Based on notes taken from a
Prototype Model for Portfolio Credit Risk Simulation

Matheus Grasselli
David Lozinski

McMaster University
Hamilton. Ontario, Canada

Issuer Simulation

Transition and Default Probabilities

The model is built from a framework where issuers are ascribed a discrete credit rating that represents an assessment of the issuer's probability of default. For instance, a rating of "4" will convey a certain default probability of the issuer.

It is assumed that every issuer of a certain credit rating has the same probability of changing credit ratings or going into default one year later. Moreover, it is assumed that observed historic frequencies of migration and default are indicative of those underlying probabilities. Where there are years of high default and low default rates, this is assumed to be a result of correlation, and movements in underlying systemic factors. Separate transition matrices for "good years" and "bad years" are provided, but are presumed to have already incorporated in them a bad systemic outcome (this will be discussed more under "Correlation").

The probabilities of migration and default are recorded in a transition matrix listing the credit state at the beginning of the year down the left hand side of the matrix, and the possible credit states at the end of the year along the top of the matrix.

Moody's 1983-2005																			
	Aaa	Aa1	Aa2	Aa3	A1	A2	A3	Baa1	Baa2	Baa3	Ba1	Ba2	Ba3	B1	B2	B3	Caa-C	Def	WR
Aaa	89.54	4.91	1.92	0.31	0.26	0.12	0.04	0.00	0.00	0.00	0.02	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.89
Aa1	2.72	78.03	8.13	6.02	1.04	0.30	0.11	0.11	0.03	0.01	0.04	0.00	0.00	0.00	0.00	0.00	0.00	0.00	3.48
Aa2	0.91	3.42	79.50	7.71	2.67	1.20	0.35	0.09	0.09	0.00	0.00	0.00	0.03	0.02	0.01	0.00	0.00	0.00	4.00
Aa3	0.13	0.51	4.11	79.04	8.28	2.60	0.61	0.20	0.15	0.07	0.01	0.03	0.03	0.02	0.00	0.00	0.00	0.02	4.20
A1	0.06	0.10	0.44	5.70	78.20	7.34	2.73	0.59	0.35	0.10	0.23	0.14	0.06	0.05	0.03	0.00	0.00	0.00	3.89
A2	0.04	0.05	0.28	0.89	5.13	78.03	7.29	2.78	0.89	0.37	0.21	0.10	0.15	0.04	0.05	0.01	0.04	0.03	3.66
A3	0.06	0.09	0.11	0.23	1.82	7.03	74.51	6.35	3.30	1.15	0.49	0.19	0.22	0.12	0.05	0.03	0.03	0.04	4.20
Baa1	0.04	0.04	0.13	0.13	0.28	2.13	6.71	73.33	7.27	3.13	0.94	0.50	0.39	0.51	0.07	0.04	0.12	0.16	4.09
Baa2	0.05	0.10	0.03	0.14	0.22	0.72	3.46	5.81	73.81	6.66	1.70	0.59	0.72	0.58	0.26	0.15	0.23	0.16	4.63
Baa3	0.04	0.01	0.03	0.05	0.17	0.40	0.69	3.21	8.56	70.07	5.57	2.64	1.40	0.68	0.42	0.25	0.63	0.32	4.87
Ba1	0.03	0.00	0.01	0.06	0.23	0.20	0.53	0.77	3.11	9.22	64.46	5.07	4.29	1.37	1.16	0.74	0.38	0.71	7.68
Ba2	0.00	0.00	0.04	0.03	0.04	0.14	0.13	0.43	0.82	2.88	8.78	62.64	7.56	2.57	2.72	1.07	0.74	0.75	8.67
Ba3	0.00	0.02	0.01	0.01	0.03	0.17	0.16	0.23	0.27	0.56	2.93	6.06	64.81	7.08	4.80	2.26	1.03	1.98	7.59
B1	0.02	0.01	0.02	0.01	0.05	0.09	0.07	0.09	0.16	0.23	0.56	2.81	6.19	63.20	9.09	4.03	2.46	3.04	7.88
B2	0.00	0.00	0.01	0.03	0.03	0.02	0.11	0.16	0.10	0.18	0.43	0.72	2.04	7.80	60.51	8.27	6.21	5.17	8.24
B3	0.00	0.01	0.06	0.00	0.01	0.03	0.08	0.09	0.08	0.16	0.11	0.35	0.78	3.26	5.88	58.66	11.23	9.07	10.13
Caa-C	0.00	0.02	0.00	0.02	0.00	0.03	0.00	0.03	0.09	0.11	0.11	0.09	0.40	0.94	1.22	3.31	59.46	10.41	23.78

When firms without a rating at the beginning or end of a period are removed from the sample, the historic frequencies as percentages can be calculated. Note that every row of the matrix will sum to 100%. There is no such requirement for the sums down the columns.

