

THE DISTRIBUTION OF DEFAULTS AND BAYESIAN MODEL VALIDATION

ABSTRACT

Quantitative rating systems are increasingly being used for the purposes of capital allocation and pricing credits. For these purposes, it is important to validate the accuracy of the probability of default (PD) estimates generated by the rating system and not merely focus on evaluating the discriminatory power of the system. The validation of the accuracy of the PD quantification has been a challenge, fraught with theoretical difficulties (mainly, the impact of correlation) and data issues (e.g., the infrequency of default events). Moreover, models—even ‘correct’ models—will over-predict default rates most of the time.

This paper addresses the challenge by presenting a means of assessing when the realized default rate will differ from the expected default rate. We first set up the standard single factor framework for determining the distribution of defaults that one would expect in a bucket of exposures which, according to the rating system output, have a uniform PD. We then extend the framework to analyze the distribution of defaults across multiple buckets and the distribution of defaults in one bucket over several time periods.

Based on this framework, we present two Bayesian approaches in relation to two common modeling issues which occur in practice. The first approach provides a set of techniques to facilitate risk assessment in the absence of sufficient historical default data. It allows us to determine the posterior distribution of a PD, given zero realized defaults, thereby providing a framework for determining the upper bound for a PD in relation to a low default portfolio. The second approach provides a means for monitoring the calibration of a rating system. It allows us to determine the posterior distribution of the aggregate shock in the macro-economic environment, given a realized default rate. By comparing this distribution to the institution’s view of the stage of the credit cycle its borrowers are in, this approach provides useful insight for whether an institution should revisit the calibration of its rating system. This method allows one to determine that a calibration needs to be revisited even when the default rate is within the 95% confidence level computed under the standard approach.

AUTHOR

Douglas W. Dwyer

Copyright© 2006, Moody's KMV Company. All rights reserved. Credit Monitor, CreditEdge, CreditEdge Plus, CreditMark, DealAnalyzer, EDFCalc, Private Firm Model, Portfolio Preprocessor, GCorr, the Moody's KMV logo, Moody's KMV Financial Analyst, Moody's KMV LossCalc, Moody's KMV Portfolio Manager, Moody's KMV Risk Advisor, Moody's KMV RiskCalc, RiskAnalyst, Expected Default Frequency, and EDF are trademarks owned by of MIS Quality Management Corp. and used under license by Moody's KMV Company.

ACKNOWLEDGEMENTS

I am grateful to Joseph McQuown, Bill Morokoff, Martha Sellers, Jamie Stark, Roger Stein, Kenneth Wee, Sarah Woo and Andrew Zhang for helpful comments and suggestions. All remaining errors are, of course, my own.

Published by:
Moody's KMV Company

To Learn More

Please contact your Moody's KMV client representative, visit us online at www.moodyskmv.com, contact Moody's KMV via e-mail at info@mkmv.com, or call us at:

FROM NORTH AND SOUTH AMERICA CALL:

1 866 321 MKMV (6568) or 415 874 6000

FROM EUROPE, THE MIDDLE EAST, AND AFRICA CALL:

44 20 7280 8300

FROM ASIA, NEW ZEALAND, AUSTRALIA AND INDIA CALL:

813 3218 1160

1 INTRODUCTION

A single number bet on an American roulette wheel pays out 36 times the initial bet with a probability of 1/38. From the Casino's perspective, they will win if after playing 36,000 times there are fewer than 1,000 losses. On average, they would expect 947.4 losses with a standard deviation of 30.4. The probability that a Casino will lose money on 36,000 spins of the wheel is 4.1%.¹ The reason a Casino can afford such a narrow spread on a bet whose payoffs are so skewed is that a roulette table is carefully designed to yield independently, identically distributed random numbers.² This fact makes it relatively easy to determine whether the roulette table is functioning properly. For example, the odds of the Casino experiencing a loss rate of greater than or equal 2.9% rather than the 2.6% expected loss rate on 36,000 bets is 1 in 10,000 for a properly functioning roulette table.

In credit risk management, the issues are fundamentally more complicated, due to two characteristics. First, the 'true default rate' of an actual exposure is more difficult to ascertain than the probability of losing on a roulette table due to the limited number of existing default events and data availability issues. Second, there is correlation in default events. Put simply, there is some uncertainty regarding the underlying default probability of any given exposure and when defaults occur they tend to be cyclical, for example, a recession will trigger a cluster of default events. As an example, Figure 1 presents the one-year default rate from 1970–2005 for investment grade issuers as rated by Moody's Investors Service. The average default rate over this time period was 0.068%. Nevertheless, in 22 of these years there were actually zero defaults. Occasionally, the default rates for investment grade credits are elevated and these periods generally correspond to recessionary periods in the economy such as in 1973, 1982, 1991, and 2000–2002.

