Measuring credit portfolio risk

Research paper

February 2010 **Written by Leah Kelly and Tony Adams**

When investing in credit, the question "How do I measure how much risk I am taking?" is not easily answered. The answer might include duration, average credit rating, tracking error number or Value at Risk (VaR) but, unfortunately, all these measures can be significantly deficient by failing to capture the true risk profile of credit investments.

This paper does not intend to diminish the value of high quality credit analysis and stock selection in order to reduce credit losses; rather it addresses the portfolio construction elements of credit portfolio management. The key risk in credit portfolio management is the risk of extreme losses (known as tail risk) and many measures fail to adequately capture this risk. This paper seeks to highlight the benefits, flaws and assumptions of each of the current approaches to measuring credit risk and promotes a push for a more standardised reporting model, which would enable a true relative comparison of credit portfolio risk. The theories and approaches discussed in this paper are well documented and firmly established.[1]

Credit risk

Risk is the chance that an investment will not generate its expected return. In the case of bonds, this may incorporate market risk and credit risk. It is important to differentiate between the two. Market risk arises when the market price of assets changes due to changing market factors, which may be interest rates or credit spreads. This paper does not focus on the market risk of bonds – for details, see Dowd (2005) and references therein. Credit (or default) risk is the risk that the specific investment or portfolio will not generate the expected returns due to the failure of the borrower to repay the principal and interest when due.

What is portfolio credit risk?

When making a credit investment there is always the risk that the borrower might default on its debt, measured by a credit rating[2] or default probability[3]. For strong investment grade borrowers this risk is very low but still exists. Therefore, in any credit portfolio, there is an expected rate of default. Portfolio credit risk is the risk that defaults, or losses due to defaults, run at a different level than is expected. Credit portfolio management is primarily about ensuring that losses due to defaults do not run at a higher rate than expected.

The key parameters impacting the credit risk of a portfolio are:

- Underlying credit risks;
- Default correlations; and
- Exposure positions (concentrations).

Correlation[4] is a term used to identify the degree to which things move together, that is the strength of the relationship between variables. Default correlation tells us, given a certain borrower has defaulted, the likelihood that another borrower will default. Often, due to industry factors, defaults can occur in clusters, such as the telecom equipment sector[5] in Europe or the oil and gas industry[6].

Portfolios that hold assets with a high default correlation are more likely to experience default clustering than those with low default correlations. Concentration refers to holding large positions in specific names. Since defaults can occur randomly, the size of an investor's holding when an asset defaults will directly impact the loss effect on the portfolio.

The following chart shows the actual historical default rates for 'BBB' rated credits since 1970.

Default frequency for 'Baa' rated credits (equivalent 'BBB' using S&P rating terminology)

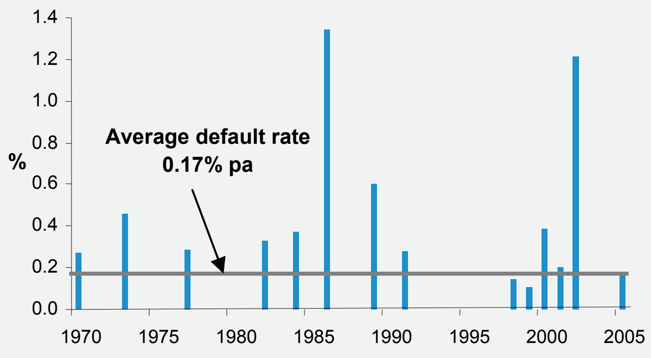

Source: Moody's

The chart shows that 'BBB' rated credits on average default at a rate of 0.17% pa, or 17 basis points per year. However, sometimes no issues default and in some years quite a few do. It is also important to note that these are averages across the entire rated universe of bonds, whereas portfolios will always own a much smaller number of bonds.

The impact of this concentration can be significant. While the average 'BBB' default rate in 2002 was 1.21%, for example, a portfolio of 10 equally weighted 'BBB' rated bonds that owned WorldCom and Enron (defaulted Dec 2001) would have had a default rate of 20%!

