An actuarial cross-dress: general insurance concepts for managing credit investment risk

Executive Summary

Actuaries are becoming increasingly specialised in practice areas but techniques from one specialism can be applied in others. This paper attempts to do so by using some of the mathematical concepts commonly applied in general insurance work to the specific case of investment risk analysis for credit (corporate bond) portfolios.

The key findings of this paper, based on industry data and tentative analysis are summarised as follows:

- Credit ratings are fairly stable in the short run, with invariant transition rates having high probabilities in general.
- A simple Markov chain analysis shows that credit quality degradation can be expected in equilibrium.
- Active management of credit portfolios is essential to ensure that default risk is controlled.
- Risk management could be further enhanced by the use of “reinsurance” e.g. by means of credit default swaps.
- Individual loss modelling based on past recovery rates is be no simple task.
- A compound distribution approach for aggregate portfolio default losses could prove useful but a simulation approach would be more suitable than an analytical one.
- Further research on the extent of capital losses given default would be useful in this context. This should facilitate aggregate loss modelling.

Introduction

Corporate bonds, corporate credit, or just “credit” for short is attracting increasing attention as an asset class in its own right. This study aims to take a more incisive look at the nature of the default risk associated with a portfolio of corporate bonds. We do this from the point of view of the investor in company debt, rather than of the issuer. Given that many investors will seek a global exposure to this sort of investment, our analysis will be global in perspective. However, the underlying principles can be generalised to more focused portfolios (and arguably should be) or investment in other forms of debt.

This paper proceeds as follows:

- In section 1 The Uncharted Waters of Credit Risk we discuss the nature and elements of the default risk in a credit investment portfolio.
- In section 2 Credit Rating Transitions we look at the dynamics of credit ratings and how these evolve over time.
- In section 3 Individual Losses we briefly discuss some features of capital losses based on recovery rates.
- In section 4 Aggregate Loss we ponder how aggregate loss distributions could be modelled.

As this paper's title implies, some of the actuarial techniques used to model general insurance problems and risks will be employed here. These include the use of Markov chains for credit rate transition analysis and the potential use of aggregate loss distribution modelling for capital loss distributions.

In terms of empirical data, and the insights that may be drawn from it with regard to the above, we shall draw on information published by Standard & Poors. Given that the empirical data for credit rate transitions is more ample than for the other aspects that interest us, this will be the primary focus of our study.

1. The Uncharted Waters of Credit Risk

Investment experience is very often modelled based on two key parameters: risk, usually interpreted as the variation in return, and the expected return. We shall focus on the analysis of risk alone, rather than return, and, in particular, on the risk of capital loss due to default. Rather than the risk of worse-than-expected performance, or capital losses due to, say, increases in market yields, this can be characterised as a structural, extreme downside risk for this kind of investment portfolio.

A portfolio of corporate debt issues is not unlike a portfolio of non-life insurance policies. Claims may or may not happen, according to some statistical frequency process. This is analogous to the frequency of default occurring.

Given that a claim on an insurance policy is made, we are interested in the size of the claim, in some cases relative to a sum insured. This can be likened to the conditional size of the credit default compared to the nominal value of the bond, be it total or partial.

We begin our analysis with a look at the frequency or probability of default. Although credit ratings are not updated in real time, they are typically taken to be a reliable indicator of the probability of default. A better indicator may be the yield of an issue given its last known credit rating, and other features of the debt issue, as yields will update on listed markets as trades are made and are a direct reflection of perceived risk. However, to keep our analysis tractable, we shall run on the assumption that credit ratings are a reasonable indicator of the propensity to default. This assumption appears to be reasonable based on a analyses of so-called Gini coefficients, as reported in a paper by Standard & Poor's (S&P)[1]. If this assumption holds, one only needs to find a mapping from credit ratings to default probabilities in order to pin down or parameterise models for the frequency of default.

