

ERASMUS UNIVERSITY ROTTERDAM

Erasmus School of Economics

Measuring credit risk of a loan portfolio:
a comparison of different probability of default models

Master thesis for the programme:

Econometrics and Management Science, Quantitative Finance

Student: Mark van Liempt

Student ID number: 312858

Academic supervisor: Karolina Scholtus

Co-reader: Erik Kole

External supervisor: Rogier Swierstra

February 21, 2017

Abstract

A portfolio consisting of loans is extremely sensitive to credit risk, thus many products have been developed over time trying to insure this credit risk. Different models trying to capture credit risk within the portfolio have been developed, often consisting of a probability of default part and a recovery rate part. This paper looks at three main forms of modelling the probability of default, a form of historical simulation, the advanced IRB model and the CreditRisk+ model. I look at the strengths and weaknesses of these models and show the three approaches all converge to a similar outcome when looking at a portfolio spread across different ratings and covers a longer period of time. Solutions to some of the weaknesses are looked at, such as using historical probability of default data to include past knowledge within the advanced IRB model. Lastly a mixed model will be looked at, which is a combination of historical simulations and the advanced IRB model. The different models are validated within this paper, showing all three approaches to be valid options and showing these models can be used by companies depending on other reasons such as using the model specifically for sector loans.

Keywords: A-IRB model, CreditRisk+ model, Historical Simulation, Probability of Default

Contents

1	Introduction	4
2	Literature	6
3	Theory	8
3.1	A-IRB Model	8
3.2	Simulated A-IRB Model	9
3.3	PGGM Model	10
3.4	Mixed Model	11
3.5	PD Estimation	11
3.6	Simplified CreditRisk+	12
3.7	CreditRisk+	13
3.8	Model Validation	14
4	Data	16
4.1	Data validation	18
5	Empirical Results	18
5.1	1 year Ba2 portfolio	19
5.2	4 year Ba2 portfolio	22
5.3	4 year Baa3 to Ba3 portfolio	26
5.4	4 year realistic portfolio	28
5.5	Estimate of the current PDs	32
5.6	Statistical validation	35
6	Conclusion	35

1 Introduction

Estimating the probability of default (PD) for the first few years after a deal is made is essential when pricing credit risk products. Current models available to estimate the PDs and the price of credit risk products tend to either focus entirely on the current state of the world or on the historical data available. When using the current state of the world the outcome can strongly be influenced by the interpretation of the current state of the world. Using historical data assumes that the historical data can directly be used to represent the future. In this paper a comparison between these different types of models will be made. Furthermore, a combination between a model based on the historical PDs and a model looking at the current PDs only is developed to see whether this would solve the weaknesses of the two individual models without losing a good portion of the strengths of these two approaches. Lastly, the economy changes over time which can be seen in bad periods such as the 2007 crisis when compared with good periods, such as the years prior to this crisis. When pricing a product which only covers the coming few years, including the current state of the world in its pricing is essential. To include this the idealised PDs and historical PDs will be replaced by estimating the current PDs through a (Pooled-) AR model.

The techniques based on historical data of the different PDs come in two main forms, the Markov chain type models and the PD over time. The Markov chain type models were the initial steps made to estimate the PDs and the credit at risk, of which the Jarrow-Lando-Turnbull model was one of its pillars. These Markov chain type models look at the different loans within the portfolio and what rating each loan has every year, thus looking at more than whether it went default or not during the relevant period. This requires a proper estimate of not only the PD for each year, but also the probability of loans to migrate to a different rating. The techniques that look at the PD over time look at the rating of a loan from its starting point, and only look at the probability of it going default within the specified number of years. The biggest difference in choice between these two main techniques is whether it is important for the financial product to know the state of the company at the end of the period the credit risk product covers. For example, a strategy that involves selling an insurance early on for a whole batch of loans and selling the insurance onwards after a certain period makes it important to know the possible states these loans can have after this period, thus making a Markov chain model suitable. However, if the model is used to value an insurance which has no option to be traded onwards estimating the current rating of the different loans provides no information relevant to the product, and at the cost of more uncertainty. Using a PD over time method is more suitable here. Since this paper focusses on an insurance that is held till maturity, the choice here is a model that looks at the PDs over time. A realistic periods for these credit risk products can be anywhere between one and five years. While the model should be valid for most types of loans, the main focus is on mid to big companies. This results in loans falling between 50,000 Euro and 20,000,000 Euro.

Since the different Basel regulations, as published by Basel Committee (2004), authority demands the model puts emphasis on the current state of the world, instead of looking at the historical events. This

is one explanation why the models looked at by the Basel committee are based on the currently assumed PDs per loan, and not just how it performed for the past recorded years. While this is straightforward for a single loan, a product covering multiple loans requires incorporation of correlation. For the models using historical data, correlation was implicitly accounted for as the different loans simultaneously have higher PDs in bad historical years and lower PDs in good historical years. To solve this issue, multiple models were developed that looked at ways of including correlation within these portfolios, such as the Advanced Internal Ratings Based approach (A-IRB), also published in Basel Committee (2004), or the CreditRisk+ model, published in Wilde (1997). This paper will look at both approaches as well as a simulated version of the A-IRB model to function better over a short period of time, instead of the longer period this model is developed for.

Two models are combined in this paper, the A-IRB model and the PGGM model. The PGGM model uses historical PDs and assumes that every individual loan follows a binomial distribution. The reason for choosing the A-IRB model is due to the correlation part of this model can be separated from the rest of the model. Alternative models, such as the CreditRisk+ model, have more indirect forms of correlations and cannot be separated so easily for use in a different model. The choice for using the PGGM model as a way of including the historical PD pattern is due to its simplicity and small amount of required assumptions. Since this technique already tries to combine two models, minimizing the number of assumptions is important. The strength of this mixed model is to no longer let one form of correlation be used to describe all correlation present, but instead looks at two forms of correlation. The first form of correlation looks at the correlation all loans have in common, independent of its rating. The second form of correlation looks at a correlation the different loans within a specific rating bracket have in common. Another strength in the mixed model is that it uses the historical data, through the PGGM model, and adjusts it to include knowledge of the current state of the world, through the A-IRB model.

A last step is to perform model validation on the different models tested in this paper. This form of validation follows the guidelines as published by Basel Validation Group (2005) which describes the current requirements in the financial sector by the Basel committee. A good example of how to perform this validation is described by Medema et al. (2009). An adjustment does need to be made, as the validation tested in the example assumes no correlation, which is a distinct element in the models tested for our portfolios. This correlation results in a longer period of data being required to perform the validation process, as an entire portfolio can now only be treated as a single test, not the number of tests equal to the number of loans in the portfolio. This leads to a less accurate validation process, but the longer period of available data compensates this.

The different models are found to give different outcomes when looking at a portfolio of short term loans, here tested one year loans, or when using a portfolio consisting of loans of similar ratings. The more spread a portfolio is going to be, the more similar the found PDs will be across the different models. One thing that does show a distinct difference between the models is the tail of the distribu-

tion, the last few percentiles. The PGGM model and a Simplified version of the CreditRisk+ model show less extreme losses when compared to the other models, potentially underestimating the probability of a severe loss in the portfolio. When looking at model validation most models will have similar outcomes when performing the statistical validation. When looking at data validation and theoretical validation two main groups are found, those that use historical PDs to model the current state of the world and those that use idealised PDs. No indication is found that one approach is superior to the other.

In this paper the described models will be tested starting with a very basic, theoretical portfolio, with the aim of looking at the different outcomes without being influenced by many different properties present in a more complex, realistic portfolio. After this, the portfolio will be expanded by including different properties. These will be extending the time to maturity, combining loans of different ratings and making the loans vary in size within a portfolio instead of equal sized loans. By doing this, an explanation for the core differences in models can be found at the theoretical portfolio, followed by a step by step look at different properties of a portfolio having influence on different models.

2 Literature

The Markov Chain type models were one of the first types of models used in pricing multiple credit risk products, and one of the most prominent models of this type is the Jarrow-Lando-Turnbull model as presented in Jarrow et al. (1997). However, multiple alternative models were developed which made use of the Markov Chain type models, such as discussed in Siu et al. (2005) or Duffie and Singleton (1999), as well as possible extensions on the Jarrow-Lando-Turnbull model, such as discussed in Millosovich et al. (2002). While especially some of the more recently developed models are getting more precise, as well as data being recorded for a longer period of time, the loss of precision by looking at the change of state per loan, instead of the direct PD per loan, comes at the cost of precision. Kiefer and Larson (2004) state that certain credit risk products can be accurately modelled by a Markov Chain type model, up to about five years.

Two of the more prominent alternatives to the Markov Chain type models are the A-IRB model, based on Vasicek (1987), and the CreditRisk+ model, based on Wilde (1997). Lütkebohmert (2009) is a good description of the CreditRisk+ model and some comparison of the positive and negative sides of this technique compared to other models, including the A-IRB model. It emphasizes the strength of the CreditRisk+ model of also being able to include other classes of risk such as sectors or country of risk present in these loans. Schechtman et al. (2004) stresses the importance of the used PDs for both the CreditRisk+ model and the A-IRB model. They also show the current state of the economy influences the short and medium period PDs, which are often the periods used for portfolios of loans. This supports the A-IRB model and CreditRisk+ when compared with alternatives such as the PGGM model or the Markov Chain type models.