Credit migration refers to the changing of a credit's rating through time. A company rated 'A' might experience some difficulty, or additional leverage, and be downgraded to a 'BBB' or 'BB' rating. Migration might also be an improvement in rating. Migration does impact the market value of assets and a less concentrated portfolio (i.e. one with a greater number of holdings) will be less susceptible to the impact of price changes due to specific names migrating. However, it must be remembered that migration is related to market risk and not, except in the extreme case where an issue migrates to a defaulted state, credit risk. This is because credit risk is the risk of default. A migration is a change in rating that may result in a change in credit spreads (which is a market risk). If an investor's investment maturity matches their investment horizon the only risk they face is credit or default risk. Since migration is a market risk, not a credit risk, it is not explicitly covered in this paper.

Traditional measures of credit risk

Average credit rating

The average credit rating is the simplest, and commonly used, tool used to measure the credit risk within a portfolio. It is also arguably the most misleading.

The average credit rating is simply a measure of the weighted average credit rating of all the individual assets in the portfolio. The first problem with this approach is that many participants use a scale that is too simplistic when calculating the average and the second is that it ignores the universally accepted benefit of portfolio diversification[7]; see, for example, Ranson (2003). Ranson describes diversification as "the only free lunch in finance". As opposed to an equity portfolio where the returns are more symmetric, in a debt portfolio receiving the principal at maturity plus interest through the life of the bond is "the upside". Ranson states that debt portfolio management "is not a search to find winners, it is a process to minimise the effects of losers". The third is that average credit rating is not actually a measure of risk at all. Risk is a measure of uncertainty, whereas an average is a measure of central tendency, a measure of what you do expect, and not a measure of how wrong you might be.

A linear scale is often used when calculating average portfolio credit ratings. Thus a 'AAA' rated bond is given a score of 1, 'AA+' a score of 2, 'AA' 3, 'AA-' 4, 'A+' 5, and so on. The portfolio credit rating is then determined by weighting the scores by the weights for each bucket.

Credit ratings indicate a certain probability of default. The following table shows the cumulative probability of default for various rating bands based on Moody's historical data[8].

Credit Rating*	10 year cumulative historical default probability
AAA	0.60%
AA	0.78%
A	1.24%
BBB	3.63%
BB	27.39%
B	50.42%

Source: Moody's *We have used S&P rating terminology for consistency

Thus 3.63% of issues rated 'BBB' have defaulted within 10 years of them being rated 'BBB'. It can easily be seen that stepping down in credit quality does not result in a linear increase in default probability. The default risk increase is significant and non-linear as the rating falls. As a consequence some industry participants will use a non-linear scale.

A portfolio of 10 'BBB' rated Australian utility securities will have an average rating of 'BBB'. A portfolio of 1,000 'BBB' globally diversified securities will also have an average rating of 'BBB'. However, the portfolio risk of each is significantly different. It will only take one default to generate a significant loss in the first portfolio (and indeed if the default is due to industry or regional factors and correlation is high, more than one default is possible) whereas it would take 100 defaults in the second portfolio to generate the equivalent negative impact. This shows that while the expected outcome of both is the same, the potential to deliver a substantially different actual outcome between the portfolios is significant.

Credit duration

Another component of credit risk is credit duration[9] (for further details, see Fabozzi (2005) and references therein). This gives a measure of how much a portfolio might lose if credit spreads widened. It is mathematically related, among other things, to the length of the exposure. Thus a portfolio with credit duration of 4 would expect to lose 4 basis points if credit spreads were to widen by 1 basis point. It is a measure of market risk rather than default or migration risk.

Credit duration, as it is typically applied, assumes that all securities have credit spreads that widen by the same amount. To illustrate this, the use of credit duration to measure credit risk assumes that if the spread on a 10 year US 'B' rated subordinated airline bond widened by 10 basis points then the spread on an AUD 'AAA' rated, 1 year Super Senior Mortgage Backed Security would also widen by exactly 10 basis points.

While credit duration gives a good 'back of the envelope' measure of what might happen to a portfolio if credit spreads were to uniformly change, it makes no assessment on the likelihood of that happening and does not, in any way, measure credit migration or default risk.

Tracking error

Tracking error[10] is the approach adopted by many participants when reviewing the performance of a portfolio against its benchmark, see Fabozzi (2005). Tracking error looks at excess return volatility as a risk measure.

