there was a lot of cross-border lending particularly by the Japanese banks.
  Japanese banks grew enormously and gathered market share; Western banks complained about Japanese banks being regulated badly.
- Basel I was an attempt to standardise the regulation governing the global banking industry.
- The heart of the Basel I norms defined "minimum required equity capital", i.e. an attempt to contain leverage. This was because banking is especially risky business:

  India: Food 1.1, Machinery 0.6, Automobiles 1.1, Auto Ancillaries 1.19, **Banking 17.6**.

  US: Manufacturing 0.25-0.35, Utilities 1.4-1.5, Trade 0.3-0.4, **Banking 15**.

Basel I, contd.

- Required equity capital was a single number calculated as a fraction of the risk weighted assets (RWA).
- RWA $= w_1 x_1 + w_2 x_2 + \ldots$, where x_1 was corporate exposure, and $w_1 = 1$.
  The weights for all the other classes of assets was set at less than 1.
- The main focus appeared to be on addressing credit risk.
- The minimum equity requirement was set through a minimum Capital Adequacy Ratio (CAR), at typically 8% of RWA.

What was right about the Basel Capital Accord

- The CAR requirement did reduce the extremely high levels of leverage in the banking industry.

What was wrong about the Basel Capital Accord

- The calculation of RWA was incorrect.
  Risks in the banking portfolio are not linear.
- Assets were classified on very broad lines.
  (Eg. OECD government bonds.)
- The focus on credit risk gave banks incentives to find new ways of bearing risk.
  (Eg. higher exposure in interest rate risk, OTC derivatives.)
- Ignored the problem of non-transparency (Eg. loans, OTC derivatives, OTC trading.)
- Ignored differences between countries. (If 8% works for the OECD, what is correct for India?)
- Negative consequences:
  - Even though these were broad recommendations, they became rigid in the hands of weak banking regulators.
  - The focus shifted from taking risks with a clear understanding of the returns, to blindly using BIS rules.

Basel II

- An attempt to move away from linear rules of thumb.
- Some of the implementation involves
  - Trying to improve upon the linear formula.
  - Reliance on credit ratings.
  - Exploit “internal models” of risk measurement in banks.
  - Taking more interest in *incentives* - of banks, of securities markets.
- This is still ‘playing the game of Basel I’ - but trying to find a better formula for equity capital.
- This translates into a better way to measure the financial risk of the bank.

Corporate governance

Separation of risk taking and risk measurement

- The person who buys a position always believes in it: "nothing can go wrong" overconfidence.
- Risk measurement is about a cold blooded + objective judgment about how much risk you have.
- It is crucial to separate out the person who adopts the position, from the person who measures the risk.
- The best positions can go wrong.

Infosys price around April 2003

- A pessimistic growth report on Infosys Technologies in April 2003 caused a 36% drop in the price.