Kupiec (2009) looks deeper into the A-IRB model and how it matches the data. The main concern of this paper is that the A-IRB model underestimates the true correlation, especially during stress periods. This is one of the problems which can, at least partially, be solved by the mixed model presented in this paper, as it includes two separate forms of correlation. Levich et al. (2012) compares the A-IRB model and its Basel alternatives, the Foundations-IRB and the Standardized Risk Weights, and concludes that the A-IRB model provided a much smaller probability of extreme losses due to credit risk, especially when compared to the Foundation-IRB model. The mixed model captures this, as one can do this by increasing the correlation either for all loans or only for the more secure loans. Another form of criticism on the Basel models, mainly the A-IRB, is in the papers written by Altman and Saunders (1997) and Altman et al. (2004), which focussed on the performance of the model over the last 20 years. One of the findings is that only using the current rating of a loan is inaccurate as loans recently given a lower rating have a higher probability of descending further in the rating, and indirectly having a higher PD than a loan that is already present in the rating bracket for a longer period.

Besides these more well-known approaches, some interesting alternatives have been looked at in literature. McNeil and Wendin (2007) applied Generalized Linear Mixed Models in a Bayesian framework to estimate credit risk. Similar to the CreditRisk+ model, this technique allows separating the loans not only based on ratings, but also other properties such as the industry involved or the country of risk. A more thorough look at the PDs and how to estimate these PDs can be found in a paper by Hanson and Schuermann (2004) where some of the base estimation techniques used by the S&P are analysed. One of the conclusions is that a proper spread in PDs for the more secure grades is present, but that some of the more risky grades tend to cover a window of PD which is too broad. Another focus within this paper is the window for the PDs estimated, here either one year intervals or five year intervals, and whether it is important to estimate it conditional on the current state of the world. For both intervals the research shows it is important to include the current state of the world. Bangia et al. (2002) shows a significant difference between PDs during bad years and during good years and discusses the importance of the underlying macroeconomic volatility. In my paper I will look at a possible solution by using a (pooled-)AR model to estimate the current PDs and thus include a use of the current state of the world. Nickell et al. (2000) shows not only the importance of current state of the world, but also other properties such as what sector and whether its underlying company is US or non-US. The CreditRisk+ and PGGM model could both be extended to also split the different loans between sectors as long as there is enough historical data available for these sectors. The A-IRB can be more challenging as it needs PDs available per sector to obtain a similar option. While this can be done, most standard ratings published are not provided in separate sectors. The mixed model thus has two options to split the model in sectors, of which the use of historical PDs would be the one more often available when choosing less commonly used sectors.

3 Theory

Credit risk models consist of different sub-models focusing on different properties of credit risk. Examples of this are PD s, a rate of recovery model or portfolio sensitivity model. These types of models can also be expanded by looking at individual loans or looking at groups of loans based on properties such as rating or sector. This paper focusses on modelling the PDs, and assumes the recovery rate to be a constant of 50%. The credit risk model, these sub-models combined, provides a loss a portfolio of loans can have due to the parties which lent the money no longer being able to pay its debt. In this paper we look at an insurance which covers the first tranche of defaults from a set of loans, which means a maximum loss of 10% of the total value of loans within the portfolio. As a result, the detach point, which describes the limit used to specify the trench, will be 90%.

To obtain the PDs for the entire portfolio a PD needs to be used for the individual loans before including correlation. PDs for individual loans can be obtained through multiple sources, of which the most established ones are Moody, S&P and Fitch, where Moody is used in this paper. The two main forms of PDs are used in this paper the current state PDs and the historically observed PDs. An example of the current state PDs are the Idealised PDs, which are the probabilities of a default over a long period scaled back to one year. These Idealised PDs are often provided by rating agencies, such as S&P, Moody's or Fitch.

3.1 A-IRB Model

The A-IRB model is based on Vasicek (1987) and uses a factor-based correlation structure,

$$x_i = a_i \cdot F + \sqrt{1 - a_i^2} \cdot Z_i, \quad (1)$$

where F is the common factor used to represent the PD all loans in the portfolio have, Z_i is a factor describing the unique features of the loan i and a_i is a weight between -1 and $+1$ for loan i . This leads to a correlation between loans x_i and x_j being equal to $a_i \cdot a_j$. Vasicek used this to create a Credit Value at Risk (VaR) where the factor F is expressed as $N^{-1}(1 - X) = -N^{-1}(X)$ where X is the percentile of the considered at VaR. Obtaining the PD unique to the individual loan, noted as $V(X,T)$, is obtained by,

$$V(X,T) = N \left(\frac{N^{-1}(Q(T)) + \sqrt{\rho} \cdot N^{-1}(X)}{\sqrt{1 - \rho}} \right), \quad (2)$$

here T is the period of time, $Q(T)$ the factor unique to the loan, $-N^{-1}(X)$ the common factor and ρ the correlation of this loan. The A-IRB model for multiple loans sums these distributions for the individual loans. Each loan requires $Q(T)$ and ρ to be obtained. To do this the loans get sorted into different categories of different PDs for these loans. These different categories will be linked to Moody categories (Aaa to Caa-C) and use the Moody idealised PDs to acquire the relevant $Q(T)$. The term ρ_i will be acquired by the use of,

$$\rho_i = 12\% \cdot \left(1 + e^{-\tau \cdot Q_i(T)}\right), \quad (3)$$

where i denotes the Moody categories and τ a factor chosen to fit the PDs, here based on the Basel committee which sets it at 50. A smaller PD, thus a more safe loan, obtains a higher correlation ρ . The reason for this is that the more risky loans tend to go default in both good and bad periods, though more likely in bad periods. The more secure loans rarely go default, and tend only run this risk during bad periods. As a result, while both risky and secure loans have correlation within their PD, this correlation tends to be higher for more secure loans. This should not be confused with the PD itself, which is higher for more risky loans by its definition.

This is not an exact deduction of the original formula presented by the Basel committee,

$$\rho_i = 12\% \cdot \left(\frac{1 - e^{-\tau \cdot Q_i(T)}}{1 - e^{-\tau}}\right) + 24\% \cdot \left(1 - \frac{1 - e^{-\tau \cdot Q_i(T)}}{1 - e^{-\tau}}\right), \quad (4)$$

but for a tau of 50, the term $e^{-\tau}$ contains a value around 10^{-22} , which result in the term $1 - e^{-\tau}$ being equal to one in most programming languages. If one replaces the term $1 - e^{-\tau}$ with 1, equation 4 can be reduced to equation 3.

A distribution for the loss of a loan can be expressed as $V(X, T) * (1 - R) * L$ where L is the size of the loan and R the recovery rate. The internal rate of return (IRR) is calculated as:

$$IRR = \left(\frac{\max(0, B - D)}{B}\right)^{1/Y} - 1, \quad (5)$$

where $B = (1 - \gamma) \cdot \sum L$, with γ (detach point) being 90% as noted earlier and L the maximum loss of a loan, D is a sum of the earlier expressed losses and Y is the number of years the product covers. As this, combined with equation 2, treats the potential losses per loan as a continuous distribution, instead of a discrete one, the probability of having no loss is zero.

3.2 Simulated A-IRB Model

The A-IRB model has the downside that there is no probability of having no loss at all. While for big portfolios with average ratings this can be reasonable to assume, smaller portfolios or portfolios containing only secure loans will have a reasonable probability to have no defaults. This makes the A-IRB model rather unfitting for these portfolios. A solution to this is using the Simulated A-IRB model, which does allow no defaults to take place over a period. The Simulated A-IRB model uses equation 1 simulated through random draw and calculate whether the draw means the loan in question goes default. This results in the following method,

$$\rho_i = 12\% \cdot \left(1 + e^{-50 \cdot Q_i(T)}\right), \quad (6)$$

$$x_{i,j} = \sqrt{\rho_i} \cdot F_j + \sqrt{1 - \rho_i} \cdot Z_{i,j}, \quad (7)$$

$$DF_{i,j} = \begin{cases} 1 & \text{if } x_{i,j} < \phi^{-1}(Q_i(T)) \\ 0 & \text{if } x_{i,j} \geq \phi^{-1}(Q_i(T)) \end{cases}, \quad (8)$$

where i is the loan, j is the simulation, $Q_i(T)$ the PD of loan i , F_j the common factor across simulation j and $Z_{i,j}$ the factor unique for loan i at simulation j . The obtained $DF_{i,j}$ states whether loan i went default in simulation j during the period of time the loan was covered, and thus has a granularity fitting to the portfolio. Further granularity results in outcomes which converge to the basic A-IRB model.

Another adjustment needs to be made for the used PDs. The Moody's long term PDs can properly represent a portfolio for the granularity of the basic A-IRB model. However, since this simulation no longer splits these loans, PDs should be used that are appropriate for the granularity of this simulated A-IRB model. An alternative to the Moody long term PDs is using the Moody's historical PDs to measure the appropriate PDs for the simulated A-IRB. The historical PDs are provided over a one, as well as multiple, year period and can thus be used without having to scale it to the appropriate number of years. To scale the long term PDs to the appropriate periods used for the basic A-IRB model, the following equation is used,

$$e^{-\lambda \cdot y} = 1 - PD_y, \quad (9)$$

where y is the number of years, PD_y is the probability of going default over y years and λ is a factor. Then the PD for y years is:

$$PD_y = 1 - e^{LN(1 - PD_Y) \cdot y/Y} = 1 - (1 - PD_Y)^{y/Y}, \quad (10)$$

where Y is the number of years used in the long term PDs and y is the number of years of the loans looked at.

3.3 PGGM Model

The PGGM model focusses on the historical defaults more directly compared to the previously used idealised PDs, through looking at the PDs of the past 30 years for the different ratings, where each year is called a cohort. This will have an advantage of creating a correlation, as the PDs of different ratings rise and fall simultaneously. The downside of this is the required assumption that historical PDs will represent the current state of the world. The PGGM model uses the following step:

1. Set the number of simulations that will be made per cohort.
2. Sort the different cohorts from worst cohort to best cohort.
3. Use a random draw from a uniform distribution for each loan for all the simulations in the different cohorts.