Tracking error may be measured ex-post (after the event) or ex-ante (before the event). Ex-post tracking error uses the actual volatility of a portfolio's active returns (excess returns above the benchmark). Therefore this measure calculates what the excess return volatility was, rather than what the current risk is, or will be. Ex-ante tracking error requires portfolio modelling to estimate the future volatility of a portfolio's excess returns based on historical trends.

The basic assumption here is that the risk of a portfolio is defined as the volatility of its returns relative to the benchmark.

This approach has two major flaws:

- The first is that it is indifferent between strategies that lower total risk and those that increase total risk if their variance (tracking error) to benchmark is the same.

– The second is that it assumes that volatility is the complete and appropriate measure of risk – thereby assuming normally distributed portfolios. The fact that returns from bonds are not normally distributed means that we need to look more closely at the distribution. Recall that a normal distribution can be completely described by its mean and standard deviation. It is symmetric and exhibits no skew. Since a bond portfolio has limited upside it exhibits a skewed return distribution, see, for example, Dynkin et al (2007), A portfolio could be constructed that has low volatility but a large tail risk and tracking error would not capture this risk.

However, the major deficiency of tracking error is that it is a market risk measure and not a credit risk measure. Even ex-ante, or forward looking, tracking error models base their calculations on historical volatilities and do not incorporate migration tables or default estimates. As a result they do not capture the migration and default risk of the portfolio they are modelling and are not, therefore, measuring credit risk. More sophisticated credit modelling tools are required to measure the credit risk within portfolios.

Best practice measures of credit risk

Measuring tail risk

There are a number of techniques available for the estimation of credit tail risk[11] including Monte Carlo simulation, see, for example, Glasserman (2003), and extreme value theory, see Embrechts, Kluppelberg & Mikosch (1997). Further, there are a number of different tools available to model portfolio credit risk including Moody's KMV's Portfolio Manager, CreditMetrics, Credit Suisse's CreditRisk+ and Algorithmics' Portfolio Credit Risk Management tool. All these tools allow the construction of portfolio loss distributions and the calculation of a variety of tail risk measures. They do build their models on varying assumptions, however, and these assumptions and their implications must be understood.

Many tail measures calculate a VaR[12] (Value at Risk) see, for example, Dowd (2005). This tells us, with a given probability, over a given time frame, the most we would expect to lose due to credit losses. Thus a 1 year, 99.5 (3 standard deviations) loss number of $x tells us that in 995 out of every 1,000 years we would expect to lose less than $x. However, it tells us nothing about the shape of the distribution around this point. For example, in those five years where we might expect to lose more than $x, how much more might we lose? And in how many of the 995 years would we lose close to, but less than, $x?

A common criticism of VaR is that it does not, in all cases, reduce in the case of diversification.[13] Take the following simple example of two bonds A and B, that default independently. Both bonds have two default states with probabilities 3% and 2%, recovering $70 and $90, respectively. The bond redeems at 100 in all other scenarios. We assume no market moves and the initial value of the bonds is equivalent to the expected value. The table below illustrates the scenarios;

Scenario	Probability	Bond A	Bond B	Portfolio (A+B)
1	0.03	70	100	170
2	0.02	90	100	190
3	0.03	100	70	170
4	0.02	100	90	190
5	0.9	100	100	200

The expected value of bond A and B is 98.9, respectively and the expected value of the portfolio is 197.8.

The table below shows the profit and loss (expected value less value under each scenario) under each scenario;

Scenario	Probability	Bond A	Bond B	Portfolio (A+B)
1	0.03	-28.9	1.1	-27.8
2	0.02	-8.9	1.1	-7.8
3	0.03	1.1	-28.9	-27.8
4	0.02	1.1	-8.9	-7.8
5	0.9	1.1	1.1	2.2

The 95% VaR for bonds A and B is 8.9. The 95% VaR for the portfolio is 27.8. That is, the VaR of the portfolio is greater than the sum of the individual VaRs for each bond. Diversification within the portfolio is not recognised under VaR.

Another criticism of the usual adopted VaR measure is that the confidence intervals used are so far down into the tail as to be meaningless for analysis. It may well make sense for a bank to maintain enough capital to sustain a 1 year in 1,000 event, but for a managed fund, which doesn't maintain capital, its clients (and business) would have disappeared long before its losses got anywhere close to this level. So what is a more appropriate measure of tail risk?

