Measures of Financial Risk

Peter Cincinelli

Financial Instruments & Markets
[Financial Investment and Corporate Finance]
Academic year 2022-2023

December 7, 2022

Outline of talk

Table of Contents

The Revolution of Risk Measures in Finance

- Financial risk is the prospect of financial loss -or gain- due to unforeseen changes in underlying risk factors;
- In this lesson we will investigate one particular form of financial risk, namely **market risk**, or the risk of loss (or gain) arising from unexpected changes in market prices (e.g., such as security prices) or market rates (e.g., such as interest or exchange rates);
- Market risks, in turn, can be classified into interest-rate risks, equity risks, exchange rate risks, commodity price risks, and so on, depending on whether the risk factor is an interest rate, a stock price, or whatever;
- Market risks can also be distinguished from other forms of financial risk, most especially credit risk (or the risk of loss arising from the failure of a counterparty to make a promised payment) and operational risk (or the risk of loss arising from the failures of internal systems or the people who operate in them);

- The theory and the practice of risk management have developed enormously since the pioneering work of Harry Markowitz in the 1950s;
- The theory has developed to the point where risk management (or risk measurement) is now regarded as a distinct sub-field of the theory of finance;
- The subject has attracted a huge amount of intellectual energy, not just from finance specialists but also from specialists in other disciplines who are attracted to it (e.g., physics), attracted not just by high salaries but also by the challenging intellectual problems it poses.

Contributory Factors - A Volatile Environment

- One factor behind the rapid development of risk management was the high level of instability in the economic environment within which firms operated;
- A volatile environment exposes firms to greater financial risk, and therefore provides an incentive for firms to find new and better ways of managing this risk. The volatility of the economic environment is reflected in various factors:
  1. Stock market volatility;
  2. Exchange rate volatility;
  3. Interest rate volatility;
  4. Commodity market volatility.

Contributory Factors - Growth in Trading Activity

- Another factor contributing to the transformation of risk management is the huge increase in trading activity since the late 1960s;
- The average number of shares traded per day in the New York Stock Exchange has grown from about 3.5 mln in 1970 to around 100 mln in 2000; and turnover in foreign exchange markets has grown from about a billion dollars a day in 1965 to $1,210 billion in April 2001;
- There have been massive increases in the range of instruments traded over the past two or three decades, and trading volumes in these new instruments have also grown very rapidly;
- New instruments have been developed in offshore markets and, more recently, in the newly emerging financial markets of Eastern Europe, China, Latin America, Russia, and elsewhere;

Contributory Factors - Growth in Trading Activity

- New instruments have also arisen for assets that were previously illiquid, such as consumer loans, commercial and industrial bank loans, mortgages, mortgage-based securities, and similar assets, and these markets have grown very considerably since the early 1980s;
- There has also been a phenomenal growth of derivatives activity. Until 1972 the only derivatives traded were certain commodity futures and various forwards and over-the-counter (OTC) options. The Chicago Mercantile Exchange then started trading foreign currency futures contracts in 1972, and in 1973 the Chicago Board Options Exchange started trading equity call options;
- Interest-rate futures were introduced in 1975, and a large number of other financial derivatives contracts were introduced in the following years: swaps and exotics (e.g., swaptions, futures on interest rate swaps, etc.);

Contributory Factors - Advances in Information Technology

- A third contributing factor to the development of risk management was the rapid advance in the state of information technology;
- Improvements in IT have made possible huge increases in both computational power and the speed with which calculations can be carried out. Improvements in computing power mean that new techniques can be used (e.g., such as computer-intensive simulation techniques) to enable us to tackle more difficult calculation problems;
- Improvements in calculation speed then help make these techniques useful in real time, where it is often essential to get answers quickly.

Table of Contents

Risk Measurement

- To understand recent developments in risk measurement, we need first to appreciate the more traditional risk measurement tools:
  1. Gap Analysis;
  2. Duration Analysis;
  3. Scenario Analysis;
  4. Portfolio Theory;
  5. Derivatives Risk Measures.