4. For each simulation, set the recovery rate that is used now. (Since the cohorts got sorted from worst to best in step 2, the first X% of recovery rate will be used in the first X% of the cohorts. Since the recovery rate is often picked worst to best, bad cohorts will be linked to a bad recovery rate and vice versa)
5. For each loan, check the relevant historical PD (so not the same PD as used in the A-IRB model) and if the relevant draw taken in step 3 is smaller than this historical PD, the relevant loan goes into default within the period looked at.
6. For each simulation calculate the Internal rate of return (IRR) in the following way,

$$IRR = \left(\frac{\max(0, B - D)}{B} \right)^{1/Y} - 1 \quad (11)$$

Where $B = (1 - detachpoint) \cdot \sum Loans$, with the detach point being 90% as noted earlier, D is the sum of all the losses obtained in this simulation, which contains a sum of the loans that went default corrected by its recovery rate, and Y is the number of years the product covers.

3.4 Mixed Model

While simulating the A-IRB model would solve the difference it has with the PGGM model on its granularity, the other difference of the PGGM model of including the historical pattern, that separates bad and good periods, is not yet included. We solved this by performing A-IRB model simulations per category of PDs and combine the different outcomes.

This type of mixed model has one drawback, the fact that correlation is now included in two separate ways, through using historical data as well as using equation 3. A solution is to lower the weight of these correlation techniques, which can easily be done in equation 3 by downscaling the function with a constant, creating the following equation,

$$\rho_i = 12\% \cdot \left(1 + e^{-50 \cdot Q_i(T)} \right) + C, \quad (12)$$

where C can be a value between 76% and -12%. This limited value on C is due to the fact that the correlation cannot exceed the value of one, nor is it realistic for the correlation to be negative. In case C is set on a minimum value of -12%, the correlation ρ_i will span between 0% and 12%, depending on what category the loan is part of. This results in taking over the property of the A-IRB model, a more heavy correlation for the “safer” loans, while keeping the base correlation properties of the PGGM model.

3.5 PD Estimation

The downside of the mixed model is that it does not take the current state of the world into account. An alternative way to including historical PDs is by estimating the idealised PDs with historical PDs. One candidate for this is using the historical PDs in an AR(k) model to estimate the current PDs per

category. Sadly the available data is only restricted to 30 years. Adding the fact that making an AR(k) model should be done for PDs over a longer period of time, as one-year PDs tend to be very spiked, and that the available data points for an AR(1) model would thus be less than 29, the estimate made can be rather inaccurate. This issue could be solved by replacing static idealised PDs with the estimates from autoregressive models. We use univariate AR(k) models for each rating category as well as a pooled approach. For a pooled-AR(k) model two models will be made, one for Aaa up to and including Baa2 and one for Baa3 up to and including B3. The reason for splitting this into two groups is due to a rather big difference between the more secure ratings and the more risky ones, which leads to a split close to the common distinction used in finance between High Yield and Investment Grade. The categories Caa-C will be excluded due to the data already being provided as one category as well as having 100% probability of default within the used time frame. The (pooled-) AR(x)-model will be estimated for the term λ as seen in equation 9, and afterwards converted back to the PD for the relevant number of years as seen in equation 10. With the available data, an AR(1)-model gave proper estimates while a higher order started to get inaccurate estimates due to the small amount of data available. The model is shown as follows,

$$\lambda_{i,t} = \beta_0 + \beta \cdot \lambda_{i,t-1} + \epsilon_{i,t}, \quad (13)$$

where i are the different categories within a pooled, only one category in case of a base AR(k) model, t the year looked at and k the number of AR components used. The estimated parameters can be used to predict the next λ and through that estimate the PDs that can be used to replace the idealised PDs.

3.6 Simplified CreditRisk+

The CreditRisk+ model has similarities with the PGGM model, especially when looking at the PDs calculated for a portfolio containing loans of a single rating. This as the CreditRisk+ model treats the PD of a single loan as a Poisson distribution, which converges to a binomial distribution used by the PGGM model for small PDs. The two main differences of the models are the way correlation is treated and how the number of defaults is determined.

The CreditRisk+ model includes the correlation through the volatility of the PDs. Here a difference is made between the CreditRisk+ model made by Credit Suisse, and the one implemented here resulting in a simplified form. For comparison between the PGGM model and the Simplified CreditRisk+ model a similar correlation is assumed within the historical data used in the PGGM model as in the simplified CreditRisk+ model. Since the CreditRisk+ model is more open for different forms of correlation, an example of this is including sector correlation, multiple forms of correlation can be assumed. To make a proper comparison the correlation used here in the Simplified CreditRisk+ model is similar to the PGGM model its historical data, and because of this a historical simulation is used for the Simplified CreditRisk+ model.

The second difference between the PGGM model and the (Simplified) CreditRisk+ model is the way the number of defaults is determined. (Simplified) CreditRisk+ model assumes a very small PD, thus the number of defaults within a portfolio can be drawn from Poisson distribution. As a result, the (Simplified) CreditRisk+ portfolio no longer looks at the individual loans. The properties of these specific defaults can be randomly drawn from the relevant rating of the portfolio looked at. The decision to run this simulation per rating is another simplification to the CreditRisk+ model, as this model uses the property of a Poisson distribution where summing individual Poisson distributions can be treated as a new Poisson distribution as long as the individual distributions are independent.

These two differences result in a Simplified CreditRisk+ model:

1. Set the number of simulations that will be made per cohort.
2. Sort the different cohorts from worst cohort to best cohort.
3. Use the historical PDs to estimate the λ , the parameter of the relevant Poisson distribution, used for the simulations in the following manner: $\lambda_{Rating,Y} = n_{Rating} \cdot PD_{Rating,Y}$, where n is the number of loans within the looked at cohort at the different ratings.
4. Use Knuth's algorithm, as published in Knuth (1969), to simulate the number of PDs per cohort per rating. Due to numbers very close to zero, use the logarithm in the simulation to mitigate the rounding errors.
5. Per rating, randomly draw the just simulated number of loans that would go default. If each loan within a rating is of similar size this draw is not required.
6. For each simulation, calculate the IRR in a way similar to equation 11.

Due to the relatively small PDs combined with a portfolio of 1,000 loans or less, the Knuth algorithm can still be properly used. If the parameter λ gets above roughly 100, a different simulator needs to be used to prevent time consuming simulations.

3.7 CreditRisk+

The Simplified CreditRisk+ model is a way of including the correlation, but the CreditRisk+ model itself also has an extension on its basic model to include correlation. Its initial step is to expand the basic models method, which used direct Poisson distribution draws, into a Negative Binomial Distribution with its parameters μ and σ known. These parameters can here be obtained by looking at the historical data of the different PDs and use its mean and variance. This results into a distribution to obtain the number of defaults.

The next step required is to determine which loans went default, as different loans have different properties such as size and rating. To determine this, the different Poisson distributions will contain a common random parameter which is obtained through a Gamma distribution. The Gamma distribution

is used to create a common factor, thus creating correlation. The parameters required are α and β , who will both be obtained through the average and standard deviation of the historical PDs.

While this model can be calculated directly, this can cause a rather high demand of calculation power for large portfolios. An alternative method to run this can be seen in Melchiori (2004). This method uses the Fast Fourier Transformation (FFT) and an Inverse FFT (IFFT), which is an algorithm used to speed up the Discrete Fourier Transformation (DFT) and the Inverse Discrete Fourier Transformation (IDFT). This method can be done through the following steps.

1. Obtain μ_k and σ_k where k is an individual loan. This is obtained by using the Moody's historical PDs for the different ratings and determining the mean and standard deviation over the available data.
2. Obtain $s = \sum \sigma_k / \sum \mu_k$, $\alpha = 1/s^2$ and $\beta = s^2/1$.
3. Obtain a vector f which individual components are $f(k) = \mu_k / \sum \mu$.
4. Calculate vector g through,

$$g_k = DFT(f_k) = \sum_{j=1}^N g_j \cdot e^{-i \cdot 2\pi \cdot k \cdot j / N}, \quad \text{for } k = 1, 2, \dots, N \quad (14)$$

where N is the total number of loans and i indicating an imaginary unit.

5. Include the correlation by adjusting g as $\hat{g} = (1 - (\sum \mu_j \cdot \beta) \cdot (g - 1))^{-\alpha}$, where α and β are the earlier determined parameters constant across the different ratings.
6. Obtain the cumulative loss distribution through IDFT,

$$f_k = IDFT(\hat{g}_k) = \frac{1}{N} \sum_{j=1}^N \hat{g}_j \cdot e^{i \cdot 2\pi \cdot k \cdot j / N}, \quad \text{for } k = 1, 2, \dots, N \quad (15)$$

where again N is the total number of loans and i indicating an imaginary unit.

The number of percentiles obtained depends on the number of loans within the portfolio. The DFT and IDFT can both be calculated for each loan by computing the equations 14 and 15 directly, but this would take quite a bit of computational power. Since the equation needs to be calculated for each loan, and each equation contains the number of components equal to the number of loans, the required computational power would be $O(N^2)$. By using the FFT and IFFT this computational power will be reduced to $O(N \cdot \log_2[N])$. This improvement to the DFT can be seen in Cooley and Tukey (1965).

3.8 Model Validation

When looking at different Credit Risk models it is currently impossible to determine one model as the best. This has multiple reasons, such as the many different assumptions made, different instruments that have credit risk and whether the different instruments can be represented by the available data.

This causes different banks and financial institutions to use a model applicable to their portfolio. An alternative is to validate the different models and show that they can be used in the type of portfolio looked at in this paper.