A proposed consistent measure of tail risk (CVaR)

A more appropriate measure of tail risk is the use of Conditional Value at Risk (CVaR)[14], see Artzner et al (1997). CVaR answers the question, given that things are bad, of how much worse can I expect them to get? More technically, it measures the extent of the expected loss that is greater than the VaR. CVaR is a measure of the tail, which helps to describe the shape of the distribution rather than just putting a value on a single point. One criticism of CVaR is that the tail is equally weighted and as such ignores to some extent the differing degrees of risk aversion of investors, see Dowd (2005).

The diagram below shows a portfolio distribution and highlights the difference between VaR & CVaR.

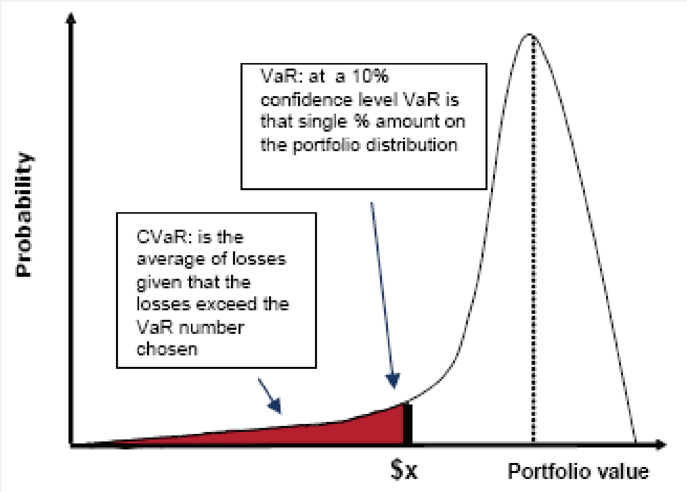

Just as tracking error has come to be a well understood and uniformly described statistic, it is necessary that a similar uniform statistic be adopted within the industry for the measurement of tail risk. Without an accepted standard measure, comparison across portfolios would not be possible.

It is important to note that, even with a standardised measure, underlying assumptions need to be understood as variances here can lead to different results.

It is suggested that a 1 year 20% tail is an appropriate market standard for this measure for the following reasons:

- Given that most credit funds in the market are sold with a suggested investment timeframe of 3-5 years, a tail that measures what might be expected to happen once every five years is appropriate.
- A measure of 20% does more than just measure the extreme event and, because it is measuring so much of the distribution, it inherently picks up some of the unexpected loss characteristics of the distribution.

Once the industry has agreed on the part of the distribution to measure, it is essential that the industry is consistent in what it is measuring. A measure of the tail is proposed that helps to describe the shape of the distribution, rather than just putting a value on a single point of the return distribution.

A real example

To demonstrate the usefulness of the tail measure against more traditional measures, we have modelled a 'low tracking error' portfolio against its benchmark portfolio. In this case the 1 month annualised tracking error was calculated as 32 basis points. The size of the tracking error portfolio is AUD2.5billion.

Using an internationally recognised global broad bond index and a low tracking error portfolio was constructed to replicate the key parameters of the index. The key parameters which were matched so as to create the tracking error portfolio are the duration, the rating of each exposure and exposure to each region and sector. The tracking error and benchmark portfolio characteristics are shown below. Average exposure is the average size of the holdings of each issuer, likewise maximum exposure is the largest holding.

	Tracking	Benchmark
Issuers	41	517
Exposures	53	1,812
Average Exposure (bps)	189	6
Maximum Exposure (bps)	740	32
Duration	4.31	4.29
Average Credit Rating	BBB	A / BBB

Now consider the credit portfolio statistics calculated from Moody's KMV Portfolio Manager:

	Tracking	Benchmark
Expected Loss[15] (bps)	14	9
Unexpected Loss[16] (bps)	175	112
1 year 20% VaR (bps)	92	63
1 year 20% CVaR (bps)	1,010	378

The following diagram shows the tails of the two distributions with the critical region highlighted.