The implications of correlation for model validation are explicitly recognized by regulators in the Basel Accord as it pushes banks to improve their estimates of standalone risk in the form of borrower PDs and account for correlations in capital via the Asymptotic Single Risk Factor (cf., Basel Committee on Banking Supervision, 2005). Further, regulators recognize that tests of PD model calibration that assume no correlation are flawed and that current tests which allow for correlation will only identify extreme issues with a model's calibration (cf., Tasche 2005).

This paper provides a means for assessing when the realized default rate will differ from the default rate forecasted by a model. We present the standard single factor framework that can be used to describe the distribution of defaults in a bucket of exposures with uniform PDs, given the correlation of exposures in the bucket and the predicted PD for the bucket.³ According to this framework, correlations play a decisive role in the testing of the quality of model calibration. This framework can be extended for computing the distribution of defaults across multiple buckets and the distribution of the default rate for one bucket observed over multiple time periods.

The paper presents two Bayesian approaches for validating the accuracy of the forecasted PD estimates generated by rating models by building on the foundations of the single factor framework. The first approach allows us to determine the posterior distribution of the PD given the realized sample default rate and an uninformed prior. Specifically, suppose we know nothing about what the PD for a bucket should be, this could imply a prior distribution for the PD that is

¹These numbers are computed using the central limit theorem to approximate a binomial distribution with a normal distribution (cf., Stein, 2006).

²One can compare a single number bet on a roulette table to a loan. The casino loans the player \$35. If the player loses, he returns the \$35 plus \$1 interest. If the player wins, he defaults on the loan and keeps the \$35. Consequently, if the player loses he is paying a spread of 286 basis point (1/35 multiplied by 10,000 which we would refer to as total spread). While the casino wins 286 basis point with probability 37/38, it loses 10,000 basis points with probability 1/38. Therefore, the expected spread is a mere 15 basis points spread. This compares with a Ba bond whose long run average one-year default rate is about 1% and the average spreads on such bonds are approximately 3%, which translates into a 200 basis point spread even assuming an LGD of 100%.

³This framework – which has been called the Vasicek Model, a Gaussian Copula or a Gaussian one-factor model -- is commonly utilized by academics, practitioners and regulators. One application is level validation of a PD model – showing that the observed default rates are consistent with the PDs produced by a model. To our knowledge, the first application of this framework to the level validation of a PD model is Kurbat and Korabiev (2002). Other examples include determining regulatory capital and pricing for CDOs. Nevertheless, the Gaussian assumption has been criticized for understating the risk of multiple extreme events relative to other distributions (cf., Embrechts, 1999). Further, multi-factor models have been used to compute economic capital with greater realism.

uniformly distributed between 0 and 1.⁴ Given an observation of zero defaults for the bucket, the posterior distribution of PDs can be solved for given a correlation assumption. Therefore, one can determine an upper bound for what the PD could be, such that we are 95% certain that the PD is below this value. This approach will prove useful in validating the level of the PD in relation to models for low risk portfolios.

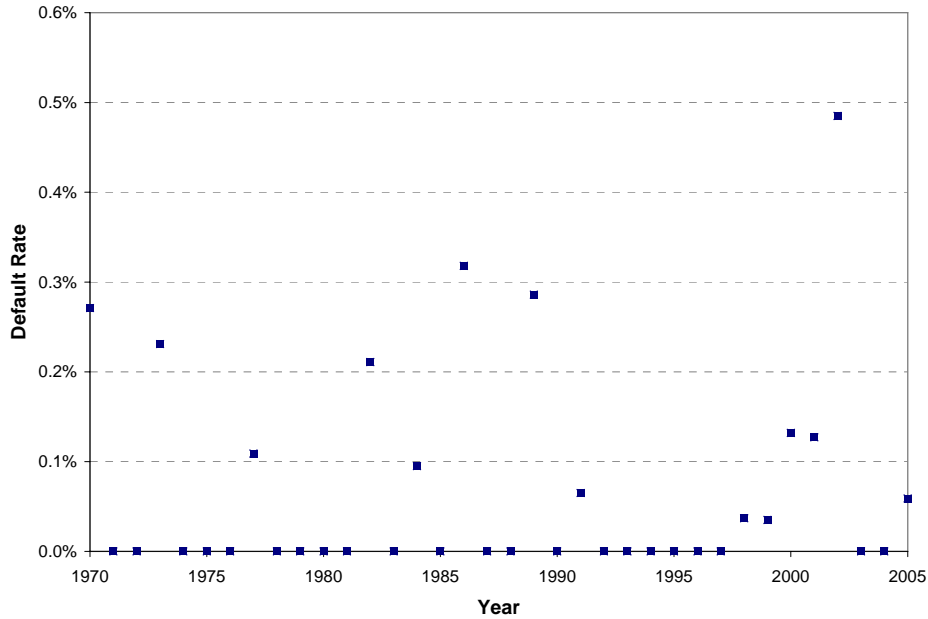


FIGURE 1 The Default Rate of Investment Grade Issues 1970–2005

Source: Exhibit 30 of Hamilton and Varma 2006.