	Aaa	Aa1	Aa2	Aa3	A1	A2	A3	Baa1	Baa2	Baa3	Ba1	Ba2	Ba3	B1	B2	B3	Caa-C	Def
Aaa	92.21%	5.05%	1.98%	0.32%	0.26%	0.12%	0.04%	0.00%	0.00%	0.00%	0.02%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa1	2.82%	80.83%	8.42%	6.23%	1.08%	0.31%	0.12%	0.12%	0.03%	0.01%	0.04%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%	0.00%
Aa2	0.95%	3.57%	82.81%	8.03%	2.78%	1.25%	0.37%	0.10%	0.09%	0.00%	0.00%	0.00%	0.03%	0.02%	0.01%	0.00%	0.00%	0.00%
Aa3	0.13%	0.54%	4.29%	82.50%	8.65%	2.71%	0.63%	0.21%	0.16%	0.08%	0.01%	0.03%	0.03%	0.02%	0.00%	0.00%	0.00%	0.02%
A1	0.06%	0.10%	0.46%	5.93%	81.37%	7.63%	2.84%	0.61%	0.37%	0.11%	0.24%	0.14%	0.06%	0.06%	0.03%	0.00%	0.00%	0.00%
A2	0.04%	0.05%	0.29%	0.93%	5.32%	81.00%	7.56%	2.88%	0.92%	0.38%	0.22%	0.10%	0.15%	0.04%	0.05%	0.01%	0.04%	0.03%
A3	0.06%	0.10%	0.11%	0.24%	1.90%	7.33%	77.77%	6.63%	3.44%	1.20%	0.51%	0.20%	0.23%	0.12%	0.05%	0.03%	0.03%	0.04%
Baa1	0.04%	0.04%	0.14%	0.13%	0.29%	2.22%	6.99%	76.45%	7.58%	3.27%	0.98%	0.52%	0.41%	0.53%	0.08%	0.04%	0.13%	0.17%
Baa2	0.05%	0.10%	0.03%	0.14%	0.23%	0.76%	3.63%	6.09%	77.40%	6.98%	1.78%	0.61%	0.75%	0.61%	0.28%	0.15%	0.24%	0.16%
Baa3	0.04%	0.01%	0.03%	0.05%	0.17%	0.42%	0.73%	3.37%	9.00%	73.66%	5.85%	2.78%	1.47%	0.72%	0.45%	0.26%	0.66%	0.34%
Ba1	0.03%	0.00%	0.01%	0.06%	0.25%	0.21%	0.57%	0.83%	3.37%	9.98%	69.82%	5.49%	4.64%	1.48%	1.26%	0.80%	0.41%	0.76%
Ba2	0.00%	0.00%	0.04%	0.04%	0.04%	0.15%	0.15%	0.47%	0.90%	3.15%	9.61%	68.58%	8.28%	2.82%	2.97%	1.18%	0.81%	0.82%
Ba3	0.00%	0.02%	0.01%	0.01%	0.03%	0.18%	0.18%	0.25%	0.29%	0.61%	3.17%	6.56%	70.12%	7.66%	5.20%	2.45%	1.12%	2.14%
B1	0.02%	0.01%	0.02%	0.01%	0.05%	0.10%	0.08%	0.10%	0.17%	0.24%	0.61%	3.04%	6.72%	68.61%	9.86%	4.38%	2.67%	3.30%
B2	0.00%	0.00%	0.01%	0.03%	0.03%	0.02%	0.12%	0.17%	0.11%	0.19%	0.47%	0.78%	2.22%	8.50%	65.94%	9.01%	6.77%	5.63%
B3	0.00%	0.01%	0.07%	0.00%	0.02%	0.04%	0.09%	0.10%	0.09%	0.17%	0.13%	0.39%	0.87%	3.63%	6.54%	65.27%	12.49%	10.09%
Caa-C	0.00%	0.03%	0.00%	0.02%	0.00%	0.04%	0.00%	0.04%	0.11%	0.14%	0.14%	0.12%	0.53%	1.23%	1.60%	4.35%	78.01%	13.65%