We follow Medema et al. (2009) and use a three-step approach to model validation:

- Theoretical validation (discussed here)
- Data validation (discussed in "Data")
- Statistical validation (discussed in "Empirical Results")

The theoretical validation focusses on what assumptions are made in the different models and whether the used theoretical background is adequate. One assumption present here is the use of historical data to estimate the current state of the world for the PGGM-model, Mixed-model and the CreditRisk+. The CreditRisk+ model does not use the historical data directly, as the long-term PDs or idealised PDs are used, but is used indirectly to estimate the volatility. The other two model make direct use of historical PDs to estimate the losses of a portfolio. The (simulated-) A-IRB model and the CreditRisk+ model assume current PDs represents the possible situations well. Both assumptions are not unrealistic, though when using the historical data the properties used by a rating agency should not change mid sample. An assumption used by all models here is that the state of a loan prior to its current state is of no importance other than what is already included in the current rating. A practical problem here could be when looking at loans that are already within a rating for a long period of time being treated similar to loans which had a volatile past period and just entered this rating bucket. While the expected values of PDs might be the same, higher moments such as the volatility of PDs are quite likely to be influenced by this. This issue will be less of a problem for bigger portfolios, but for smaller portfolios this needs to be taken into account. To counter this problem a bank or financial institute with a smaller portfolio needs to not only look at the current state of the companies which the loan is from, but also how it performed the past few years and whether the PDs need to be altered for this specific portfolio.

The data validation will look at the used data and whether it represents the portfolios properly or not. The statistical validation tests how the different models fit compared to each other, and does this through analysing the discrimination. We use Brier score in-sample testing. Brier score, BS , can be obtained as:

$$BS = \frac{1}{N} \sum_{i=1}^N (f_t - o_t)^2, \quad (16)$$

where N is the number of years tested, f_t is the probability which was estimated and o_t the probability of default observed that year. BS can fall between zero and one, where the closer the BS gets to zero, the more accurate the prediction is.

in %	Aaa	Aa1	Aa2	Aa3	A1	A2	A3	Baa1	
PD	0.01	0.1	0.2	0.4	0.7	1.2	1.8	2.6	
in %	Baa2	Baa3	Ba1	Ba2	Ba3	B1	B2	B3	Caa-C
PD	3.6	6.1	9.4	13.5	17.7	22.2	27.2	34.9	47.7

Table 1: Moody’s idealised PDs

4 Data

Loans are linked to a PD by grouping the loans into different rating categories, from Aaa to C-Caa. We obtain the PD data from Moody’s. This is provided in two ways, the idealised PDs and the historical PDs. Table 1 provides the Moody’s idealised PDs over a 1 year period which span between 0.01% and 47.7%. As can be seen, the categories do not cover an equal width of PDs, but instead cover bigger width the more risky the rating category is. The reason the categories got separated in such a way is that all categories have a specific percentage of available loans per category. In case a category get over- or under-represented, the PDs will be refined.

Besides idealised PDs, there are also historical PDs available. The historical PDs of Moody’s are available between the years 1983 up to 2012, thus having 30 years of data available. A historical PD is obtained by looking at the different ratings each year, and looking at the percentage of loans that went default in the following years. This means a loan that was down-ranked before going default still contributes to historical PDs in the category it was in prior when looking at the PDs over longer periods of time. The probability to go default in the first year is, on average, significantly smaller when compared to the second year for more secure loans. This shows the importance of using historical PDs on short term loans, as the idealised PDs tend to overestimate the PD for these more secure loans on the short term.

Figure 1 shows the probability of going default within one year per cohort. The categories A2, Ba2 and B2 are depicted. It can be seen that the historical PDs seem to indeed be related to good and bad periods. For example, between 2004 and 2008 very few defaults were present in these categories, while during the early 2000s recession and the 2008 Global Financial Crisis the number of defaults went up significantly. The most severe year however, is found around the 1990s, which also paired a global crisis that was present late 1980s and early 1990s.

Figure 2 shows the more long-term historical PDs. There are two big differences when compared to the short-term PDs in figure 1, a less spiked and more cyclical PD as well as more severe PDs for safer categories in the long term. The reason that the long term PDs are more cyclical and less spiked is the fact that severe years will now have influence over multiple years, four here. The reason why long-term PDs have more severe PDs at the safer categories is the fact that it becomes much more difficult to predict a default far in advance. Often a company in such a situation gets struck by some bad news, which is then followed by moving the company to a lower category. The company can generally survive




Figure 1: 1 year PDs per cohort as provided by Moody's rating over the period 1983-2011. Safer ratings, here A2, are expected to have lower PDs than less secure ratings, here B2. Also, across ratings a correlation is expected, as both safe and unsafe loans are expected to have a higher PD during bad years.




Figure 2: 4 year PDs per cohort as provided by Moody's rating over the period 1983-2008. Safer ratings, here A2, are expected to have lower PDs than less secure ratings, here B2. Also, across ratings a correlation is expected, as both safe and unsafe loans are expected to have a higher PD during bad years.

a period after the sudden bad news by selling parts of the company or arranging deals with counter parties. Sadly, this often postpones the bankruptcy, and through this results in more bankruptcy taking place in the long-term when looking at safer loans.

The above properties of more long-term PDs make it easier to estimate the long-term PDs based on the one of previous years, something similar to a property discussed by Cantor and Falkenstein (2001) where they discuss the properties of PDs and one of the main conclusions is that long term PDs are better to estimate properly when compared to short term.

One question which might arise is the choice of Aaa-Baa2 and Baa3-C, discussed earlier in "Theory". Often PDs are split into two groups, called investment grade and high yield, which are respectively Aaa-Baa3 and Ba1-C. The reason for this choice can be seen when looking at the historical PDs for the different ratings. In the available data, Baa3 was the first rating to truly contain big spikes of a PD, reaching a chance of close to 10% to go default within four years. This is the first big leap up in PD, as the highest four year PD in Baa2 was less than 5%. The possible difference in choice might be caused by the difference in periods looked at. The split between investment grade and high yield is made with a much longer period in mind than four years. Since the PD looked at here is often used for products that cover two to five years, a choice based on historical four year PDs seems more appropriate.

4.1 Data validation

When looking at the portfolios used to test the different models, some theoretical portfolios are used to determine step by step what the different properties of a realistic portfolio would cause. The theoretical portfolios are chosen in such a way that it slowly turns into a more realistic portfolio by including properties such as including different ratings in one portfolio, looking at a four year period versus a one year period and looking at different sizes of portfolios. The last portfolio is a realistic one, based on a portfolio provided by PGM. This inclusion of testing the models with this internal portfolio is in line with the article of Basel Validation Group (2005), as it stresses the importance of validating the models by testing it on past, realistic portfolios. When looking at the used PDs, the data source is Moody's, which is one of the best representatives in this field. Over the available period the PDs were registered properly and there are no known errors in the data.

5 Empirical Results

The simulations are run over different portfolios, starting at a very basic portfolio in which there is little variation in the PD from individual loans and the two basic techniques are used. This will then be expanded over different, more complex, portfolios as well as the proposed expansions of the basic models.




Figure 3: CDF of 1 year Ba2 portfolio

5.1 1 year Ba2 portfolio

Table 2 shows some of the base properties of the A-IRB model and the PGGM model when looking at a portfolio of Ba2 loans, where the PGGM model has some variations depending on the number of loans it covers. The A-IRB model does not have variation in the number of loans it covers. The reason for this is because the A-IRB model makes use of an analytical approach and thus its granularity is independent of the number of loans, given the loans have the same probability of default. This difference can be seen when comparing the percentiles in table 2 between the A-IRB model and the number of loans used in the PGGM model. It shows a small number of loans in the PGGM model results in a bigger chance of having no loss at all, while simultaneously having a bigger chance of having extreme losses.

When comparing the results of the A-IRB model with the PGGM ($n=1000$) model the one can see that the expected IRR is much lower for the PGGM model. Figure 3 show that there is a huge chance of having no loss at all in the PGGM model, which can explain why the expected IRR is much lower for the PGGM model. There are two main causes for the high probability to have no loss at all for the PGGM model. In the previous paragraph it is explained that the A-IRB model is not sensitive to a change in sample size due to not treating a loan as a single loan anyway. This difference has the direct result that the chance of having no default at all is infinitely small, thus rounded off to zero given that the PD is unequal to zero.

in %	A-IRB	PGGM (n=10)	PGGM (n=100)	PGGM (n=1000)	Simpl. CR+ (n=1000)	CreditRisk+
Expected IRR	-6.5	-3.1	-3.0	-3.0	-3.0	-9.3
Exp. Shortfall 95%	-17.2	-54.0	-19.9	-19.2	-19.1	-24.8
Taildepth 95%	-8.0	-46.9	-12.0	-11.5	-11.5	-11.4
0.01 th percentile	-24.5	-100.0	-55.0	-32.5	-32.5	-40.9
5 th percentile	-14.5	-50.0	-15.0	-14.5	-14.5	-20.7
20 th percentile	-9.7	0.0	-5.0	-6.0	-6.0	-13.6
50 th percentile	-6.0	0.0	0.0	0.0	0.0	-8.1

Table 2: Ba2 portfolio for a 1 year period. The PGGM model will change depending on the number of loans in the portfolio, where a bigger number of loans converges to a more continuous distribution. Both the CreditRisk+ model and the A-IRB model provide a continuous distribution as outcome, while the simplified CreditRisk+ model and PGGM model are discrete.