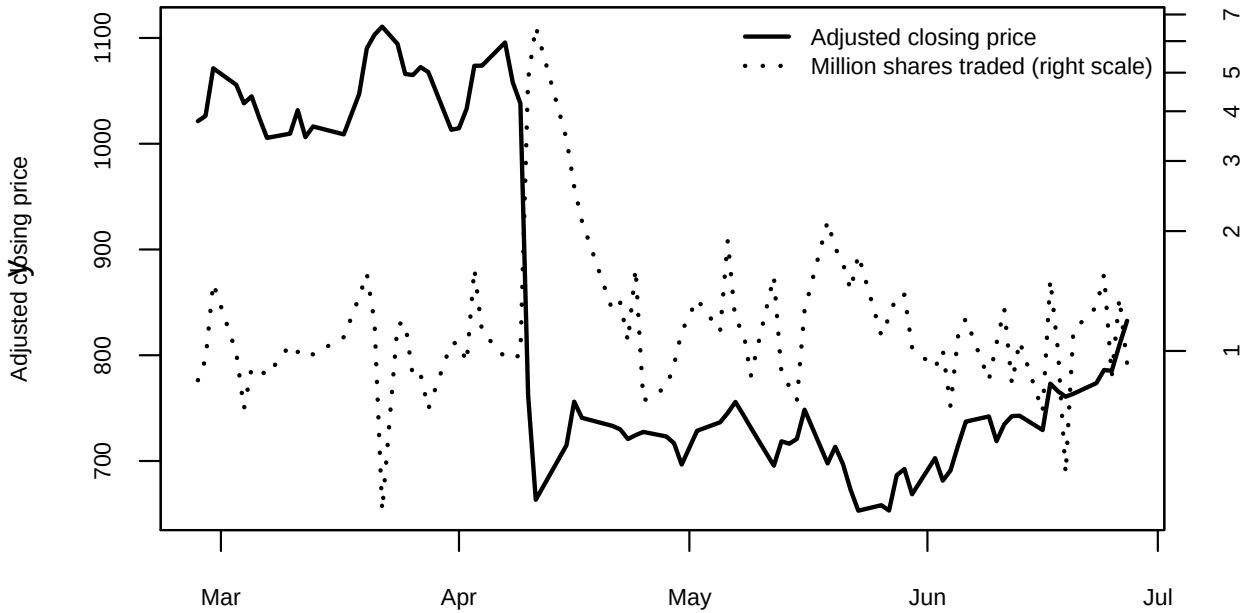

- Need to fear the market.

A new framework for risk measurement

Risk measurement

- The most practical measure of risk is to answer the following question:
  "If things go wrong, how much will I lose?"
- What can go wrong for a financial firm?
  1. Prices can drop
  2. Market impact cost can spike - so you get into a position and it seems to do well but when you sell off the impact cost burns return
  3. Operations risk - outright fraud within the organisation.
- Ideally we want risk measurement which encompasses all this so as to give top management an accurate picture of what is going on.
  To give top management an answer to the question "If things go wrong, how much will I lose?"
- One approach to answer this question, emerged after a series of financial firms failed in the US in the nineties. That is **Value at Risk**, or **VaR**.

Risk measurement, step 0 – Valuation

Mark to market (MTM)

- Risk management is what happens in the future.
  But before we discuss the future, we need to understand what has **already** happened.
- The best option to measure the present: Marking to Market (*MTM*).
  MTM is recognising what has happened.
- Accounting data is incorrect.
  What you want to know is **not** what you paid for a security, **but** what you will get if you try to sell it.
- *Ideally*, we should use (market price - impact cost) for valuation.
  Impact cost is a measure of liquidity. It varies (liquidity risk) and may not always be observable.
- *At least*, we should use market price.

Problems with marking a portfolio to market

- Regulations sometimes allow you to NOT do MTM at various situations.
  For example, the Hold-To-Maturity portfolio of banks.
  This is bad practice. 100% MTM should be the religion.
- The real bottleneck very often is a lack of IT systems and databases that are plugged into real-time price feeds.
  For example, many firms do not have the information properly organised so as to do a full MTM frequently.
- Before risk management systems, there **has** to be an MTM system.

The alternative – Mark to Model (MTMod)

- Often the securities held by a financial firm is either
  1. Not traded – such as loans
  2. Is illiquid – OTC derivatives / GOI bonds
- In such situations, MTM is not possible.
- For most fixed income instruments such as bonds, we can use models like the Yield Curve to reprice the asset more frequently.
  This would be a more accurate reflection of the current price of the bond compared with the current practise of using the book value, or of using the LTP.

Pressures to improve modelling volatility

From the ’80s, there has been increased focus on research in modelling financial market volatility. Reasons?

1. **Markets trading options (mid ’80s).** Volatility is a direct input to pricing options. A better forecast for volatility gets a “more accurate” options price. If the model is correct, you can make money in options markets.
2. **Risk management (late ’80s).** Better volatility models means more efficient deployment of risk capital to set aside to cover potential losses – Value at Risk models.
3. **Regulatory compulsions (Basle II).** In the second take on BIS norms, firm with “better, proven” risk models are allowed more flexibility with how they can deploy capital.
4. **Statistical coherence.** When financial market volatility is modelled as changing, it mostly (but not fully) explains non-normality of returns.

