FINANCIAL TIMES **Guides**

BANKING

GLEN ARNOLD

FT PUBLISHING
FINANCIAL TIMES

ALWAYS LEARNING
PEARSON

The Financial Times Guide to Banking

Income from selling loans

Many banks have become adept at originating loans but then shortly afterwards selling the right to receive interest and principal to another financial institution. These are often bundled up with other loans, e.g. the right to receive payments on 1,000 mortgages from householders for the next 20 years could be sold in the active mortgage-backed securities (MBS) market. The bank can make a profit by selling the 1,000 mortgages for more than was originally loaned. In other words, mortgagees are charged a higher interest rate than the MBS' holders receive. Thus the asset of 1,000 loans is removed from the balance sheet and cash is put there instead, at least temporarily until an alternative use is found, such as originating another 1,000 mortgages. (There is more on this repackaging of debt (securitisation) in Chapter 11.)

Trading financial instruments

Banks trade in a wide range of financial instruments on behalf of clients as agents, as market makers allowing buyers and sellers to trade and as proprietary traders trying to make profits from taking positions buying and selling on their own account in the markets (see Chapter 11). Also greater volatility in financial markets has led bankers to attempt to manage interest rate risk, foreign exchange movement risk and other risks through the use of futures, options, swaps and other derivative products. The availability of these liquid markets has also led to a greater propensity to speculate, to out-guess the markets. The huge bets made expose banks to enormous risks from time to time. There have been some spectacular failures as a result: (Barings Bank in 1995, Lehmans in 2008) and large losses e.g. UBS in 2011.

Value at risk used to estimate the overall risk of on- and off-balance-sheet exposures

Because banks hold a very wide range of assets from corporate loans to complicated derivatives, and they bear a number of obligations, senior managers can lose track as to the extent to which the firm as a whole is exposed to risk. One division might be building up large holdings of bonds while another is selling options and swaps, and yet another is packaging up mortgage bonds and selling MBS. Perhaps what one division is doing will offset the risk that another is taking on. On the other hand, it might be that risk is merely compounded by the combination of positions. Each day the mix of assets and liabilities changes and therefore the risk exposure changes.

Back in the 1990s some bankers[2] thought it would be a good idea to produce a single number that encapsulated the overall risk profile of the bank each evening. The senior managers could look at that and be reassured that they were not taking excessive risk. If the number started to look dangerously high then they could instruct a reweighting of assets and obligations until a safety margin was restored. The measure that they came up with is called **value at risk**, or **VaR**, which asks: 'If tomorrow is a bad day (e.g. different asset classes, such as shares and bonds, fall in market price significantly) what is the minimum that the bank will lose?' VaR is an estimate of the loss on a portfolio over a period of time (usually 24 hours is chosen) that will be *exceeded*[3] with a given frequency, e.g. a frequency of one day out of 100, or five days out of 100.

Another way of looking at the frequency element is called the **confidence level**. Thus with a 99% confidence level set the VaR might turn out to be $100m. Therefore for 'one-day VaR' there is a 1% chance that the portfolio could lose more than $100m in 24 hours. A 95% confidence level means that there is a 95% chance that the loss will be less than the derived figure of say $16m for a day, and a 5% chance that it will be greater than $16m.

Calculating VaR

So, how does a bank calculate VaR estimates? It needs some numbers and some assumptions. One assumption often made is that returns on a security (share, derivative, bond, etc.) follow a particular distribution. The usual assumption is the **normal distribution** where there is a large clustering of probabilities of returns around the average expected return and then very small probabilities of the extremes ('thin tails'). This is rather like the distribution in the heights of 14 year olds: most are clustered around 5' to 5'6'', with decreasing numbers of children at 4'9'', 4'8'', 4'7" and very few in the 'tail region' of 4' on the downside; and decreasing numbers at 5'7'', 5'8'' and 5'9'' with very few in the tail region of greater than 6'1''. And the pattern of decrease is symmetrical around the mean observation, say 5'3''.