Risk Measurement - Gap Analysis

- One common approach was (and, in fact, still is) gap analysis, which was initially developed by financial institutions to give a simple, albeit crude, idea of interest-rate risk exposure;
- Gap analysis starts with the choice of an appropriate horizon period (1 year), or whatever;
- We then determine how much of our asset or liability portfolio will re-price within this period, and the amounts involved give us our rate-sensitive assets and rate-sensitive liabilities;
- The gap is the difference between these, and our interest-rate exposure is taken to be the change in net interest income that occurs in response to a change in interest rates. This in turn is assumed to be equal to the gap times the interest-rate change:

- This in turn is assumed to be equal to the gap times the interest-rate change:

$$\Delta NII = (GAP)\Delta r \tag{1}$$

- where ΔNII is the change in net interest income and Δr is the change in interest rates;
- Gap analysis is fairly simple to carry out, but has its limitations:
  1. it only applies to on-balance sheet interest-rate risk, and even then only crudely;
  2. it looks at the impact of interest rates on income, rather than on asset or liability values;
  3. results can be sensitive to the choice of horizon period.

- Another method traditionally used by financial institutions for measuring interest-rate risk is duration analysis. The Duration D of a bond (or any other fixed-income security) can be defined as the weighted average term to maturity of the bond' cash flows, where the weights are the present values of each cash flow relative to the present value of all cash flows:

$$D = \sum_{i=1}^{n} (i * PVCF_i) / \sum_{i=1}^{n} PVCF_i \tag{2}$$

- where $PVCF_i$ is the present value of the period i cash flow, discounted at the appropriate spot period yield. The duration measure is useful because it gives an approximate indication of the sensitivity of a bond price to a change in yield:

$$\Delta Price \approx -D * \Delta y / (1 + y) \tag{3}$$

- where y is the yield and Δy the change in yield. The bigger the duration, the more the bond price changes in response to a change in yield;
- The duration approach is very convenient because duration measures are easy to calculate and the duration of a bond portfolio is a simple weighted average of the durations of the individual bonds in that portfolio;
- It is also better than gap analysis because it looks at changes in asset (or liability) values, rather than just changes in net income;

Risk Measurement - Duration Analysis

- However, duration approaches have similar limitations to gap analysis:

  1. they ignore risks other than interest-rate risk;
  2. they are crude, and even with various refinements to improve accuracy, duration based approaches are still inaccurate relative to more recent approaches to fixed-income analysis;

- Moreover, the main reason for using duration approaches in the past is no longer of much significance, since more sophisticated models can now be programmed into micro-computers to give their users more accurate answers rapidly.

Risk Measurement - Scenario Analysis

- A third approach is scenario analysis (or *"what if"* analysis), in which we set out different scenarios and investigate what we stand to gain or lose under them;
- To carry out scenario analysis, we select a set of scenarios, or paths describing how relevant variables (e.g., stock prices, interest rates, exchange rates, etc.) might evolve over a horizon period;
- We then postulate the cash flows and/or accounting values of assets and liabilities as they would develop under each scenario, and use the results to come to a view about our exposure;
- Scenario analysis is not easy to carry out. A lot hinges on our ability to identify the 'right' scenarios, and there are relatively few rules to guide us when selecting them;

- We need to ensure that the scenarios we examine are reasonable and do not involve contradictory or excessively implausible assumptions, and we need to think through the interrelationships between the variables involved;
- We also want to make sure, as best we can, that we have all the main scenarios covered;
- Scenario analysis also tells us nothing about the likelihood of different scenarios, so we need to use our judgement when assessing the practical significance of different scenarios;
- In the final analysis, the results of scenario analyses are highly subjective and depend to a very large extent on the skill or otherwise of the analyst.

Risk Measurement - Portfolio Theory

- A somewhat different approach to risk measurement is provided by portfolio theory;
- Portfolio theory starts from the premise that investors choose between portfolios on the basis of their expected return, on the one hand, and the standard deviation (or variance) of their return, on the other;
- The standard deviation of the portfolio return can be regarded as a measure of the portfolio' risk. Other things being equal, an investor wants a portfolio whose return has a high expected value and a low standard deviation;
- These objectives imply that the investor should choose a portfolio that maximises expected return for any given portfolio standard deviation or, alternatively, minimises standard deviation for any given expected return;

Risk Measurement - Portfolio Theory

- A portfolio that meets these conditions is efficient, and a rational investor will always choose an efficient portfolio. When faced with an investment decision, the investor must therefore determine the set of efficient portfolios and rule out the rest;
- Some efficient portfolios will have more risk than others, but the more risky ones will also have higher expected returns. Faced with the set of efficient portfolios, the investor then chooses one particular portfolio on the basis of his or her own preferred trade-off between risk and expected return;
- An investor who is very averse to risk will choose a safe portfolio with a low standard deviation and a low expected return, and an investor who is less risk averse will choose a more risky portfolio with a higher expected return.