There are two significant problems with using these frequencies directly as probabilities. One is that there are a large number of zero entries. This is met by replacing those entries with a percentage found using geometric extrapolation from observed non-zero results. The second significant problem is that one expects to see a certain monotonicity in the way the probabilities decrease as one moves away from the main diagonal. That is, for most ratings, the highest probability is that the rating will not change, as shown in bold down the main diagonal of the above matrix. The probability is expected to grow steadily smaller as one considers moving one credit rating up or down, then 2 ratings, 3 ratings and so on. While the general pattern is there in the raw data, the specific results show some inconsistencies. This can be observed by plotting the logarithm of the probabilities of migrating one rating, 2 ratings, 3 ratings etc. An example for the graph of the logarithms of moving one notch up in credit quality (the top line), two notches (the second line) and so on appears as follows:

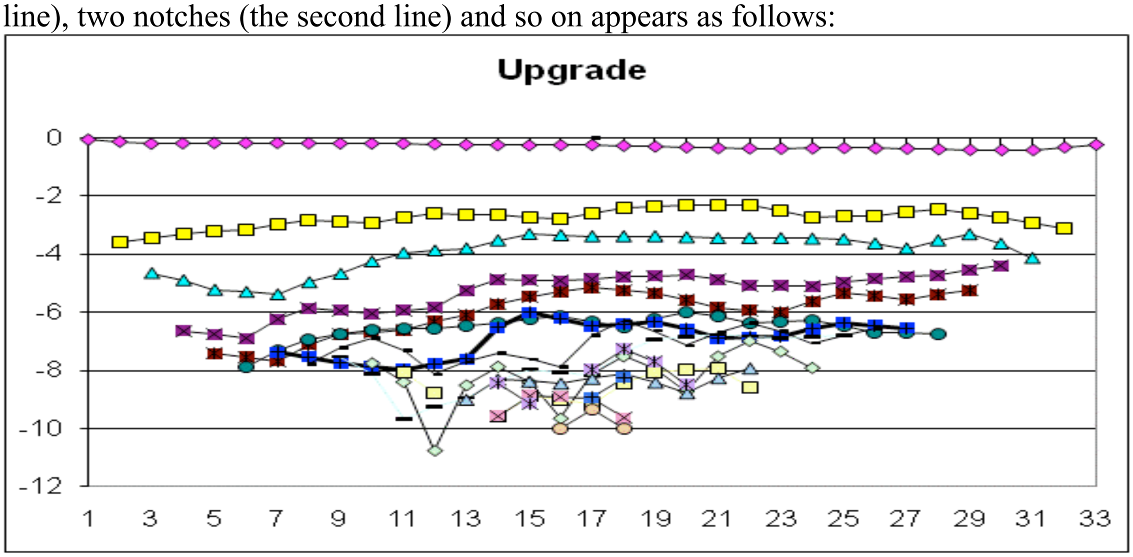

Generally the plot shows the expected pattern. However, where the lines cross, that monotonicity has been violated. The comparable graph for the downgrades is as follows:

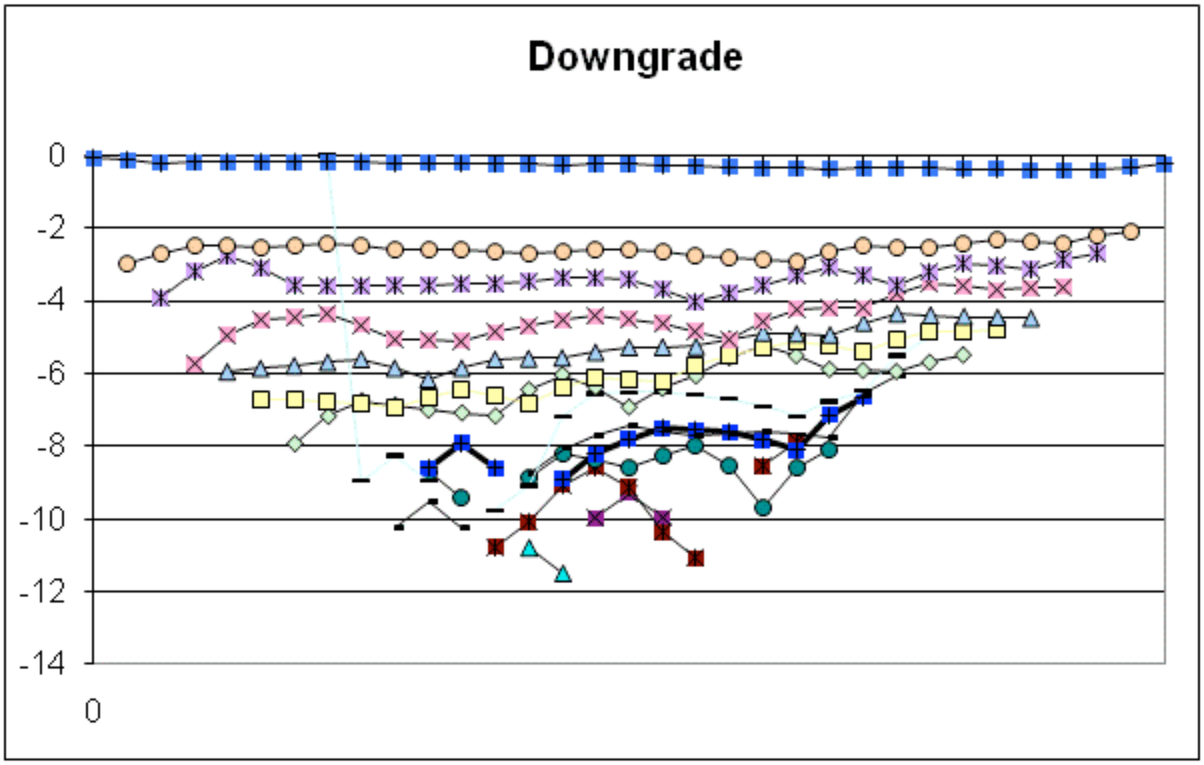

Those zero entries that were replaced with extrapolated values can be readily identified as the perfectly flat lines in the graph. In order to "rationalize" the probabilities, the probabilities are adjusted manually to achieve monotonicity with the intent of dong a minimum amount of alteration. Migration probabilities were generally **increased** in order to achieve this, anticipating that this is the more conservative approach. When this was done, the "untangled" graph for upgrades appears as follows:

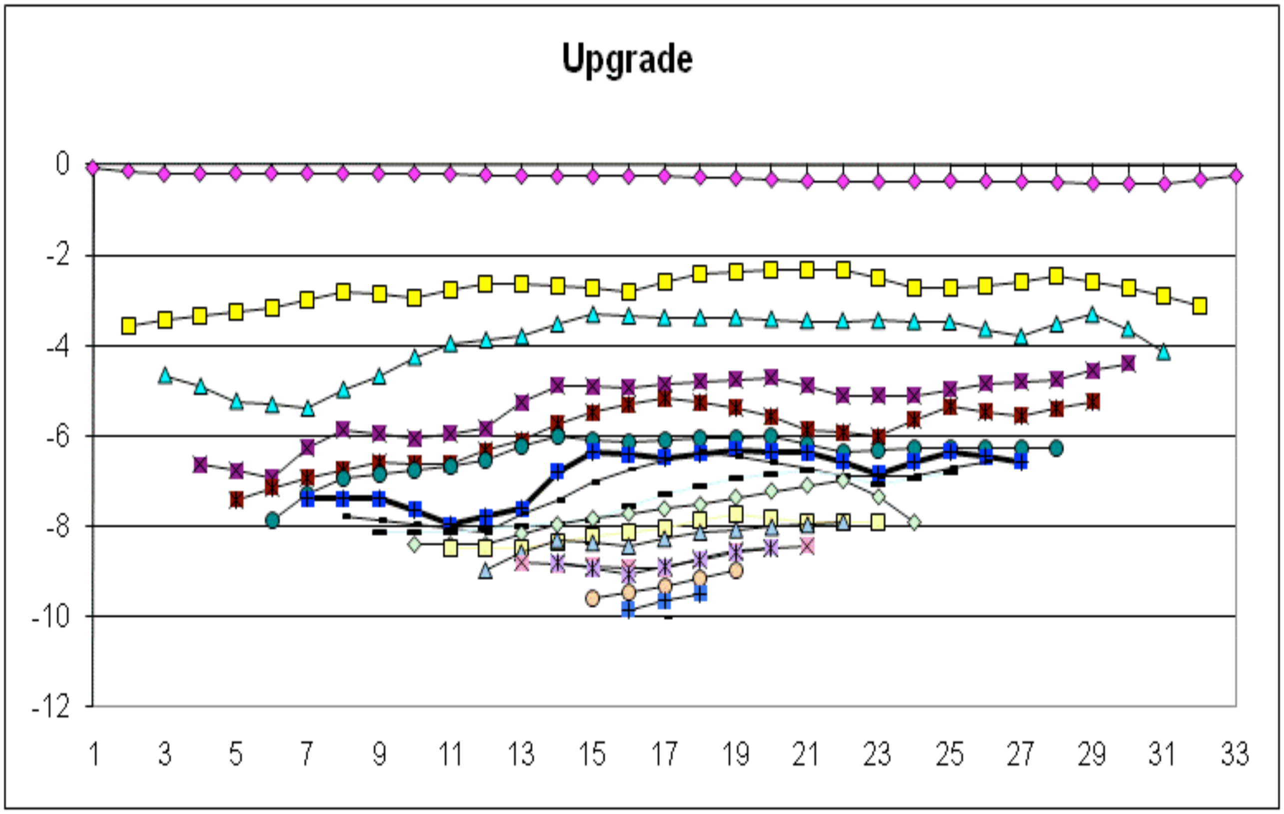

While the downgrade graph is:

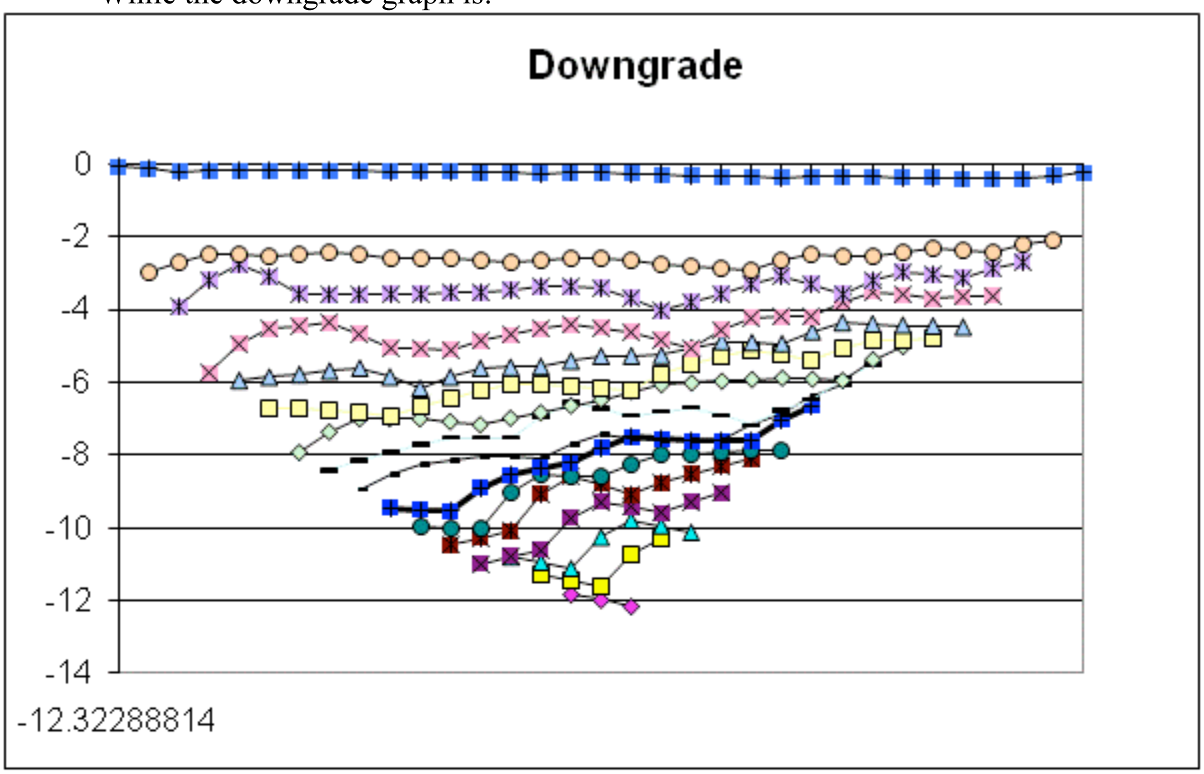

From Moody's published migration frequencies for senior unsecured debt, the average transition frequencies were