For banks the distribution of possible outturns is symmetrical about the mean too – there is the same chance of being, say, £3m above the average expected return as of being £3m below it. See Figure 8.3 for a normal distribution of probabilities, a bell-shaped symmetrical curve. The usual source of data, whether combined with a normal distribution assumption or not, is a long time series of an historical data set of daily return data for the securities. Then the mathematicians assume that this represents the future distribution of returns. Another important source of information is the calculation of the extent to which asset returns move together – e.g. do corporate bond prices move in step with prices for MBS? Then it is assumed that these correlations remain true for future estimations.

Figure 8.3 shows a possible output from using VaR. On the right-hand side the return numbers increase but the probability of earning those high returns decrease significantly the further we move away from the average expected return (as would the probability of finding a 6'3'' 14 year old). The probabilities for returns below average are symmetrical with those above average in this case where we assume 'normality'. (If we have used real past return data the distribution may not be quite normal and the maths for calculating the confidence level becomes more complicated, but it can still be handled. A skewed distribution just creates more fun for the mathematicians.) You can read off the 99% chance of not losing more than the amount marked by the dashed line. If the dashed line is at –$200m then one day out of 100 we would lose more than –$200m, in theory.

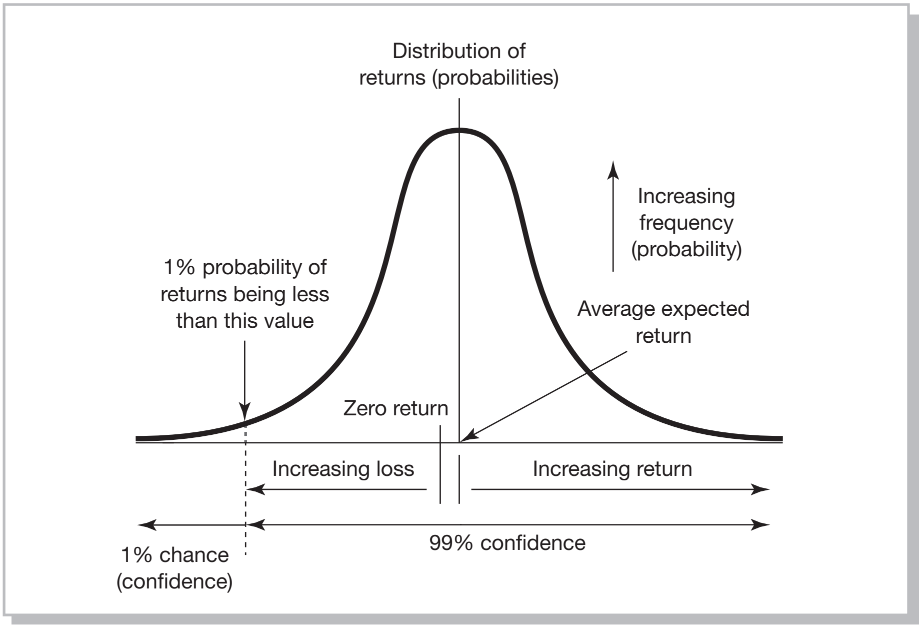

Figure 8.3 **A VaR analysis assuming a normal distribution of the probabilities of return on an asset or collection of assets**

Is past data enough?

A key assumption is that the past data has a very close bearing on the future probabilities. Even using the normality method you need past data to estimate the size of the various probabilities. So, you might gather data from the previous three years. If unusual/infrequent events are not present in that data set you might be missing some extreme positive or negative possibilities. This event might have a massive impact on risk, but might be missed if the data set is limited to a period of stability. Indeed it might be missed even if the data set is extensive if the event is very rare. An influential writer on derivatives and market players, Nicolas Taleb[4] was warning us long before the crisis that the mathematicians were not allowing for the possibility of extreme events (dubbed **black swans**) – that the tails of the distribution are in fact much bigger, 'fatter', than generally supposed, because extreme things hit us more often then we anticipate. There are 'high-impact rare events', often caused by human reactions to apparently insignificant triggers; reactions such as exuberance, fear and panic. Asset returns and liabilities which appeared to be uncorrelated suddenly all move together to a much greater extent than short-memory players in financial markets expect. There are risks out there that we did not know existed, until it is too late. The more experienced old-hands know that unexpected and unimaginable events shape the markets, leaving them with question marks concerning the extent to which they can trust historical data to be the only guide.