- One of the key insights of portfolio theory is that the risk of any individual asset is not the standard deviation of the return to that asset, but rather the extent to which that asset contributes to overall portfolio risk;
- An asset might be very risky (i.e., have a high standard deviation) when considered on its own, and yet have a return that correlates with the returns to other assets in our portfolio in such a way that acquiring the new asset adds nothing to the overall portfolio standard deviation;
- Acquiring the new asset would then be riskless, even though the asset held on its own would still be risky. The moral of the story is that the extent to which a new asset contributes to portfolio risk depends on the correlation or covariance of its return with the returns to the other assets in our portfolio;

- Portfolio theory provides a useful framework for handling multiple risks and taking account of how those risks interact with each other;
- It is therefore of obvious use to portfolio managers, mutual fund managers and other investors. However, it tends to run into problems over data;
- The risk-free return and the expected market returns are not too difficult to estimate, but estimating the betas is often more problematic. Each beta is specific not only to the individual asset to which it belongs, but also to our current portfolio. To estimate a beta coefficient properly, we need data on the returns to the new asset and the returns to all our existing assets, and we need a sufficiently long data set to make our statistical estimation techniques reliable;
- The beta also depends on our existing portfolio and we should, in theory, re-estimate all our betas every time our portfolio changes;
- Using the portfolio approach can require a considerable amount of data and a substantial amount of ongoing work.

- When dealing with derivatives positions, we can also estimate their risks by their Greek parameters:
- The **delta** Δ, which gives us the change in the derivatives price in response to a small change in the underlying price;
- The **gamma** Γ, which gives us the change in the delta in response to a small change in the underlying price (or, if we prefer, the second derivative of the derivatives price with respect to a change in the underlying price);
- The **rho** ρ, which gives us the change in derivatives price for a small change in the interest rate;
- The **vega** ν, which gives us the change in derivatives price with respect to a change in volatility;
- The **theta** Θ, which gives us the change in derivatives price with respect to time;

- In using these measures, we should also keep in mind that they make sense only within the confines of a dynamic hedging strategy: the measures, and resulting hedge positions, only work against small changes in risk factors, and only then if they are revised sufficiently frequently;
- There is always a worry that these measures and their associated hedging strategies might fail to cover us against major market moves such as stock market or bond market crashes, or a major devaluation;
- We may have hedged against a small price change, but a large adverse price move in the wrong direction could still be very damaging: our underlying position might take a large loss that is not adequately compensated for by the gain on our hedge instrument.

Table of Contents

Value at Risk - VaR - Origin and Development

- In the late 1970s and 1980s, a number of major financial institutions started work on internal models to measure and aggregate risks across the institution as a whole;
- They started work on these models in the first instance for their own internal risk management purposes, as firms became more complex it was becoming increasingly difficult, but also increasingly important, to be able to aggregate their risks, taking account of how they interact with each other, and firms lacked the methodology to do so;
- According to industry legend, the chairman of JP Morgan, Dennis Weatherstone, asked his staff to give him a daily one-page report indicating risk and potential losses over the next 24 hours, across the bank's entire trading portfolio;
- This report, the famous *"4:15 report"*,was to be given to him at 4:15 each day, after the close of trading.

- In order to meet this demand, the Morgan staff had to develop a system to measure risks across different trading positions, across the whole institution, and also aggregate these risks into a single risk measure;
- The measure used was Value at Risk (or VaR), or **the maximum likely loss over the next trading day**, and the VaR was estimated from a system based on standard portfolio theory, using estimates of the standard deviations and correlations between the returns to different traded instruments;
- While the theory was straightforward, making this system operational involved a huge amount of work: measurement conventions had to be chosen, data sets constructed, statistical assumptions agreed, procedures determined to estimate volatilities and correlations, computing systems established to carry out estimations, and many other practical problems resolved.