One problem with this approach is that credit ratings are not stable: they evolve over time. In the S&P paper mentioned above, the historic rates of transition between different credit ratings are presented in considerable detail, by time period, sector and by geography. The next section will look at global rates and their properties in more detail.

It follows to investigate losses given default. On default it is rarely the case that the capital loss on a given bond will be total, as efforts will be made to compensate bondholders in extreme events such as issuer wind up by distributing assets on which the issues are secured (if any). In less extreme circumstances, it may be that a single coupon is not paid on time. Severity distributions for potential losses can be defined as $1 - R$ where R is the recovery rate, itself a random variable.

Once good models for both the frequency and severity of default have been defined, they can be combined to build an aggregate loss model. Such a model ought to work well to the extent that the credit investments in the portfolio are reasonably homogeneous. It is proposed that such modelling should provide a more incisive tool with which to measure and manage default risk than the simple, single distribution models that might otherwise be used. The analogy with modelling aggregate claims on a portfolio of general insurance policies is clear.

---

1 See pp. 6, 61-66 and Appendix III, Vazza, D. et al *2009 Annual Global Corporate Default Study and Rating Transitions*, Standard & Poor's, March 2010.

2. Credit Rating Transitions

As mentioned above, ample data is available on credit rating transition rates. Our first question relates to the stability of these rates over time. If they are not reasonably stable, then modelling default frequencies for a portfolio on an ex-ante basis may prove problematic.

S&P has calculated average credit transition rates for the period 1981-2009. Their stability can be investigated in different ways. These include measuring the variation in these rates over time and verifying that single-year rates are consistent with longer term rates. We shall also look at the possibility that historic rates might imply some kind of Markov equilibrium in rating transitions.

The following table extracted from the S&P paper gives global, average transition rates for corporate credit over the period mentioned. The numbers in parentheses give the standard deviation of the rates, as calculated by S&P.[2]

Global Corporate Average Transition Rates, 1981-2009 (%)

From/to	AAA	AA	A	BBB	BB	B	CCC/C	D	NR
One year									
AAA	88.21	7.73	0.52	0.06	0.08	0.03	0.06	0.00	3.31
	(5.09)	(4.84)	(0.87)	(0.18)	(0.26)	(0.20)	(0.40)	(0.00)	(2.41)
AA	0.56	86.60	8.10	0.55	0.06	0.09	0.02	0.02	4.00
	(0.54)	(4.87)	(3.99)	(0.75)	(0.26)	(0.25)	(0.07)	(0.08)	(1.92)
A	0.04	1.95	87.05	5.47	0.40	0.16	0.02	0.08	4.83
	(0.14)	(1.16)	(3.47)	(2.13)	(0.50)	(0.36)	(0.07)	(0.12)	(1.96)
BBB	0.01	0.14	3.76	84.16	4.13	0.70	0.16	0.26	6.68
	(0.07)	(0.24)	(2.34)	(4.44)	(1.80)	(1.05)	(0.25)	(0.27)	(1.86)
BB	0.02	0.05	0.18	5.17	75.52	7.48	0.79	0.97	9.82
	(0.06)	(0.16)	(0.40)	(2.44)	(4.94)	(4.78)	(0.93)	(1.06)	(2.92)
B	0.00	0.04	0.15	0.24	5.43	72.73	4.65	4.93	11.83
	(0.00)	(0.13)	(0.38)	(0.34)	(2.59)	(5.25)	(2.64)	(3.27)	(3.07)
CCC/C	0.00	0.00	0.21	0.31	0.88	11.28	44.98	27.98	14.37
	(0.00)	(0.00)	(0.74)	(1.05)	(1.34)	(7.86)	(12.81)	(12.90)	(7.57)

One observation that can be made immediately is that the credit ratings appear to be very "sticky", as the one year invariance rates along the lead diagonal are all quite close to one. Further, they are stickier the better the credit rating is, as the rates gently reduce down the diagonal. The chances that a rate will change by one full notch or more over a single year (off-diagonal entries) are correspondingly very small, although they increase as the initial rating worsens.