The second approach allows us to test the forecasted PD against our beliefs regarding the distribution of the aggregate shock in the macro-economic environment given the realized default rate in the bucket—assuming that the PD model is correct. This distribution can be compared with knowledge regarding the state of the economy over which the model was tested. For example, suppose one observes a 3% default rate in a bucket of 1,000 exposures that had a PD of 1%. With a correlation of 0.2, one will observe such an elevated level of defaults more than 5% of the time. Nevertheless, we show that this level of default is only consistent with the predicted PD of 1% if there was a large negative shock that occurred during this time period. If the general business conditions surrounding the episode are not indicative of elevated default rates, then one should consider revising the calibration of the model that produced the PD of 1%. The standard approach of computing the p-value of the default rate for a given correlation assumption would not reach such a conclusion.

In the appendix, we present an application of this approach to the default rate time series for investment and speculative grade issuers from 1970–2005 as computed by Moody’s Investors Service. We show that the model can be calibrated to these time series. Further, the model is consistent with many properties of these time series. Finally, we show how the approach could be used to reject the hypothesis that the speculative grade default rate is 1.4% using data from either 2004 or 2005 even though the p-value associated with this hypothesis is less than 90 percent in both years under the standard approach.

⁴ In Bayesian inference, one combines a prior distribution with observed outcomes to determine a posterior distribution of the parameter in question. The prior distribution is intended to represent one’s knowledge of the problem at hand. There is always a certain amount of judgment involved in this choice. One justification for an uninformed prior is that it imposes as little of the statistician’s prejudice onto the data as possible. Further, from a technical standpoint a uniform prior is analytically convenient, which makes it a common choice for a parameter that is known to have both an upper and lower bound (cf., Gelman, et. al., 1995, Chapter 2). In Appendix C on page 26, we show how the results change for an alternative prior.

The framework employed in this paper has been used widely throughout the credit risk literature. The framework perhaps began with Vasicek (1987). The main contribution of this paper is the application of Bayesian methods to model validation. For example, many authors have used classical statistics to derive the set of possible PD s that would not be rejected given N observations and zero defaults.⁵ The first Bayesian application in this paper, in contrast, solves for the posterior distribution of the PD given N observations, 0 defaults and a prior distribution of the PD , from which a Bayesian style confidence interval can be derived. The second Bayesian application in this paper uses the realized default rate to derive the posterior distribution of the aggregate shock to the portfolio—given a PD , a sample size, and a correlation parameter. This distribution can then be compared to one’s general knowledge of the factors impacting the economy at that time, which in turn, provides a useful check on the plausibility of the model’s calibration.

In the context of correlated defaults, explicit analytic solutions are often not available. While others have employed analytic approximations and bootstrapping or both (cf., Tasche, 2003 and Hanson and Schuermann, 2005), our approach is to rely on numerical integration and Monte Carlo methods. For our taste, numerical integration and Monte Carlo methods allow for more direct derivations, albeit at the cost of requiring more sophisticated software implementations.

This paper is organized as follows. The next section shows how the standard one factor framework can be employed to compute the distribution of defaults using Monte Carlo techniques. The third section shows how to use Bayesian approaches to compute a posterior distribution for the PD in a zero default portfolio and to compute a posterior distribution for the aggregate shock given an observed sample default rate. The latter technique can be used to reject a model calibration even when the sample default rate falls within the 95% confidence level implied by the standard approach. A fourth section summarizes some data issues that must be addressed when applying the technique. The implications for model building and calibration are discussed in the conclusion. There are three appendices. The first shows the relationship between the standard single factor framework and regulatory capital requirements, and discusses what constitutes a reasonable assumption for the correlation parameter within this context. The second appendix relaxes the assumption of a uniform prior distribution for PD s. The third presents an application of the technique to the default rate time series for investment and speculative grade issuers as provided by Moody’s Investors Service.

⁵ See, for example, Christensen, Hansen and Lando (2004). Pluto and Tasche (2005) also use this technique in the context of a Gaussian one-factor model.