- The subsequent adoption of VaR systems was very rapid. First among securities houses and investment banks, and then among commercial banks, pension funds and other financial institutions, and non-financial corporates;
- The state of the art also improved rapidly;
- Developers and users became more experienced. VaR systems were extended to cover more types of instruments; and the VaR methodology itself was extended to deal with other types of risk besides the market risks for which VaR systems were first developed, including credit risks, liquidity risks and cash-flow risks.

Table of Contents

Measuring VaR

- There are three basic approaches (such as: **Variance-Covariance Method**, **Historical Simulation**, **Monte Carlo Simulation**) that are used to compute Value at Risk, though there are numerous variations within each approach;
- The measure can be computed analytically by making assumptions about return distributions for market risks, and by using the variances in and covariances across these risks;
- It can also be estimated by running hypothetical portfolios through historical data or from Monte Carlo simulations.

Measuring VaR

- Suppose we are working to a daily holding or horizon period;
- At the end of day $t - 1$, we observe that the value of our portfolio is P_{t-1}. However, looking forward, the value of our portfolio at the end of tomorrow, P_t, is uncertain;
- Ignoring any intra-day returns or intra-day interest, if P_t turns out to exceed P_{t-1}, we will make a profit equal to the difference, $P_t - P_{t-1}$;
- If P_t turns out to be less than P_{t-1}, we will make a loss equal to $P_{t-1} - P_t$. Since P_t is uncertain, as viewed from the end of $t - 1$, then so too is the profit or loss (P/L);
- Our next-period P/L is risky, and we want a framework to measure this risk.

Measuring VaR - The Normality Assumption

- The traditional solution to this problem is to assume a mean–variance framework;
- We model financial risk in terms of the mean and variance (or standard deviation, the square root of the variance) of P/L (or returns);
- As a convenient starting point, we can regard this framework as underpinned by the assumption that daily P/L (or returns) obeys a normal distribution;
- Strictly speaking, the mean–variance framework does not require normality, and many accounts of it make little or no mention of normality. Nonetheless, the statistics of the mean–variance framework are easiest understood in terms of an underlying normality assumption, and viable alternatives (e.g., such as assumptions of elliptical distributions) are usually harder to understand and less tractable to use.

Measuring VaR - The Normality Assumption

- A random variable X is normally distributed with mean μ and variance σ^2 (or standard deviation σ) if the probability that X takes the value x, $f(x)$, obeys the following probability density function (pdf):

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} exp\left[-\frac{1}{2}((x-\mu)/\sigma)^2\right] \tag{4}$$

- where X is defined over $-\infty < x < \infty$. A normal pdf with μ=0 and σ=1, known as a standard normal, is illustrated in Figure 1.

Measuring VaR - The Normality Assumption

Figure 1: The normal probability density function.

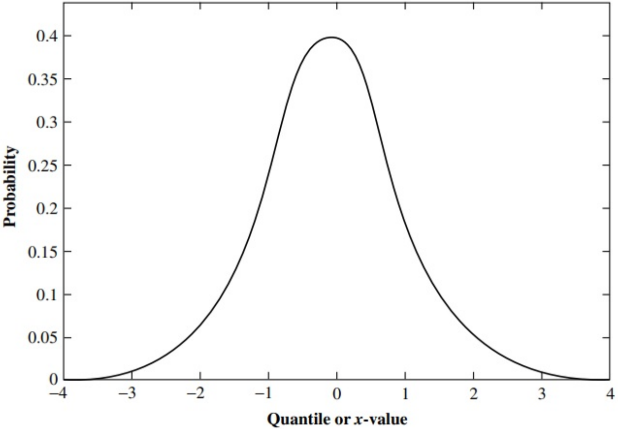

- This pdf tells us that outcomes are more likely to occur close to the mean μ;
- The spread of the probability mass around the mean depends on the standard deviation σ: the greater the standard deviation, the more dispersed the probability mass;
- The pdf is also symmetric around the mean: X is as likely to take a particular value $x - \mu$ as to take the corresponding negative value $-(x - \mu)$;
- Outcomes well away from the mean are very unlikely, and the pdf tails away on both sides: the left-hand tail corresponds to extremely low realisations of the random variable, and the righthand tail to extremely high realisations of it;
- In risk management, we are particularly concerned about the left-hand tail, which corresponds to high negative values of P/L, or big losses.