The standard deviations of the invariance transition rates do not appear to be grossly disproportionate. However, we note that the coefficients of variation are not small in general, particularly for smaller cross-rating transition rates e.g. AA to AAA (coefficient of variation greater than ½).

At a cursory glance, all of this might tend to suggest that the significant transition rates and rating themselves should remain fairly stable over time.

The actuarial eye will immediately see this matrix as a representation of a Markov chain. The states of the Markov chain are simply the S&P credit ratings. If we invert the matrix, so that the columns, rather than the rows, sum to one, then standard Markov analysis can be carried out. We shall do this rounding all of the historical average rates to the nearest 0.01 to avoid spurious

---

2 Ibid, p. 53.

accuracy and, where the columns do not sum to 1, we adjust the transition rates to NR (not rated) to ensure they do. The S&P study does not give rates for transition away from the states D (default) or NR, so we shall assume for now that all rates away from these are zero and the invariant rate is one in both cases (i.e. an issue that has defaulted remains in default indefinitely and non-rated issues remain non-rated). The following transition matrix results:

*Transition matrix, **T***

	FROM								
TO	*AAA*	*AA*	*A*	*BBB*	*BB*	*B*	*CCC/C*	*D*	*NR*
AAA	0.88	0.01	0	0	0	0	0	0	0
AA	0.08	0.87	0.02	0	0	0	0	0	0
A	0.01	0.08	0.87	0.04	0	0	0	0	0
BBB	0	0.01	0.05	0.84	0.05	0	0	0	0
BB	0	0	0	0.04	0.76	0.05	0.01	0	0
B	0	0	0	0.01	0.07	0.73	0.11	0	0
CCC/C	0	0	0	0	0.01	0.05	0.45	0	0
D	0	0	0	0	0.01	0.05	0.28	1.00	0
NR	0.03	0.03	0.06	0.07	0.10	0.12	0.15	0	1.00

We can apply the classic no claims discount problem from general insurance here, where we wish to solve for the equilibrium proportions in each NCD state. Instead, we wish to solve for the vector **Π** of equilibrium proportions of credit issues in the portfolio in each rating state. Recall that this is done by solving the system of equations **T.Π** = **Π** and imposing the condition that the sum of the elements of **Π** must equal one.

This problem can be solved by finding the inverse of the matrix **T'**, which is identical to **T** except that one is deducted from each of the entries along the lead diagonal and one of the rows is replaced by a row of ones, as shown:

Modified Matrix representing Markov equilibrium equations, **T'**

-0.12	0.01	0	0	0	0	0	0	0
0.08	-0.13	0.02	0	0	0	0	0	0
0.01	0.08	-0.13	0.04	0	0	0	0	0
0	0.01	0.05	-0.16	0.05	0	0	0	0
0	0	0	0.04	-0.24	0.05	0.01	0	0
0	0	0	0.01	0.07	-0.27	0.11	0	0
0	0	0	0	0.01	0.05	-0.55	0	0
0	0	0	0	0.01	0.05	0.28	0	0
1	1	1	1	1	1	1	1	1

The vector **Π** will appear in the last column of the inverse of **T'**, which, in this case, is a column of zeroes with a final entry of 1. This implies, of course, that all issues will end up being non rated in the long run – arguably a foregone conclusion!

This was, of course, an exercise in futility: transition rates of 1 appearing in the matrix where they do act as “traps” for migrating issues. Setting these to 1 in the first place is highly questionable. For example, an issuer might default on a single coupon once and then never default again and a given issue may remain temporarily non-rated for any number of reasons.