Long-Term Average Moody's Transition Matrix Rationalized								
	Aaa	**Aa**	**A**	**Baa**	**Ba**	**B**	**Caa-C**	**Def**
Aaa	92.164%	7.349%	0.420%	0.042%	0.009%	0.004%	0.000%	0.012%
Aa	1.299%	92.344%	5.965%	0.281%	0.067%	0.025%	0.003%	0.016%
A	0.054%	2.758%	90.850%	5.557%	0.571%	0.155%	0.027%	0.029%
Baa	0.031%	0.212%	5.171%	88.093%	5.050%	0.964%	0.258%	0.222%
Ba	0.016%	0.103%	0.631%	6.620%	81.921%	8.606%	0.861%	1.244%
B	0.009%	0.056%	0.229%	0.450%	5.099%	80.506%	7.310%	6.341%
Caa-C	0.005%	0.04%	0.10%	0.29%	0.86%	7.18%	77.88%	13.65%

This can now be compared to the 2003 data from Moody's:

2003								
	Aaa	**Aa**	**A**	**Baa**	**Ba**	**B**	**Caa-C**	**Def**
Aaa	100.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%	0.000%
Aa	1.314%	90.471%	8.215%	0.000%	0.000%	0.000%	0.000%	0.000%
A	0.211%	0.527%	92.624%	6.744%	0.000%	0.000%	0.000%	0.000%
Baa	0.000%	0.000%	0.528%	91.658%	5.597%	2.006%	0.106%	0.000%
Ba	0.000%	0.000%	0.328%	2.620%	81.550%	12.336%	2.074%	1.201%
B	0.000%	0.000%	0.000%	0.331%	3.867%	83.867%	9.392%	2.541%
Caa-C	0.000%	0.000%	0.000%	0.000%	0.000%	7.250%	70.000%	22.750%

In comparing Moody's long-term results to the Moody's results from 2003, it can be seen that the 2003 data shows a worse economic year. Downgrade frequencies are about 130% of their long-term average, while the upgrade frequencies are smaller by about the same factor (1/1.3).

One of Moody's worst reported years for downgrade and defaults was 2002. By applying the appropriate multiplicative factor, a matrix for a bad year is produced. Based on the Moody's results, upgrades were multiplied by 0.7, while downgrade frequencies of more than 10% were divided by 0.7. Downgrade frequencies for those frequencies less than 10% were multiplied by a representative factor of 2. In these calculations, an additional restriction was implemented that no off-diagonal frequency could be larger than the frequency on the main diagonal.

Time Horizon:

The above matrices represent a model for one year transition probabilities. To run the model for longer time periods, the transition matrices can be re-used to simulate subsequent years. The credit rating at the end of one year becomes the credit rating at the beginning of the next. To run the model for periods of time of less than a year, it is necessary to first make the matrix square by adding a line for the transition probabilities of a company in default. The conventional model is to make that a row of zeroes save for 100% in the last column, ensuring that default is an absorbing state (In truth, some firms do re-emerge from bankruptcy. However, where that occurs, the emerged firm is often considered to be a completely new company). The square matrix can then be diagonalized using its eigenvectors. Then, if the transition matrix is A and it is diagonalized as

$$A = P\, D\, P^{-1}$$

a transition matrix for some other time period, t, in years is given by

$$A_t = P\, D^t\, P^{-1}$$

where D^t is the diagonal matrix whose entries are the entries of D raised to the power t.