Measuring VaR - The Normality Assumption

- A pdf gives a complete representation of possible random outcomes: it tells us what outcomes are possible, and how likely these outcomes are. Such a representation enables us to answer questions about possible outcomes and, hence, about the risks we face. These questions come in two basic forms:

    1. The first are questions about likelihood or probability. We specify the quantity (or quantile), and then ask about the associated probability. For example, how likely is it that profit (or loss) will be greater than, or less than, a certain amount?
    2. The others are questions about quantiles. We specify the probability, and then ask about the associated amount. For example, what is the maximum likely profit (or loss) at a particular level of probability?

- These questions and their answers are illustrated in Figure 2.

Measuring VaR - The Normality Assumption

Figure 2: Normal quantiles and probabilities.

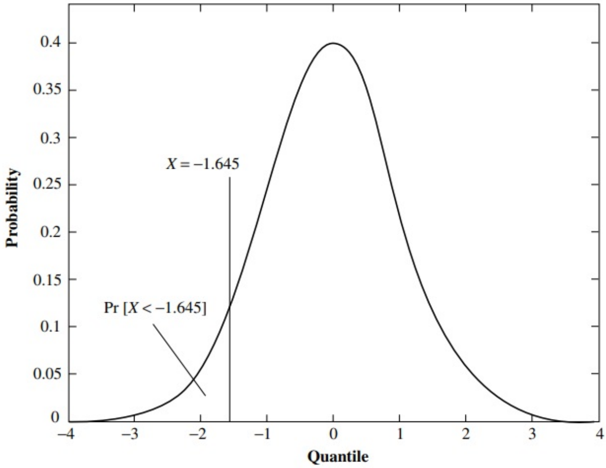

Measuring VaR - The Normality Assumption

- Figure 2 shows the same normal pdf, but with a particular $X - value$, equal to -1.645;
- We can regard this value as a profit of -1.645 or a loss of 1.645;
- The probability of a P/L value less than -1.645 is given by the left-hand tail, the area under the curve to the left of the vertical line marking off X=-1.645;
- This area turns out to be 0.05, or 5%, so there is a 5% probability that we will get a P/L value less than -1.645, or a loss greater than 1.645;
- Conversely, we can say that the maximum likely loss at a 95% probability level is 1.645. This is often put another way: we can be 95% confident of making a profit or making a loss no greater than 1.645;
- This value of 1.645 can then be described as the value at risk (VaR) of our portfolio at the 95% of confidence.

Measuring VaR - The Normality Assumption

- The assumption that P/L is normally distributed is attractive for three reasons;
- The first is that it often has some, albeit limited, plausibility in circumstances where we can appeal to the central limit theorem;
- The second is that it provides us with straightforward formulas for both cumulative probabilities and quantiles, namely:

$$Pr[x \leq X] = \int_{-\infty}^{X} \frac{1}{\sigma\sqrt{2\pi}} exp\left[-\frac{1}{2}((x-\mu)/\sigma)^2\right] dx \tag{5}$$

$$X_{cl} = \mu + \alpha_{cl}\sigma \tag{6}$$

Measuring VaR - The Normality Assumption

- Where *cl* is the chosen confidence level (e.g., 95%), and α_{cl} is the standard normal variate for that confidence level (e.g., α=0.95=-1.645). α_{cl} can be obtained from standard statistical tables or from spreadsheet functions;
- Equation 5 is the normal distribution (or cumulative density) function, which gives the normal probability of x being less than or equal to X, and enables us to answer probability questions;
- Equation 6 is the normal quantile corresponding to the confidence level *cl* (i.e., the lowest value we can expect at the stated confidence level) and enables us to answer quantity questions;
- The third advantage of the normal distribution is that it only requires estimates of two parameters, the mean and the standard deviation (or variance), because it is completely described by these two parameters alone.
- Nonetheless, the assumption of normality also has its limitations. You will happy to see them in other course.

Table of Contents

References

- Dowd, K. (2003). *An Introduction to Market Risk Measurement*. John Wiley & Sons.