Risk measurement step 1: how to measure of risk?

Challenges in measuring risk

- **Biggest challenge**: how to measure risk?
- Assets are traded and prices are observed. But what can we use to *observe* risk?
- Commonly used statistical inference assume normality of returns and homoscedasticity.
- However stylised facts from all financial markets sugest:
  - Non-normality, and/or
  - Heteroskedasticity
- What measures are these facts based on?

Variance vs. Standard Deviation

- Volatility can be measured as σ^2 of returns.
- More often, volatility is denoted as the standard deviation of returns instead, σ.
- Given a time series of returns with T observations, we calculate it as:

$$\sigma = \sqrt{\sum_{t=1}^{t=T} (r_t - \bar{r})^2 / (T-1)}$$

  To create one value of σ, we need many observations of returns.
- Problem: How do we then create a time series of σ? Traditionally, create time series of r_t^2, where every day's returns r_t was squared as the estimate of that day's volatility.

Alternative measures for time series of volatility

- Traditional approaches:
  1. Moving window average σ,
  2. Range of prices as a measure of σ.
- Recent approaches using new markets and improved market information:
  1. Implied volatility from options markets,
  2. Realised volatility using high-frequency data.
- These become the inputs to creating models of volatility at the level of individual assets that get used in (a) forecasts of volatilty or (b) measures of portfolio volatility like VaR.

Value at Risk or VaR

Defining VaR

- The VaR is the loss that will be exceeded with a known probability. To calculate VaR, we need:
  1. **A well defined portfolio:** The VaR of a portfolio is not the same as the sum of the VaR of individual assets.
  2. **A well-defined period of time:** Most investors calculate a VaR over a day. Ie, what is the loss the portfolio could incur at the end of the next trading day?
     Firms with assets that do not have a daily marktomarket valuation technology could calculate their VaR over a longer horizon, say, a month.
  3. **A defined confidence level:** how frequently would the portfolio loss be greater than what you calculate?
     For example, 95% VaR is the minimum amount of loss you would expect to make on 5 days out of a 100day investment period.
     99% VaR is the miminimum amount of loss you would expect to make on one day out of a 100day investment period.
     It is not sensible to calculate 100% VaR.

Additional inputs to calculate VaR

- Once we have the portfolio, we also need:
  1. the MTM value of the portfolio.
  2. the risk of the portfolio, measured by the standard deviation or the variance of the portfolio returns.
- For traded securities, the MTM value is readily calculated. If there are non-traded or OTC assets in the portfolio, then we need to calculate the mark-to-model value of the assets.
- Similarly, for traded securities, the variance of the returns can be readily calculated.
  For non-traded and OTC assets, the variance of these assets will have to be modelled seperately.

Mathematical definition

- In mathematical parlance, if the oneday rupee profit on a portfolio is x and it has has a probability distribution function (pdf) $f(x)$, then the VaR v at a 95% level is:

$$\int_{-\infty}^{v} f(x)dx = 0.05$$

Calculating VaR, example 1

- We have an investment of Rs.1000 crore in the NSE-50 ("Nifty") stock index portfolio.
- Using the past levels of Nifty, we can calculate the variance of returns. This was about 1.26% daily upto the end of 2005.
- If we want to hold this investment for the next three years, at a confidence level of 99%, the VaR is calculated to be Rs.30 crore.
- We interpret this as follows:
  - This investment would see a single-day loss less than Rs.30 crore on 742 out of the next 750 trading days.
  - On 8 out of the next 750 trading days, this investment could see a single day loss of more than Rs.30 crore.
- This number becomes a decision tool for the management: are they willing to take this kind of risk in their investment portfolio or not.