The second reason why the PGGM model contains a high probability to make no loss is related to the use of historical PDs. Historical PDs vary over time and are based only on how many loans in the relevant categories have gone default. However, during economically good periods it is not uncommon for none of the loans to go default at all, especially in the more secure categories. As a result, simulations from that year will automatically result in 0% IRR. This result is more prominent when using loans from a single category, here Ba2, as a year only needs a 0% PD for Ba2 rating. When spreading a portfolio across different ratings a 0% IRR becomes less likely to occur, as it is very unlikely to have 0% PD for multiple rating in the same year. However, even though it is less likely, a spread portfolio still has a probability for no default as the draws will always have a chance to draw no defaults when using a discrete distribution. In contrast, the A-IRB model simulates from one set of PDs, of which all categories have a certain probability bigger than 0% to go default.

The Simplified CreditRisk+ model has quite a few properties similar to the PGGM model. First of all, fewer loans in the portfolio results into a larger chance of either having an IRR of zero, or having an IRR of one. This due to similarities discussed in "Theory". Secondly, the distribution given a portfolio is almost identical when comparing the PGGM model and the Simplified CreditRisk+ model. With a small difference from the fourth percentile onwards, the PGGM model and the Simplified CreditRisk+ model are similar in shape of the distribution. This is to be expected, given this type of portfolio and period. Under these conditions, the PGGM model results into a simulation from a binomial distribution per cohort, while the Simplified CreditRisk+ model has that of a Poisson distribution. Since the portfolio we are looking at has a sample size of 1,000 and the PD for Ba2 rating is either 0%, or very close, the Poisson distribution is going to have a mean equal to that of a binomial distribution used by the PGGM model, thus also equal expected IRR to the PGGM model, and higher moments in this case have a negligible effect. For example, considering the highest one year PD for Ba2 is 4.19% in the year 1998, the kurtosis of the binomial distribution will be 0.019 while that of the Poisson distribution is 0.024. Keeping in mind that this is the worst year, and that more than half the number of years has a PD of zero, resulting in equal simulations for both techniques, one can expect both distributions to be very similar for this portfolio. As a result, the differences between the A-IRB model and the Simplified CreditRisk+ can be explained in a similar way as discussed earlier for the PGGM model and the A-IRB model.

The CreditRisk+ model its 50th and 20th have bigger losses compared to the other models, of which the A-IRB model comes closest to the CreditRisk+ model. The reason why is that similar to the A-IRB model, the CreditRisk+ model does not have a probability of no loss and thus always treats a portfolio as a highly granular portfolio. At higher percentiles the CreditRisk+ model keeps a relatively high loss, while the A-IRB estimates relatively small losses compared to the other models. The reason for this is the stronger presence of correlation in the CreditRis+ model, as seen in the lower percentiles as well, combined with the Gamma distribution used by the CreditRisk+ model for its correlation component being dependant on the historical PDs. If the historical data contained better and more stable years, the CreditRisk+ model would come closer to the A-IRB model.

5.2 4 year Ba2 portfolio

Table 3 shows the previous models for estimating IRR over loans covered for four years instead of one, all other properties of the loans are not changed. Compared to table 2, one can observe that the probability of going default has shifted to the right, as expected, but that especially the probability of no defaults for the PGGM model has decreased tremendously. While covering a portfolio of a 1000 one year loans, the probability of not having any defaults was 58.1%. In contrast, this probability for four year loans is 0.12%. The main reason for this outcome is that the historical PDs no longer contain years where the chance to go default is equal to 0%, as the minimum PD over four years for Ba2 is equal to 0.85%.

Another change when comparing the two tables is the expected IRR for the PGGM model when only having ten loans. When looking at one year loans the expected IRR stayed within a 0.1% range however, table 3 shows an expected IRR of a 10 loans portfolio having almost doubled that compared to the 100 and 1000 loans ones. The difference of years covered is the direct cause of this. While the expected losses over multiple years should, on average, stay the same, the volatility of the losses should increase when the number of loans is decreased. This evens out in a linear function, but the expected IRR goes to the power $1/(Nr.of\ years)$ when calculating the IRR across multiple years. In contrast, this stays linear when the number of years is equal to one, but as table 3 uses a four year period, this no longer holds for this longer period. This results in no longer observing the same expected IRR being equal for the different number of loans.

While increasing the number of years decreased the dominance of the PDs equal to zero, the variation in historical PDs over multiple years is still influenced heavily by the difference in bad and good years. Figure 4 shows the Cumulative Distribution Function (CDF) of both the PGGM model and the A-IRB. In the figure, one can see that the A-IRB model only exists of one steep increase, which is expected as the A-IRB uses one estimate of the PDs. The PGGM model contains two separate steep ascending areas. Here the first steep area, between -40% to -20%, can be linked to bad economic periods of time while the second steep area, between -10% to -2%, mimics good economic periods.

The mixed model includes two forms of correlation, and thus without a correction for this the model would contain a way too high correlation. This can be seen in table 3, where the mixed model without any correction has a big chance of having an extreme loss. This is paired with a bigger chance to have a very low loss over this period. The mixed model has a 12% chance to have a loss of less than 1%, while the PGGM model only has the chance of 2% to have such a small loss, as can be seen in the percentiles in table 3.

Even with this larger chance of small losses present for the mixed model, the expected IRR is still a bigger loss for the mixed model, 15.3%, than the PGGM model, 11.2%. As seen in the 4-year tests versus the 1-year tests, this increase of years causes the expected IRR to no longer be a linear function, and thus letting higher moments influence the expected IRR. When the added correlation from the A-IRB




Figure 4: CDF of 4 year Ba2 portfolio

in %	A-IRB	Sim. A-IRB (n=100)	Sim. A-IRB (n=1000)	PGGM (n=10)	PGGM (n=100)	PGGM (n=1000)	Simpl. CR+ (n=1000)	CreditRisk+	Mixed (C=0%)	Mixed (C=-6%)	Mixed (C=-12%)
Expected IRR	-13.0	-14.1	-13.6	-20.3	-12.9	-11.2	-11.6	-14.1	-15.3	-14.5	-11.8
Exp. Shortfall 95%	-78.3	-99.9	-80.5	-100.0	-71.7	36.0	-37.1	-84.0	-100.0	-99.5	-41.9
Taildepth 95%	-24.5	-38.6	-26.9	-79.7	-30.9	-21.5	-21.5	-29.9	-84.7	-58.9	-23.1
0.01 th percentile	-100.0	-100.0	-100.0	-100.0	-100.0	-65.0	-100.0	-100.0	-100.0	-100.0	-100.0
5 th percentile	-37.5	-52.7	-40.5	-100.0	-43.8	-32.7	-33.1	-44.0	-100.0	-73.4	-34.9
20 th percentile	-16.9	-18.1	-17.9	-15.9	-20.5	-24.2	-24.5	-18.3	-19.5	-19.7	-23.6
50 th percentile	-8.1	-8.5	-8.5	0.0	-6.9	-7.1	-7.4	-8.4	-5.9	-6.6	-7.4

Table 3: Ba2 portfolio for a 4 years period. The mixed model uses a portfolio with 1000 loans (n=1000). C is a correction used to reduce the correlation obtained through the A-IRB part in this model, which is used to not completely retain both the PGGM model and A-IRB model correlation. The PGGM model with a portfolio of 10 loans is too small for most realistic portfolios, but stresses the importance of size a portfolio contains.

	PGGM model	Simpl. CreditRisk+
Mean	159	159
Variance	133.7	159
Skewness	0.058	0.080
Kurtosis	0.001	0.006

Table 4: PD distribution properties of the PGGM model, which has a binomial distribution, and the Simplified CreditRisk+ model, which has a Poisson distribution

model is decreased, all the measures in table 3 seem to converge for the mixed model to the PGGM model. With a portfolio existing of only a single category of loans, here Ba2, the mixed model can even become the exact same as the PGGM model as long as C is set to be equal to $-12\% * (1 + e^{(-50 * Q_i(T))})$.

When looking at the Simplified CreditRisk+ model over a longer period, such as four years here, one expects a larger difference when compared to the PGGM. Not only does the IRR now contain a power factor, as discussed in the earlier comparisons, the previously observed similarity between the Poisson distribution and the binomial distribution is less present when the PD increases. The most severe PD in this portfolio, under a four year period, is 15.9%, much higher when compared to the maximum of 4.2% when looking at the one year period PDs. Similarly, none of the PDs in the different cohorts are equal to 0% when looking at the 4 year PDs, which explains why a probability of no defaults present is also not significantly off of 0%.

The biggest difference between the Simplified CreditRisk+ model and the PGGM model lies within the tail. At the 0th percentile, one can see that the lowest IRR in the simulation for the PGGM model is 65.0%, while this is 100.0% for the Simplified CreditRisk+ model. It is important however, that while these extreme situations are important for stress testing portfolios and looking at tail risk in general, they are also much less accurate as these percentiles tend to fit less properly with the distribution. When looking at the number of defaults for the IRR of -100%, one can see more than 200 defaults present for multiple draws and peaking at 220 defaults, while the maximum number of defaults within the simulated PGGM model is 197. Due to the detach point being 90%, the difference between the most severe losses in both simulations is only 35%, which would have been bigger in case a higher detach point was taken. Similar to the one year portfolio, the cause here is the difference within the higher moments between the two models. Looking again at the worst year in the historical data, here the year 1983, leading to a 15.9% expected loss, the properties of the relevant distributions can be seen in table 4.

Again, keeping in mind that this is the worst year in mind, the higher variance, skewness and kurtosis explain the higher outliers. For the other percentiles one sees only a minor difference, which is normal considering the fact that this part of the distribution will be more dependent on the various years during a normal and positive state of the world, and thus expecting a smaller difference between the two distributions.