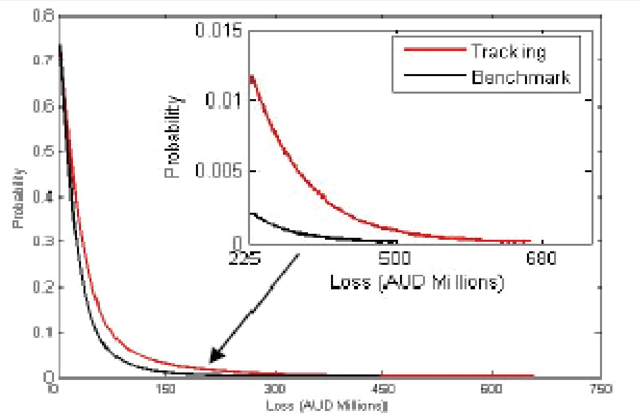

It can be seen that the tracking portfolio, despite its overall characteristics being consistent with the benchmark, and despite it running an ex-ante tracking error of only 32 basis points, is much more risky than the benchmark.

It is not suggested that many managers would have such a concentrated portfolio, but that a portfolio with a very low tracking error might be carrying a very large amount of risk relative to the benchmark.[17]

It must be remembered that in looking at this portfolio going forward, and in tracking its performance, that it may not experience any defaults and may well track the index portfolio very well. However, this does not mean that it did not take significant credit risk in doing so.

Summary of risk measures

The following table presents a brief summary of the risk measures discussed in this paper and highlights their relative advantages.

Conclusion

There are many tools currently used to measure credit risk. Many of these, while benefiting from simplicity of calculation, are significantly inadequate in answering the question "How much credit risk am I taking?". Some, such as average credit ratings, are not really risk measures at all. Credit duration and tracking error are measures of market risk rather than credit default risk. While Credit VaR is a major improvement, it fails to tell us anything about the fatness or length of the tail that it is measuring.

This paper has proposed the use of CVaR as a market standard for the measurement and, importantly, reporting of credit risk. It proposes that CVaR be measured at the 1 year 20% level as an appropriate level for the funds management industry. Such a benchmark is necessary to enable consistency across the industry and improve reporting transparency and overall credit risk management practices.

Measure	Measure incorporates:				
	Default probability	Concentration effects	Term to maturity	Effect of diversification	Extreme event risk
Average rating	✓	×	×	×	×
Credit duration	×	×	✓	×	×
Tracking error	✓	×	✓	×	×
VaR	✓	✓	✓	×	×
CVaR	✓	✓	✓	✓	✓

References

Artzner, P., Delbaen, F., Eber, J. & Heath, D. (1997), “Thinking coherently”, RISK, 10 (11), 68-71

Caouette, J. B., Altman, E. I. and Narayanan, P. “Managing Credit Risk”, New York, 1998, John Wiley & Sons, Inc

Dowd, K. “Measuring Market Risk”, John Wiley & Sons, 2005

Dynkin, L., Gould, A., Hyman, J., Konstantinovsky, V. and Phelps, B. “Quantitative Management of Bond Portfolios”, Princeton University Press, 2007

Embrecht, P., Kluppelberg, C. and Mikosch, T. , “Modelling Extremal Events for Insurance and Finance”, Berlin, Heidelberg: Spinger Verlag, 1997

Fabozzi, F. J. (Ed), “The Handbook of Fixed Income Securities”, McGraw Hill, 2005, 7th Edition

Glasserman, P., “Monte Carlo Methods in Financial Engineering”, Berlin, Heidelberg, Springer-Verlag, 2003

Ranson, B. , “Credit Risk Management”, Thomson/Sheshunoff, 2003.

End notes

[1] Many examples exist, see:

- Artzner, P., Delbaen, F., Eber, J. & Heath, D. (1997), "Thinking coherently", RISK, 10 (11), 68-71
- John B. Caouette, Edward I. Altman, and Paul Narayanan, "Managing Credit Risk", New York, 1998, John Wiley & Sons, Inc
- Dimitris Chorafas, "Credit Derivatives & The Management of Risk", New York, 2000, Prentice Hall
- Credit Derivatives Insights, Single Name Instruments & Strategies, Morgan Stanley 2006, various authors

[2] *Rating:* An alpha scale showing an assessment of the likelihood of default usually published by an independent rating agency such as Moody's Investor services, Standard & Poor's (S&P) or Fitch. The following table outlines the various scales and gives a brief description of each.