Using VaR

- The above examples show how VaR can be used to make a direct comparison of the risk of various investments.
- VaR is a measure using which top management can take a decision on whether they are taking on too much risk or not.
  A firm with equity capital of Rs.1 crore would readily make a Rs.100 crore investment in Nifty because the 99% VaR is only Rs.3 crore
- VaR can also be used to set limits for traders and fund managers.
  A trader with limits of Rs.1 crore, and an investment variance of 3.5% cannot take a position greater than Rs.23 crore.
- VaR is used in clearing corporations to do risk-management.

Caveats and issues in VaR

Issues in implementing VaR

- VaR is a concept, a good theoretical construct.
  It needs to be well implemented to be useful.
- Ideal world:
  1. We will have the capacity to integrate price-risk, liquidity-risk and operational-risk into a single model for risk.
  2. The model will be correct.
  3. The model will be applicable to a portfolio of different assets.

  Then the VaR will be a “correct” estimate.

Issues in implementing VaR, contd.

- The real world calculation of VaR is very different. Two main problems are:
  1. There are imperfect information systems.
     For instance, in India, we do not have an integrated approach to price risk (of equity, commodities, interest rates, currency, credit). We do not even have systems with an integrated approach to give prices of all securities.
  2. There are imperfect models.
     - For example, there is evidence to indicate that returns are not normally distributed.
       However, the VaR is calculated using the normal distribution.
     - For example, the risk of the portfolio should be a forecast of the risk in the future, rather than the past.
       There are several models to forecast risk: how do we choose the "correct" one?

Issues in implementing VaR, contd.

- Therefore, in the real world, we have to constantly
  1. Test the VaR system
  2. Adjust for model errors

  We use engineering approximations of the theoretical concept.
- These can be good for many practical purposes. But we have to constantly adjust – either explicitly or implicitly – for the imperfections of VaR estimates.
- The Basle norms, while prescribing that banks develop their own VaR systems, also include incentives for constantly improving it.
  In Basle-II, the approach taken is that a bank obtains greater flexibility in doing their business if they can demonstrate to the regulator that their VaR system is "good".
- This is done on a regular basis, so that banks are constantly testing and improving their VaR systems.

What the VaR does not capture

Abusing the VaR

- In the deployment of capital at a bank or financial firm, it is not economical to have a VaR system which can never be exceeded.
- For example, setting the VaR of a GOI bond portfolio to be 40% of gross value of the investment.
- Given the low variance of interest rates, this is not likely to ever be exceeded.
- This is an extremely sound VaR system.
- However, it is an inefficient use of capital.

Misusing the VaR

- For example, if the board allocates USD 10 billion of equity capital for a bond trading business.
  1. A limit set by the VaR of 40% of the gross position means that the gross position of the trading desk is limited to USD 25 billion.
  2. The limit is not sensitive to whether it is a long or a short position, or a combination of the two.
  3. Given that the variance of interest rates are considerably lower, the limit should be set to 4%.
  4. This would imply that the trading desk ought to be able to take a gross position of USD 250 billion.
- In a good VaR system, there **has** to be some violations.

Extreme tail loss

- VaR captures the minimum loss that is exceeded at a given level of confidence.
- However, when the VaR **is** exceeded, what is the expected amount of loss?
- This is not answered by VaR.
- There is a new strand of literature called “Extreme Value Theory” (EVT) which does exactly this job.
- VaR assumed a normally distributed set of returns from which the expected minimum loss was chosen.
- EVT looks at the behaviour of just the left hand tail of the distribution of prices, and asks “if the VaR is exceeded, what is the average amount of loss for this security?”
- The behaviour of the EV can be different from that of a normal distribution because of:
  - High volatility and clustering of volatility
  - Lower liquidity
  - Higher correlation (lower diversification)
- Example, LTCM

Sources of financial market risk

Recent focus on modelling risk: risk factors

- Market value (price/market risk),
- Default (credit risk), and
- After the 2008 crisis, the markets ability to support transactions (liquidity risk).