The CreditRisk+ model its percentiles are much closer to the Simplified CreditRisk+ model. The

cause is similar as earlier explained for the A-IRB model versus the Simulated A-IRB model, where no longer having years with no defaults is present. That said, the CreditRisk+ model still has a bigger difference compared to the Simplified CreditRisk+ model, which can be seen in the expected IRR. Looking at the percentiles the Simplified CreditRisk+ model seems to only have a lower IRR at the 20th percentile. A cause could be a higher correlation in the CreditRisk+ model, and that the high 20th percentile for the Simplified CreditRisk+ model is caused by the present crisis in the historical PDs.

5.3 4 year Baa3 to Ba3 portfolio

Up to now, the used portfolios consisted of loans from a single category, Ba2. In reality, portfolios tend to contain loans spread over different categories, usually between A1 and B3. The next step is to perform the same comparisons with portfolio spread across a small number of different categories while keeping the rest constant, and see how this influences the differences between the different models. This spreading of the portfolio would also provide a better look at the mixed model, as it would allow the A-IRB part of the mixed model to focus on including different levels of correlation for different categories.

Table 5 shows results for a portfolio consisting of loans with rating Baa3 to Ba3, with each 250 equal sized loans. Using the PGGM model as a base, the A-IRB model comes quite close to the PGGM model when specifically looking around the 5th percentile. This can be explained by the fact that the A-IRB model puts emphasize on the tail, while the PGGM model tries to fit best at the whole distribution. This emphasize on the tail from the A-IRB results in the A-IRB model being more off the mark outside the tail, and thus also expected to be further off the PGGM model.

When looking at the simulated A-IRB model one can see that the expected IRR is closer to the PGGM model, but has a big difference in the 95% expected shortfall. The cause for this can be seen from the 0.01th percentile, which shows that the simulated A-IRB model finds a probability of a 100% loss, while the PGGM has a 46.0% loss as 0.01th percentile. The chance a 100% loss occurs is 0.34%, which has a significant influence on the 95% expected shortfall and thus creating this big difference here with the PGGM model.

The A-IRB model again shows a distinct difference when compared to the PGGM model, Mixed model and CreditRisk+ model. Solving this with a simulated A-IRB model seems to help create a more similar IRR distribution, but keeps a significant difference present in the tails.

When looking at the mixed model, the base properties found in the Ba2 portfolio seem to again be present in this portfolio. The most important change found in this portfolio is the mixed model given a correction of -6% on the A-IRB correlation. In the Ba2 portfolio, this resulted in huge differences in the looked at percentiles, as well as double the tail depth in the mixed model with a -6% correction compared to the PGGM model. When looking at the percentiles in this Ba3-Baa3 portfolio, the percentiles are reasonably close to each other until the 20th percentile, after which a significant difference is present.

in %	A-IRB	Sim. A-IRB	PGGM	Simpl. CR+	CreditRisk+	Mixed (C=0%)	Mixed (C=-6%)	Mixed (C=-12%)
Expected IRR	-13.0	-11.2	-10.6	-11.0	-15.5	-14.5	-13.1	-11.0
Exp. Shortfall 95%	-75.0	-37.4	-29.2	-29.5	-77.2	-100.0	-75.6	-31.5
Taildepth 95%	-24.0	-14.2	-15.4	-15.0	-26.0	-85.5	-27.4	-15.6
0.01 th percentile	-100.0	-100.0	-46.0	-51.6	-100.0	-100.0	-100.0	-100.0
5 th percentile	-37.0	-25.4	-26.0	-26.0	-41.5	-100.0	-40.5	-26.6
20 th percentile	-16.9	-15.3	-18.6	-18.8	-20.3	-18.8	-18.3	-18.6
50 th percentile	-8.5	-9.4	-8.4	-8.9	-10.8	-7.1	-8.0	-8.9

Table 5: Baa3 to Ba3 portfolio for 4 years with 250 equal sized loans per rating, thus leading to a portfolio of 1000 loans in all cases. Probability of having no default is between 0% and 0.5%, thus no longer being significantly present, likely caused by a spread across different loans and thus not having no defaults any more when one category has a PD of 0% for a cohort.

The mixed model ($C=-12$) has another interesting property when considering the base value of the A-IRB correlation is removed. Looking at the two tables, the mixed model seems to be closer to the PGGM model when the loans are more spread. This is especially visible for the 95% expected shortfall. In the previous portfolio, an additional correlation of 0.4% through the A-IRB correlation was present. With the new portfolio this correlation is now between 0.016% for the Ba3 category to 3.6% for the Baa3 category. This shows that the A-IRB correlation is less significantly present for the more risky loans, which tend to also be the loans that have more impact on the extreme losses. Since the more risky loans are corrected less by the A-IRB correlation, the distribution will be influenced less in its tails, which tends to be the biggest difference between the PGGM model and mixed model.

When spreading the portfolio for the Simplified CreditRisk+ model one would expect a distribution closer to the PGGM model when compared to a portfolio with an equal average PD but less of a spread across different categories. The reason for this is while there is a correlation present between the different PDs, history does show a difference in worst year for the different categories. For example, the worst year for the rating Ba2 with a four year loan is 1982 in the used data. If one would compare this with the spread portfolio here, the worst year is 1987 while the year 1982 takes an eight spot of the 27 different years. The reason for this is that the Ba3 rating has a most severe year in 1987 for a four year period, and since this rating its PD is much bigger than Ba2, results in this rating being more prominently present in the cause of loss due to defaults.

Similar to previous comparisons, the IRR will show a more significant difference when looking at the last and first few percentiles, again due to the difference in higher moments between the Poisson distribution and a binomial distribution. An interesting note here is the distribution for the Simplified CreditRisk+ model has no zero loss in the simulations ran, while the PGGM model does have a small PD of no loss. This seems to be counter intuitive to the skewness which the Simplified CreditRisk+ model has compared to the PGGM model. While the difference is not significant, a possible explanation is the difference in the pseudo random number generator. The other models can be compared by using the same seed, thus the simulations across these models use similar draws. This is not possible for the Simplified CreditRisk+ model, as it uses a different number of random draws, namely one per draw from the Poisson distribution and one per required number of defaults to pick. The other models look at defaults per individual loan, thus the number of random draws is equal to the size of the portfolio times the number of simulations.

5.4 4 year realistic portfolio

Next a more realistic portfolio is used based on PGGM scenarios, which generally means a portfolio consisting of loans from different categories. The portfolio tested here falls between A3 and B1 and contains 730 loans across these portfolios. Furthermore, loans will now have a difference in value, falling between 20,000 and 25,000,000, instead of an equal value.

in %	A-IRB	Sim. A-IRB	PGGM	Simpl. CR+	CreditRisk+	Mixed (C=0%)	Mixed (C=-6%)	Mixed (C=-12%)
Expected IRR	-7.7	-7.6	-6.6	-6.9	-9.5	-8.0	-7.4	-6.8
Exp. Shortfall 95%	-38.1	-36.1	-17.2	-17.1	-32.7	-48.0	-31.9	-18.4
Taildepth 95%	-8.7	-12.3	-13.0	-8.6	-12.7	-16.5	-13.0	-8.9
0.01 th percentile	-100.0	-100.0	-28.9	-31.2	-100.0	-100.0	-100.0	-42.1
5 th percentile	-20.0	-20.6	-15.3	-15.4	-22.1	-24.5	-20.4	-15.8
20 th percentile	-10.6	-10.7	-10.8	-11.1	-13.3	-11.5	-11.2	-11.0
50 th percentile	-5.6	-5.6	-5.9	-6.2	-7.7	-4.8	-5.4	-6.1

Table 6: A3 to B1 portfolio for 4 years.

Both the A-IRB model and the simulated A-IRB model have a very similar value to that of the PGGM model for the 50th and 20th percentiles, and starts to have a bigger negative value for the 5th and 0.01th percentiles, as can be seen in table 6. Figure 5 shows the CDF for these distributions compared to that of the PGGM model, and indeed shows a slower decline from the 20th percentile onward. The reason a more fat tail is present in the (simulated) A-IRB model is the stronger presence of correlation. This higher correlation should not only lead to a higher probability of big losses, but also a higher probability to have very small losses. This higher probability of small losses however, is not present in the (simulated) A-IRB compared to the PGGM model. The cause of this is the property of historical PDs used by the PGGM model to contain years of (close to) no defaults in the categories of this portfolio. While these historical PDs are also used to estimate the PDs used by the (simulated) A-IRB model, these estimated PDs are set prior and not varied during the simulation nor used in the analytical solution. As a result, neither the A-IRB model or the simulated A-IRB model will obtain a PD of 0%. While having no defaults in a simulation for the simulated A-IRB is possible, a PD of 0 would result in no defaults in that category at all in that specific year and thus having a stronger correlation given these extremely small probabilities of default. This should work for the opposite as well, given a PD of 100% for a specific category at specific years should lead to more extreme negative losses in the PGGM model. However, in the historical PDs, a 100% PD over 4 years is not present, nor is a PD remotely close to that in the categories B1 to A3, where the biggest PD in the historical PDs is 33.52%.

Table 6 also shows the mixed model compared to the PGGM model for this A3-B1 portfolio. The table contains a lot fewer extreme values compared to previous tables involving the mixed model. To explain this it is important to realize where the addition of the category correlation has most influence when looking at defaults. While the added correlation is higher for more safe categories, these safe categories are rarely the most dominant categories when looking at extreme losses, which are instead dominated by years with high PDs for the lower categories. At worst, this extra correlation for secure categories amplifies simulations that are already hit severely in the more risky categories. The reason for this is the fact that different categories tend to move to higher PDs in similar years. As a result, in a lot of years the higher correlation in more secure categories has no influence, as the PD in the many good years is 0% for these secure categories. It is not until the more risky categories have relatively big PDs that the secure categories start contributing to the losses.