Be that as it may, longer term empirical statistics produced by S&P do suggest a very strong “pull” toward default and non-rated status. Average 20-year global corporate transition rates from 1981 to 2009 show significant proportions in default and very large proportions settling in non-rated territory

– above 0.5 for all initial states other than AAA and CCC/C.3 So while the analysis above may be spurious, there is some empirical confirmation that its results make some sense.

We can simplify our problem by focusing on the preoccupations of most institutional investors: ensuring that their portfolio, or at least the bulk of it, remains rated investment grade i.e. BBB- or better. We therefore need to truncate the matrix T and in doing so make another assumption i.e. that all of the transition rates below BBB can be aggregated into one "super state" – sub-investment grade – of BB+ or worse. For simplicity, this super state includes the ratings NR and D in the following simplified transition matrix:

*Truncated matrix, **U***

	AAA	*AA*	*A*	*BBB*	*Sub-IG*
AAA	0.88	0.01	0	0	0
AA	0.08	0.87	0.02	0	0
A	0.01	0.08	0.87	0.04	0
BBB	0	0.01	0.05	0.84	0.05
Sub-IG	0.03	0.03	0.06	0.12	0.95

Solving for the Markov equilibrium state probabilities as described above renders the vector:

$[0, 0.01, 0.08, 0.24, 0.67]^t$

In other words, a portfolio left to run its own course is expected to experience a marked degradation in credit quality in the long run, based on ex-post transition statistics. About two thirds of the portfolio can be expected to end up in junk bond or non-rated territory on this basis.

Fanciful number-crunching of this kind has its weaknesses but it does provide some support for claims often made by professional fund managers. For example, they insist on active management and specialist analysis for credit portfolio, emphasising that the first priority of good credit management lies not so much in chasing the upside but rather in avoiding downside events. If there is something in the naïve numerical analyses above, these claims would seem to be justified.

Modelling the incidence of default based on credit rating evolution is clearly cut out to be somewhat more complicated than projecting from the past. Additional assumptions about the ongoing stability of a portfolio's unique credit rating distribution would be needed. The relevance of past transition rates for the future also has to be called into question as does the strength of the assumption that the transition rates themselves are reasonably stable.

In practice, it is unusual for a portfolio to be constructed and then managed on a buy-and-hold basis. As issues mature or fall out of favour they would need to be replaced by other issues in a much more dynamic way. However, the use of focused, skilful active management of such portfolios, especially where based on clear and enforced construction rules, may justify the use of equilibrium credit rating exposure assumptions. If so, mapping from credit rates to default probabilities is a simple task, as corresponding default probabilities are well documented and seem to exhibit fairly consistent patterns over time e.g. very low rates of default for investment grade ratings which become material and increase significantly as the rate sinks to BB and lower.[4]

---

3 Ibid, p. 82.

4 See for example Ibid, p. 9.

On another note, risk management techniques for credit portfolios might also take some inspiration from those used in non-life insurance. Where risks become “bad risks”; that is, where the probability of claim becomes large, an insurer is likely to avoid them, decline to renew coverage of them or reinsure them. The lesson for portfolio management is that default probabilities both now and in future be modelled and monitored. Risk avoidance would take the form of deciding not to buy the issue, declining renewal would correspond to selling the offending issue and using reinsurance might be mirrored by the purchase of loss mitigating derivatives such as credit default swaps, where the price is right. Further analysis here is beyond the scope of this paper.

3. Individual Losses

As indicated above, the severity of capital loss can be investigated with reference to recovery rates given default. A detailed analysis of the distribution of potential capital losses 1-*R* would require access to plentiful data. This might allow severity distributions for individual issues to be calibrated, in the same way one might for claims distributions on insurance policies.