Simulating Ratings Changes:

Structural Models

- Aim to simulate cause of credit changes (particularly default), typically from movement of a firm's assets
- Most applications model asset returns as normally distributed
- Sections of the normalized distribution of asset returns are aligned with probabilities of default or migration
- Two of the more common approaches involve:
  - use of a transition matrix (Credit Metrics)
  - use of a copula; typically normal or Gaussian

Transition Matrix

- Set of probabilities for discrete changes in credit state

	AAA	AA	A	BBB	BB	B	CCC	Default
AAA	88.65%	10.29%	1.02%	0.00%	0.03%	0.00%	0.00%	0.01%
AA	1.08%	88.71%	9.56%	0.34%	0.15%	0.15%	0.00%	0.01%
A	0.06%	2.87%	90.20%	5.91%	0.74%	0.18%	0.01%	0.03%
BBB	0.05%	0.34%	7.10%	85.28%	6.08%	1.01%	0.09%	0.04%
BB	0.03%	0.08%	0.55%	5.64%	83.50%	8.03%	0.53%	1.64%
B	0.01%	0.05%	0.18%	0.68%	6.89%	82.92%	2.88%	6.38%
CCC	0.00%	0.00%	0.66%	1.05%	3.06%	6.14%	62.97%	26.11%

- Empirical description of a discretized distribution

Transition Matrix

- Probabilities of migration and default subdivide the normalized asset distribution

	Default	CCC	B	BB	BBB	A	AA	AAA
BBB	0.04%	0.09%	1.01%	6.08%	85.28%	7.10%	0.34%	0.05%

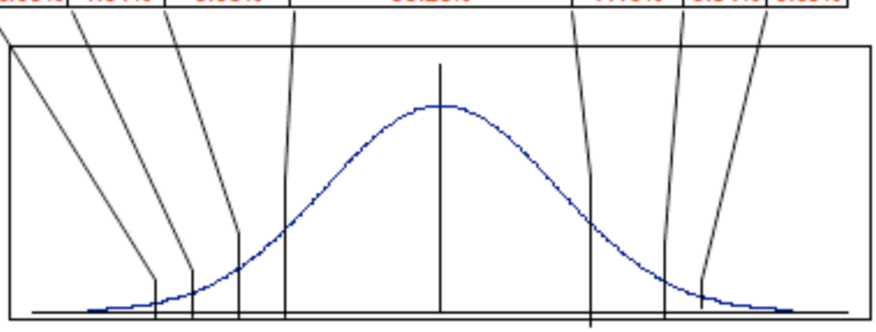

- Draws for different firms are easily correlated
- Process can be repeated to simulate multiple time steps

Correlation:

In modeling changes in the quality of a firm using a normal distribution, including correlation is simply a matter of introducing correlation into those normal draws for the asset returns. However, asset returns are not directly observable. Nonetheless, changes in the value of a firm's assets are well approximated by changes in the value of their equity. By computing the correlations between returns from companies' equity data, those values can be applied to the normalized asset returns.

Equity returns are only available for those companies with public stock trading on an exchange. In order to produce correlations for other firms, those correlation results are used to establish rules for the correlations between other firms.

The correlation rules are based on firms' industry classification and size. The correlation of asset returns (equity returns) observed between two companies in different industries tends to lie in the range of 20% to 35%. The correlation between firms in the same industry will be much higher, in the range of 45% to 65%. The highest observed correlations tend to arise for larger banks within the same country. This gives rise to considerations of firm size when setting correlations. Empirical evidence has supported the assertion that firms of larger size tend to display higher correlations. Generally, larger firms are more closely linked to systemic movements in the market, while smaller firms behave more idiosyncratically. Where there is uncertainty in the precise level of correlations, that gives rise to a range of values that can be mapped to the range of firm sizes. The smallest companies will acquire the smaller correlations in the model's range of correlations, while the larger firms are attributed the larger correlations.

The correlated standard normal variables used in the prototype model are built in a straightforward manner out of underlying independent standard normal variables in order to produce the model correlations. Consider firm "i" in industry "j". It's normalized asset return x_i is N(0,1) and found as:

$$x_i = p_i\, y_j + \mathrm{sqrt}(1 - p_i^2)\, z_i$$

where y_j is an N(0,1) random variable representing changes within the entire industry "j", z_i is an independent N(0,1) random variable representing idiosyncratic changes specific to firm "i", and p_i is a weighting factor between the two. The larger the value for p_i, the more the firm moves with its industry. Smaller firms are expected to have smaller values for p_i. The intent is that the variable z_i is independent of all other random variables. However, the industry factors y_j are themselves correlated with one another. This is achieved by building them as a weighted combination of a global systemic variable, say w_0, and an idiosyncratic component as

$$y_j = q_j\, w_0 + \mathrm{sqrt}(1 - q_j^2)\, w_j$$

where w_0 and the w_j are independent N(0,1) variables. The variable w_0 is used by every industry and can be interpreted as producing global systemic changes, while the w_j are random factors specific to each industry.