When looking at the mixed model with a correlation correction of -6% its statistics are very close to the ones of the simulated A-IRB. In fact, the whole distribution is very similar with only the outliers, the worst 0.36% that a 100% loss occurs in the simulated A-IRB, as an exception. This indicates the difference between the simulated A-IRB model and the PGGM model consisting of an increase of general correlation, as well as further enhancing the correlation between more secure loans.

The simulations of this A1-B3 portfolio shows that for a more diverse portfolio, the simulated A-IRB




Figure 5: CDF of 4 year realistic portfolio

model becomes much closer to the PGGM model. This is caused by less volatility in the historical PDs when looking at more secure categories. It also indicates that the PGGM model assumes less correlation when compared to the simulated A-IRB. For the more risky loans, a difference between the two models seems to be more pronounced. Due to the dependence of the PGGM model on the history being a representation of upcoming economic situations, the question would be whether history indeed fits the current upcoming economic state. If not, one could think of a simulated A-IRB model as a more careful alternative to the PGGM model.

The comparison between the Simplified CreditRisk+ model and the PGGM model is more difficult here, as the different sizes in loans prevent treating the whole loss through defaults as a direct Poisson distribution for Simplified CreditRisk+ and a binomial distribution. The number of defaults should still have the same distributions and as this is an underlying distribution of the IRR, a very similar distribution is still expected with a small difference of the Simplified CreditRisk+ being more severe.

Looking at table 6 this is indeed the case, where even the worst simulation for the Simplified CreditRisk+ model is only 2.3% worse, even less than expected. A possible explanation is the variation in size of the different loans. Given a number of defaults during a bad scenario, it is possible to be worse off when the size of loans varies, but the chance of this occurring is extremely slim. The reason is that worst draws are already filtering the total number of draws by worst years when looking at the worst

	β_0	β_1	λ_i	PD_i
Aaa	1.2e-4 (1.1e-4)	-0.040 (0.205)	1.2e-4	4.6e-4
Aa1	5.8e-4 (5.8e-4)	-0.040 (0.205)	5.8e-4	0.002
Aa2	4.3e-4 (9.5e-4)	0.561 (0.179)	0.002	0.009
Aa3	7.3e-4 (4.0e-4)	0.355 (0.190)	0.001	0.005
A1	0.001 (6.3e-4)	0.537 (0.176)	0.004	0.014
A2	7.7e-4 (3.5e-4)	0.459 (0.179)	0.002	0.006
A3	8.4e-4 (3.9e-4)	0.510 (0.173)	0.002	0.008
Baa1	0.001 (6.6e-4)	0.455 (0.187)	0.004	0.016
Baa2	0.002 (9.3e-4)	0.358 (0.190)	0.003	0.012
Baa3	0.002 (1.4e-3)	0.554 (0.161)	0.004	0.015
Ba1	0.007 (3.2e-3)	0.527 (0.173)	0.014	0.053
Ba2	0.004 (3.0e-3)	0.709 (0.130)	0.008	0.031
Ba3	0.006 (5.2e-3)	0.808 (0.122)	0.018	0.070
B1	0.009 (7.1e-3)	0.779 (0.130)	0.026	0.098
B2	0.013 (9.8e-3)	0.787 (0.130)	0.043	0.159
B3	0.021 (0.016)	0.764 (0.138)	0.051	0.183
Aaa-Baa2	8.3e-4 (1.7e-4)	0.467 (0.055)	X	X
Baa3-B3	0.004 (2.0e-3)	0.885 (0.032)	X	X

Table 7: β s estimates for the (Pooled-)AR(1) model on the PDs (Variance in parentheses) based on Moody’s historical PDs. See equation 13 to see the parameters needed to estimate λ_i where i is the rating. See equations 9 and 10 on how to convert λ_i into a PD_i , where again i is the rating.

percentile. For this simulation, the worst year is 1988, followed closely by 1987, so this worst percentile is already filtering out 93% of the draws. This explains why the number of simulations done here, namely 27,000 draws, will likely not contain a proper draw for the worst percentile. When the research puts emphasis on stress scenarios, it would be better to draw simulation per year separately, to increase the number of draws and obtain a better estimate of the worst percentile.

The CreditRisk+ model has a percentiles that seem to all be shifted to slightly more severe losses, roughly 2%, when compared to the (simulated) A-IRB model. Since the expected IRR is also roughly 2% lower for the CreditRisk+ model compared to the (simulated) A-IRB model, it implies the CreditRisk+ model does not necessarily has a higher correlation, but could also have a more negative look at the portfolio in general. Keeping in mind that the CreditRisk+ model uses the historical PDs in a more indirect way, only using it to calculate the parameters in its Gamma distribution, this can mean the historical PDs have a lower average PD for this portfolio.

5.5 Estimate of the current PDs

The different λ estimated with (pooled-) AR(1)-models and the parameters obtained can be seen in table 7 (variance between parentheses), as well as the relevant PDs. The λ and PDs are not available for the pooled-AR(1) model, as this model estimates multiple λ s and PDs for a single estimate of the two β s. See table 8 for the estimated PDs through a pooled-AR(1) model.

With the exception of Aaa and Aa1, the different categories have a small β_0 that is not significantly different from zero paired with a significant β_1 . Also, the higher PDs tend to get paired with both a higher

in %	Aaa	Aa1	Aa2	Aa3	A1	A2	A3	Baa1	Baa2
Idealised	0.04	0.4	0.8	1.6	2.8	4.7	7.0	10.0	13.6
Historical	0.04	0.2	0.3	0.4	0.8	0.5	0.6	1.0	1.4
AR	0.05	0.2	0.9	0.5	1.4	0.6	0.8	1.6	1.2
Pooled-AR	0.3	0.3	0.9	0.6	1.2	0.6	0.8	1.4	0.7
in %	Baa3	Ba1	Ba2	Ba3	B1	B2	B3		
Idealised	22.3	32.6	44.0	54.0	63.4	71.9	82.0		
Historical	2.4	5.8	7.0	13.3	16.9	23.9	32.3		
AR	1.5	5.3	3.1	7.0	9.8	15.9	18.3		
Pooled-AR	2.4	5.6	3.3	6.7	8.7	14.0	14.0		

Table 8: Moody’s idealised PDs compared with the historical PDs and the two estimated PDs, one through an AR-model and one through two Pooled-AR models, one for Aaa-Baa2 and one for Baa3-B3. Caa-C is separated into a third class with its base PDs found in the historical PDs.

β_0 and β_1 , not just β_1 which might be expected if the β_0 is not significant. The reason why Aaa and Aa1 have an estimate of β_1 that is negative and much closer as well as not significantly different from zero is due to the lack of defaults in the historical data for these two categories. In both cases there was only one case where a default was seen within four years, namely in the starting year 1983. This results in a lack of data to estimate these parameters properly, thus resulting in these unexpected negative estimates.

When looking at the PDs one can see that roughly an ascending PD for the categories can be seen. This is similar to both the idealised and historical PDs. However, compared to both the idealised and historical PDs the estimated PDs seem to be lower, especially for the categories Ba2 up to and including B3. One possible explanation for this is the fact that the estimated PDs take into account the current state of the world through the AR(1)-model. The last year used here for the recorded PDs is 2008, which looks over the period 2008-2011. As discussed earlier in the data, the PDs around the Euro-crisis were much smaller when compared with the 2000 crisis or the 1990 crisis. As a result, the estimated PDs will be smaller as more value is attached to the current state of the world when compared with the idealised and historical PDs. The reason for this is that historical PDs are estimated with equal weights to the past years with available data, and the idealised PDs are estimated with long term PDs in mind, thus attaching less value to the current state of the world.

The last option looked at is the pooled-AR(1) model. The β_0 and β_1 estimates for the pooled-AR(1) model are all significant at 5% level, which cannot be said for the β_0 and β_1 estimates by the AR(1) model. One of the culprits for the AR(1) model can be the lack of available data. This explains why the PDs that are smaller are less significant when compared with bigger PDs, as a lower PD would need more available data before it can be estimated significantly off of the null hypothesis of a not significant parameter. Another problem with the AR(1) model is the estimated β_1 for Aaa and Aa3, as there are not enough registered defaults in the available data for these two categories. By using a pooled-AR(1) model, this lack of defaults no longer offsets the estimated β_1 .

Table 8 shows a more direct comparison between the PDs obtained in the different techniques used. The AR and pooled-AR estimated PDs are ascending at a much more modest speed when compared to

in %	Sim. A-IRB	Sim. A-IRB (AR)	Sim. A-IRB (P.-AR)
Expected IRR	-7.6	-4.6	-4.2
Exp. Shortfall 95%	-36.1	-23.1	-14.4
Tail depth 95%	-13.0	-9.3	-6.7
0.01 th percentile	-100.0	-100.0	-100.0
5 th percentile	-20.6	-13.9	-10.9
20 th percentile	-10.7	-6.8	-6.3
50 th percentile	-5.6	-3.2	-3.5

Table 9: Using the estimated PDs for the Simulated A-IRB model. First the outcome using the historical PDs, followed by the AR estimated PDs and the Pooled-AR estimated PDs.

the idealised and historical PDs. A second observation is that while the estimated PDs are ascending in general for the four different PD techniques, only the idealised PDs are ascending perfectly. This is caused by the fact that idealised PDs are not estimated. The different loans are estimated to fall under one of the relevant PDs instead. The other three PDs are estimated based on historical data, thus vulnerable for error in the estimate. For the AR and pooled-AR model a second cause is the extra weight of the current state of the world. This means the historical PDs of last year influences the estimates that are made. If a safer category performed worse when compared to a more risky category in the most recent used PDs, this can cause the same lack of ascending order.