Led astray

VaR led many senior bankers into a false sense of security prior to the crisis. They were thus emboldened to double and triple their bets on securities such as mortgage-backed bonds and the related derivatives because they appeared to produce high profits without raising VaR much. Many regulators insisted on the disclosure of VaR. Indeed, under the supervisory capital rules for banks VaR could be used as an argument for lowering capital requirements. Banks jumped at the chance of lowering capital buffers so as to increase return on equity (see Chapter 7) – they leveraged themselves up. They particularly liked to stock up on 'high-quality' mortgage-based derivatives because the model told them that they had trivial VaR and so they were required to hold only trivial capital reserves. Yes, the regulators have much to answer for, too. They placed far too much faith in these mathematical models. They even allowed banks to do their own calculations of risk exposure, and more or less accepted these for setting capital limits.

The crisis

Although we may have moved on from the global financial crisis which began in 2007 and the economic climate is very different, it is still an important event in any discussion on banking and finance. Prior to the summer of 2007 banks tended to use historical evidence going back between one and five years to estimate VaR. By then a large part of their exposure was to the mortgage market. In the mid-2000s this had been as placid as the sea was when the Titanic was crossing the Atlantic – no trouble encountered for day after day. It would seem that VaR gave little indication of real likely losses. For example, in the autumn of 2007 Bear Stearns reported an average VaR of $30m[5] which is tiny for a bank with so many assets. It could withstand days and days of such losses and barely notice. So according to VaR it was hardly at risk at all. And yet within weeks it was bust, losing $8bn of value. As the complete failure of the bank approached, the VaR number did rise slightly to $60m because it started to incorporate data for the most recent days when securities became more volatile as the market mayhem started – but it still had lots of older placid data pushing the number down.

It turned out that 'highly unlikely events' such as a correlated fall in house prices all over the USA, and the fall in market prices of derivatives and bonds, can all happen at the same time. The average person on the street could have told them that, but these mathematicians and financial economists could not see past their complex algebra. The guys over at Merrill Lynch were perplexed: 'In the past these AAA CDO securities [mortgage based derivatives – see Chapter 11] had never experienced a significant loss in value.'[6] But how long was that 'past' – if they could find a decade of data they would have been lucky, because these instruments were so young.

As the crisis got under way most of the large banks experienced losses that should only happen once in 1,000 years (according to their models) day after day. In their language these were six-sigma (standard deviation) events, which were virtually impossible from their faithful-to-the-algebra perspective. Indeed, in August 2007 the market experienced several 25-sigma events – these should happen only once every 14 billion years. How confusing for them – the model let them down when it really mattered. Triana[7] likens VaR to buying a car with an air bag which protects you 99% of the time, that is when conditions are moderate, but if you have a serious crash it fails.

Bear Stearns understood the problems with VaR. Take this statement from their filing with the Securities and Exchange Commission on February 29 2008:

> VaR has inherent limitations, including reliance on historical data, which may not accurately predict future market risk, and the quantitative risk information generated is limited by the parameters established in creating the models. There can be no assurance that actual losses occurring on any one day arising from changes in market conditions will not exceed the VaR amounts shown below or that such losses will not occur more than once in 20 trading days. VaR is not likely to accurately predict exposures in markets that exhibit sudden fundamental changes or shifts in market conditions or established trading

> relationships. Many of the Company's hedging strategies are structured around likely established trading relationships and, consequently, those hedges may not be effective and VaR models may not accurately predict actual results. Furthermore, VaR calculated for a one-day horizon does not fully capture the market risk of positions that cannot be liquidated in a one-day period.[8]

Despite this list of doubts they had used it because everyone else did (and the regulators allowed them to get away with it).

Once faith in VaR had evaporated those lending to banks became afraid that the risk metrics they publicly announced under-reported their real exposure. They took the action that you or I would take when told that borrowers might default on what they owe: stop lending any more money and try to call in old loans. The problem is the entire banking system was a complex web of loans to each other and once confidence had gone the whole system collapsed. Notice the key words here are not quantifiable – 'afraid', 'confidence' – these fuzzy things are just as important to understand about banking as the maths, if not more so.

There is more on the causes of the financial crisis in *The Financial Times Guide to Financial Markets* by Glen Arnold (Pearson: 2012).