The only parameters that are free to be set are the weights p_i and q_j. These are selected to generate the appropriate correlations. For two firms "i" and "k" within the same industry, their asset correlation is

$$\mathrm{corr}(x_i, x_k) = p_i\, p_k$$

As a result, if two firms from the same industry are about the same size and are to have a correlation of 0.65, one would ascribe to each a weight of

$p_i = \text{sqrt}(0.65) \sim 0.806$

In that way, knowing what correlation should arise between any two firms within an industry, the weights p_i for every firm in that industry can be determined. The weights (and their resulting correlations) can be adjusted up or down by the size of the firm as appropriate.

It is the correlation between firms in different industries that determines the values for the weights q_j. For two firms "i" and "k" each in different industries, say "m" and "n", their correlation is given by

$\text{corr}(x_i, x_k) = p_i\, p_k\, \text{corr}(ym, yn) = p_i\, p_k\, q_m\, q_n$

Thus, in a similar manner, where firms from two different industries are to have a correlation of, say 0.35, this would suggest that the industries should have factors such that

$q_m\, q_n = 0.35/(p_i\, p_k)$

The simplest model would ascribe to all average-sized firms the same value p_i, and then the industry weights would be

$q_m = \text{sqrt}(0.35)/p_i$

Section 2: Asset Valuation

Pricing

Performing assets are priced as if they are fixed coupon bonds of some fixed notional (clean price). Each asset is given a coupon rate. If a coupon rate is not provided at input, the asset is assumed to begin the simulation at par, with a coupon matching the yield. Each asset is priced once for each possible final credit state at the specified time horizon. For the defaulted value a price equal to the asset's recovery is used (par minus the loss given default). Subsequent modeling improvements will allow the recovered funds to be reinvested into a new asset trading at par.

Yield

The yield on each asset is assumed to be equal to the appropriate interest rate off a base curve plus a spread that is a function of the credit rating. If an asset is downgraded in a simulation, the spread will increase, and the asset price will be lowered. The magnitude of the reduction in price depends significantly on the remaining maturity. As the remaining maturity diminishes, the asset value is pulled to par provided the issuer does not default.

In the prototype model, the base curve and the spread curves are taken to remain unchanged through time. The model focuses on credit risks only. Market risks arising from general movements in interest rates and market appetite for credit risk of a given quality are assumed to be managed separately.

Recovery:

The simplest first approach for any model is to simply assume a fixed recovery value for each asset. A common assumption based on observed recoveries for corporate bonds is to assume one recovers 50% of the face value of the asset. Future improvements can build more involved models for recovery allowing uncertainty and even correlation with the systemic variables. Although there are many factors that impact recovery, more advanced modeling is typically hampered by a lack of current and historic data on the details of the asset's support.

Section 3: Portfolio Valuation

Value:

Simply enough, for each scenario, the final credit rating for each issuer is read in. Then each asset is priced looked up using the value computed for the issuer's simulated rating. All the assets in the portfolio are summed to produce the portfolio value.

Loss:

The issue of defining a loss is not a trivial one. Comparing the price of the portfolio computed at the time horizon to its initial price (likely par) is the obvious approach. However, this neglects to account for all the cash flows that were generated during the life of the simulation. In addition to the portfolio profit, those cash flows also embody reparations for the expected losses that occur during the simulation. That is, even in the best of times, a certain percentage of the portfolio can be expected to default. The simulation will include those expected losses, but will not offset them with the returns that were to pay for them. In addition, the portfolio simulation will include any "pull to par". That is, if any assets begin the simulation at values different from par (if a coupon is specified), at the end of the simulation, if nothing else happens, the value of the asset will be closer to par due to the shorter maturity.

Distribution:

By collecting the portfolio values from every scenario, a distribution of the portfolio value is created, allowing for an assessment of portfolio losses. Statistics of interest include expected loss, and measures of unexpected loss including loss at specific outlying percentiles and standard deviation.