When looking at the estimated losses based on the PDs estimated with a pooled-AR model, as seen in table 9, one can see a smaller expected IRR and 95% shortfall, as well percentiles closer to zero, with the exception of the 0.01th percentile. This is expected as the estimated PDs were done for a time where the last period looked did not contain a huge number of defaults in comparison to previous years, as was discussed earlier.

What does look surprising is the 0.01th percentile, which means a 100% loss of a 10% coverage is still possible with a reasonable probability, even though the estimate was made over a relatively good state of the world. Looking at the simulation, the probability of this -100% IRR is 0.008%. Looking at the simulations with an IRR of -100%, the number of defaults is on the higher side but not the biggest number of defaults. One explanation is the use of loans with a variation in size of the loans, which could mean the bigger loans were hit simultaneously. The reason why the PGGM model does not have this occur is due to it using a different method to add correlation to the PDs per simulation. The PGGM model assumes the correlation present in the historical PDs is enough to capture it. A possible problem here lies in not yet having enough data available that captures this higher correlation in safer loans during bad years.

The AR(1) model its IRR distribution also has smaller losses compared to the historical PDs, but is closer to compared to the pooled-AR(1) model. Due to a small sample size available for estimating the AR(1) parameters, inaccurate estimates were expected to result in a distribution with more extreme losses for different ratings. Table 8 already showed the AR(1) model to not have extreme estimates, and thus no extreme values in the IRR distribution is no surprise.

in %	PGGM	Sim. A-IRB	CR+	Mixed (C=-6)
Brier in-sample (5y)	7.98	8.31	7.88	8.24

Table 10: This table contains a Brier score, between 0% and 100%, for the four main models used here. The closer the value is to 0%, the better the model performs.

5.6 Statistical validation

The Brier score presented in table 10, which is done for the different models under the realistic portfolio using the historical PDs. The historical PDs were used to make this test comparable between the different models, as using the estimated PDs would make it difficult to see whether the model performs better or if including the estimated PDs gives an advantage. Only the simulated A-IRB model could use the estimated PDs and thus its Brier score cannot be compared to other models without this problem. Looking at table 10, none of the models performs significantly better, which supports the practical approach different banks and financial institutions have of picking a personal preference most fitting to their portfolio. It also makes it questionable whether the extra use of data through historical PDs instead of idealised PDs would be worth the improvement. The PGGM model does contain a lower Brier score than the simulated A-IRB model, but only by a small margin.

6 Conclusion

This paper looked at the three of the more commonly used Credit Risk models, a form of historical simulations, the A-IRB model and the CreditRisk+ model, and how these performed for an insurance which covers a portfolio's first 10% tranche. The paper also looked at some small adjustments to the different models, such as a simplified CreditRisk+ model which makes use of historical PDs to create correlation or the Simulated A-IRB model which tries to include the probability of having no defaults. Furthermore a combination of the A-IRB model and the PGGM model, called Mixed model, was developed and tested if this could solve some of the weaknesses these models have. Lastly a validation was done for the different models, to see whether the different models are applicable and how they performed relative to each other.

It can be concluded that using a direct A-IRB model has weaknesses for short period portfolios or portfolios consisting of a small number of loans. This due to a lack of inclusion of a split between bad years and good years within the A-IRB model as well as a lack of possibility to include a probability to have no defaults within a time frame. A solution to the second problem can be found through a simulated A-IRB model. While this simulated A-IRB still underestimates the chance to be in a bad state, it does contain a more realistic chance to have a “worst case”-state by having a probability to lose the entire portion covered.

A small step has been made into including a split between bad years and good years, through looking at a mixed model, which would support this problem as raised by papers such as Bangia et al. (2002). This mixed model tries to combine the advantages of the A-IRB model with those of the historical simu-

lations. While it shows promise for an alternative to a historical simulation, this mixed model demands the availability of proper historical data as well as a good description of the current state of the world for the different ratings. While more historical data will be available in the future, solving this demand, capturing the current state of the world will always be difficult.

An alternative way of including the bad and good state of the world is estimating the PDs through an (pooled-)AR model. The outcome shows a more modest PD, as expected when looking at the year the PD was estimated for compared to some of the other crises. When comparing the AR model with the pooled-AR model, the pooled-AR model seems a better fit. One of the causes for this is the lack of data to properly estimate an AR model for each PD, but a more thorough comparison can be made once more data is available in the future.

A problem still occurring is the estimation of PDs of relatively safe loans, as can be seen in the outcome of the PDs estimated through the (pooled-)AR model. One expects higher ranked PDs to have a lower PD, yet some PDs, such as Aa2, have contained a higher PD when compared to some of the PDs ranked as less safe. The culprit can again be seen in the lack of available data. Proper estimates of the PDs of safer loans requires a proper amount of bad years, as this gives more defaults to work with, and the current data still lacks that. Once more data is available, this issue is expected to be less prominent.

When looking at model validation none of the models performs significantly better. While this allows the option for using the different models open, one can ask whether the more complex models are worth the effort. A rather basic model, such as the A-IRB model supported by the Basel committee, performs similarly without using complex techniques with many assumptions attached.

Further research can be done towards a relation between the recovery rate and the current state of the world. While a more spread recovery rate seems to have less influence on the IRR, as the portfolios tend to cover a good number of loans, thus evening out the recovery rate, a correlation between a bad state of the world and a lower recovery rate could influence the IRR. Since papers such as Altman et al. (2004) show evidence for this correlation, further research in this field for the models is important. Further research can also be made around finding alternative ways to include good and bad periods. One way could be to use two forms of PDs in a simulated A-IRB model to represent good and bad years, and combine these distributions to obtain a final distribution including this split between bad and good years. Having said that, a more spread portfolio seems to suffer less from this lack of bad and good years and thus for portfolios used in practice, it is questionable whether an adjustment to the simulated A-IRB model is required. Furthermore, the simulation across multiple years reduces the direct influence bad and good years have.

References

- Altman, E., Resti, A., and Sironi, A. (2004). Default recovery rates in credit risk modelling: a review of the literature and empirical evidence. *Economic Notes*, 33(2):183–208.
- Altman, E. and Saunders, A. (1997). Credit risk measurement: Developments over the last 20 years. *Journal of banking & finance*, 21(11):1721–1742.
- Bangia, A., Diebold, F. X., Kronimus, A., Schagen, C., and Schuermann, T. (2002). Ratings migration and the business cycle, with application to credit portfolio stress testing. *Journal of Banking & Finance*, 26(2):445–474.
- Basel Committee (2004). International convergence of capital measurement and capital standards, a revised framework.
- Basel Validation Group (2005). Studies on the validation of internal rating systems. (14).
- Cantor, R. and Falkenstein, E. (2001). Testing for rating consistency in annual default rates. *The Journal of Fixed Income*, 11(2):36–51.
- Cooley, J. W. and Tukey, J. W. (1965). An algorithm for the machine calculation of complex fourier series. *Mathematics of computation*, 19(90):297–301.
- Duffie, D. and Singleton, K. J. (1999). Modeling term structures of defaultable bonds. *Review of Financial studies*, 12(4):687–720.
- Hanson, S. G. and Schuermann, T. (2004). Estimating probabilities of default. *FRB of New York Staff Report*, (190).
- Hull, J. C. (2006). *Options, futures, and other derivatives*. Pearson Education India.
- Jarrow, R. A., Lando, D., and Turnbull, S. M. (1997). A markov model for the term structure of credit risk spreads. *Review of Financial studies*, 10(2):481–523.
- Kiefer, N. M. and Larson, E. C. (2004). *Testing simple Markov structures for credit rating transitions*. Comptroller of the Currency.
- Knuth, D. E. (1969). The art of computer programming: Seminumerical algorithms, volume 2, chapter 4.
- Kupiec, P. H. (2009). How well does the vasichek-basel airb model fit the data? evidence from a long time series of corporate credit rating data.
- Levich, R. M., Majnoni, G., and Reinhart, C. (2012). *Ratings, rating agencies and the global financial system*, volume 9. Springer Science & Business Media.
- Lütkebohmert, E. (2009). The creditrisk+ model. *Concentration Risk in Credit Portfolios*, pages 53–60.

- McNeil, A. J. and Wendin, J. P. (2007). Bayesian inference for generalized linear mixed models of portfolio credit risk. *Journal of Empirical Finance*, 14(2):131–149.
- Medema, L., Koning, R. H., and Lensink, R. (2009). A practical approach to validating a pd model. *Journal of Banking & Finance*, 33(4):701–708.
- Melchiori, M. R. (2004). Creditrisk+ by fast fourier transform. *YieldCurve*, August.
- Millossovich, P. et al. (2002). An extension of the jarow-lando-turnbull model to random recovery rate. Technical report, Working Paper Dipartimento di Matematica Applicata, Università di Trieste.
- Nickell, P., Perraudin, W., and Varotto, S. (2000). Stability of rating transitions. *Journal of Banking & Finance*, 24(1):203–227.
- Schechtman, R., Garcia, V. S., Koyama, S. M., and Parente, G. C. (2004). Credit risk measurement and the regulation of bank capital and provision requirements in brazil-a corporate analysis. *Banco Central do Brasil Working Paper*, (91).
- Siu, T.-K., Ching, W.-K., Fung, E. S., and Ng, M. K. (2005). On a multivariate markov chain model for credit risk measurement. *Quantitative Finance*, 5(6):543–556.
- Tasche, D. et al. (2000). Calculating value-at-risk contributions in creditrisk+.
- Vasicek, O. (1987). Probability of loss distribution. Technical report, Technical report, KMV Corporation.
- Wilde, T. (1997). Creditrisk+: a creditrisk management framework. *Credit Suisse First Boston*.