Assessment	Moody's scale	S&P & Fitch scale	
Extremely Strong	Aaa	AAA	Investment Grade
	Aa1	AA+	
	Aa2	AA	
Very Strong	Aa3	AA-	
	A1	A+	
	A2	A	
Strong	A3	A-	
	Baa1	BBB+	
	Baa2	BBB	
Adequate	Baa3	BBB-	
	Ba1	BB+	Speculative Grade or High Yield
	Ba2	BB	
	Ba3	BB-	
	B1	B+	
	B2	B	
	B3	B-	
Currently highly vulnerable / defaulted	Caa / D	CCC / D	Distressed / Default

[3] *Default probability:* the likelihood of any borrower defaulting in a given period (usually one year). Thus, based on historical observations, a 'BBB' rated security has a 0.17% likelihood of defaulting in the next year

[4] Correlation: a measure of how two factors move together. Thus default correlation measures the impact of one borrower defaulting on another borrower

[5] The telco sector experienced defaults in 2001. Globally this sector accounted for approximately 25% of defaulted debt by volume (USD17 billion) and 16% by issuer count, this is compared to an average default rate of 3.7% for all issuers .See Moody's Investor Services "Default and Recovery Rates of Corporate Bond Issuers" February 2002. Also see, Moody's Investor Services "Default Rates: Yes, it's as Bad as it Looks", July 2001.

[6] The oil and gas sector experienced defaults in 1999. 18 issuers defaulted. See Moody's Investor Services " Historical Default Rates of Corporate Bond Issuers, 1920-1999", January 2000.

[7] Diversification: the spreading of investments across a variety of different securities or borrowers.

[8] Moody's Investor Services "Default and Recovery Rates of Corporate Bond Issuers, 1920-2004" January 2005.

[9] Duration: a measure of the average life of a fixed income investment incorporating coupon and principal flows. The duration of a security provides a good guide to its expected price change given a change in its yield.

[10] Tracking error: A measure of how closely a portfolio is expected to track (or historically has tracked) its benchmark. Mathematically it is the standard deviation of the excess returns of the portfolio over the index.

[11] Tail Risk: The risk of an extreme negative event.

[12] VaR: Value at Risk (VaR) measures the maximum potential loss on a group of securities over some time period, given a specified probability. In other words, once a probability or degree of confidence has been set, VaR is the amount that represents the statistical maximum loss for a single security or group of securities.

[13] Technically, this is violating the sub-additivity property. The sub-additivity property is $VaR(A)+VaR(B) \geq VaR(A+B)$

[14] CVaR: Conditional Value at Risk (CVaR) measures the expected loss given that VaR has been exceeded. In layman's terms it answers the question "given that things are bad, how bad can I expect them to get?". A unified measure of tail risk defined as:

$CVaRq = VaRq + E(X - VaRq|X > VaRq)$ where VaR is Value at Risk at the q% probability level.

[15] Expected loss: the amount an investor would expect to lose, over the longer term, by investing in credit. It is calculated as the probability of default multiplied by the amount of loss in the event of a default. Thus if the expected loss given a default of a 'BBB' rated issue was 50c in the $, the expected loss would be 0.17% x 50%, thus 0.085% or 8.5 basis points.

[16] Unexpected loss: The volatility of loss around the expected loss.

[17] To model the funds the following assumptions were made:
- The investment horizon is 1 year
- The probability level of interest is 20%
- Expected maturity dates are used allowing the modelling of migration risk
- Default probabilities are derived from the issuer ratings. A Moody's preference and if unavailable, then S&P

For further information cfsinstitutional@colonialfirststate.com.au

Regional Managing Director, Australia and New Zealand

Joanna Davison	+61 2 9303 7007

Institutional Business Development and Consultant Relationships

Matt Russell	+61 2 9303 7024
Ali Karmali	+61 2 9303 6070
Tracey McNaughton	+61 2 9303 2695

Institutional Relationship Management

Liz Krajewski	+61 2 9303 2927
Daniel Bristow	+61 2 9303 6311
Helen Squadrito	+61 2 9303 6142
Hugh O'Neill	+61 2 9303 6116
Nadene Moore	+61 3 8628 5615
Simon Good	+61 3 8628 5681
Lucas Rooney	+61 7 3328 5875