Ex-post analysis performed by S&P shows that the losses on default have some predictable features. Firstly, the losses are not always as severe as one might expect, as recovery rates are fairly high on average. From 1987 to 2009, average nominal recovery rates have oscillated between about 30% and 90%. Unsurprisingly, poorer recovery rates have been observed in times of economic stress.[5] In addition, recovery rates appear to have a negative relationship with default rates. That is, the higher the rate of default, the lower the recovery rate and, correspondingly, the higher the rate of capital loss.[6]

What is perhaps less obvious is the observation that recovery rates appear to have a bimodal distribution, with '20% of instruments experiencing recovery rates of 0-10%, while 14% of instruments have recoveries of par or higher.'[7] This, combined with the putative relationship of capital loss to default rate makes for interesting potential for statistical modelling. It should be noted that, logically, recovery rates will have a systematic dependency on the type of debt instrument held; in particular on both its repayment priority (senior or subordinated) and whether and how the debt is secured on issuer assets.

Further analysis here would be very useful but is not attempted here.

5 Ibid, p. 68.

6 Ibid, p. 67.

7 Ibid, p. 68.

4. Aggregate Losses

Elegant, theoretical solutions for modelling aggregate loss distributions have been formulated for general insurance. These involve combining a claims arrival or frequency process with a claim severity distribution. Models such as the compound Poisson processes will be familiar to readers. We could imagine that default frequencies could be modelled using a Poisson process for rare events and the capital losses on default approximated by either exponential or gamma distributions, as per textbook theory. The distribution of the aggregate loss

$$\sum_{i}^{N} L(i)$$

where *L(i)* is the i^{th} of *N* losses, would then have a gamma distribution.

However, our discussion above would seem to call for more a more precise analysis. The heterogeneity of default rates by credit rating, issuer and issue type calls for separate loss modelling for different parts of the portfolio.

In addition, we have seen that the individual loss on default is, in fact, contingent on the default rate, represented in the frequency process. The theory underlying compound probability distributions requires that the frequency and severity distributions be independent. The nature of the risks at hand would therefore need to be simulated taking account of these dependencies in a more sophisticated way. Clearly, no two portfolios can be modelled in the same way, as the contents of each portfolio will call for a unique approach to aggregate loss modelling.

Single period (i.e. single year) analysis should be feasible. For example, calculations of aggregate, single year value-at-risk due to default could be simulated without too much trouble, provided that the portfolio is not unduly complex. Multi-period analysis may prove more challenging if portfolio structure cannot be assumed to be stable. Our observations of credit rate transitions show just how radically portfolio composition might change without active intervention.

Conclusion

Our observations of credit investment, partly from an actuarial perspective, have uncovered the following insights:

- Credit ratings themselves appear to be fairly stable in the short run, with invariant transition rates having high probabilities in general. The chances of ratings moving by a full notch or more over a single year are generally low.

- The error in average, ex-post transition rates appears to be reasonably low for invariant rates but it is comparatively high for the others. This is of less significance in Markov analysis where the larger transition rates will dominate equilibrium calculations.

- A simple Markov chain analysis indicates that a significant slide to poor credit quality on a laissez-faire basis can be expected to occur in long-run equilibrium. This appears to be confirmed by longer term transition statistics.

- This validates claims that active management of credit portfolios is essential to ensure that quality criteria are met and default risk is adequately controlled.

- As in general insurance, risk management could be further enhanced by the use of “reinsurance” e.g. by means of credit default swaps or other derivatives where viable.

- Empirical data suggest that individual loss modelling based on past recovery rates would be no simple task. This is mainly because recovery rates are correlated with default rates and the distribution of overall recovery rates is non-trivial (bimodal).

- A compound distribution approach for aggregate portfolio default losses could prove useful, particularly for one-year value-at-risk calculations or similar. However, it is likely that a detailed simulation approach would be more suitable than an analytical one, and that separate modelling for reasonably homogeneous sub-portfolios would be necessary.

We would suggest that further research on the extent of capital losses given default would be useful in this context, provided that good data to support analyses of them can be obtained. This could facilitate aggregate loss modelling, especially if it turns out that severity distributions can be modelled reasonably well using simple probability density functions.