2 THE STANDARD SINGLE FACTOR GAUSSIAN FRAMEWORK

Consider a situation where there is a bucket of 1,000 exposures with a uniform PD of 1% (according to the output of a PD model in a rating system) and 5 defaults have been observed. Is the model output wrong? How does one evaluate the quality of the calibration of the model? What if one observes 15 or 25 defaults instead?

To address this issue, we seek to apply the work of Kurbat and Korablev (2002), Tasche (2005) and Stein (2006) to develop a single factor framework that determines the distribution of defaults in a bucket.⁶ In the later part of this section, we will present a number of analytic extensions to the framework. Appendix A: provides a recipe and corresponding SAS code to simulate the distribution of defaults in a bucket given the number of exposures, the PD and the correlation assumption.

A single factor Gaussian model can be motivated from the perspective of a structural model.⁷ Under a structural model, asset value follows a Brownian motion. At the horizon—in this case one-year—if the asset value falls below the boundary default occurs, where the default boundary can be interpreted as a required debt payment. Therefore, the likelihood of default is determined by leverage adjusted for volatility which is summarized as ‘distance to default.’ Given a PD, one can set the volatility of assets to one and solve for the level of leverage that produces this PD.

This single factor framework essentially relates changes in the distance to default, i.e., default probabilities, to asset return correlations. It decomposes PD into a function of a single factor that models the systemic risk, with a residual representing the idiosyncratic risk of an individual exposure. We start by assuming N exposures in a bucket. These exposures have a uniform probability of default denoted by PD . Assuming an underlying structural model, one can compute an *adjusted distance to default* (DD) as:

$$DD = -\Phi^{-1}(PD) = \Phi^{-1}(1 - PD)$$

where Φ^{-1} denotes the inverse of a cumulative normal distribution.⁸ In order to simulate a distribution of defaults, one uses a correlation model to determine changes to DD .

Suppose that

$$r_j = \sqrt{\rho} r_m + \sqrt{(1 - \rho)} e_j$$

$$r_k = \sqrt{\rho} r_m + \sqrt{(1 - \rho)} e_k$$

Where r_j and r_k are the asset returns of two firms, j and k , r_m is the return of the market portfolio, ρ is the percent of the variation of the firm’s returns explained by the market, which is assumed to be constant for the two firms. r_m , e_j and e_k are idiosyncratic draws, or shocks, from a standard normal distribution. The correlation between the asset returns of the two firms is given by:

⁶ This section will be written from the perspective of classical statistical inference. Specifically, we will take that $PD=1\%$ as the null hypothesis. Under classical statistics, one will reject the null hypothesis if the probability that an outcome as or more extreme than the observed outcome is below some critical value such as 5% or 1% given the null hypothesis. If one observes an outcome that is not extreme given the null hypothesis, then one ‘fails to reject’ the null hypothesis. A failure to reject a null hypothesis is not equivalent to confirming the null hypothesis with any degree of certainty—it does not show that the PD is correct.

⁷ In our view, greater realism can be achieved with a multi-factor model. For example, Moody’s KMV Portfolio Manager™ uses multiple factors representing industry, country, and region risk in addition to global economic risk.

⁸ Note that this approach differs from the approach used to map a DD to an EDF that is employed by the Moody’s KMV Public Firm Model. The Moody’s KMV Public Firm Model uses an empirical mapping.

$$\rho(r_j, r_k) \equiv \frac{\text{Cov}(r_j, r_k)}{(\text{Var}(r_j)\text{Var}(r_k))^{1/2}} = \rho$$

In a pure Merton model, the 1-year distance to default is defined as:

$$DD = \frac{\log(V_A / X) + (\mu - \sigma_A^2 / 2)}{\sigma_A}$$

Where V_A is asset value, X is the default point, μ is the expected return on assets (adjusted for cash outflows) and σ is the volatility of assets. Assuming that X , μ and σ remain constant over the period, the change in DD is given by:

$$\Delta DD = \frac{\log(V_{At}) - \log(V_{At-1})}{\sigma_A} = \frac{r}{\sigma_A}$$

where r is the return on the firm's assets. Therefore, a *portfolio outcome* can be determined by drawing one aggregate shock and N idiosyncratic shocks from a standard normal distribution. For a given PD, we set the asset volatility to 1 and solve for the degree of leverage that yields the appropriate PD.⁹ Therefore, the change in DD for exposure j is simply:

$$\sqrt{\rho} r_m + \sqrt{1 - \rho} e_j.$$

Default for exposure j occurs if the new DD is negative:

$$\sqrt{\rho} r_m + \sqrt{1 - \rho} e_j - \Phi^{-1}(PD) < 0$$

One can simulate a portfolio outcome by drawing one aggregate shock and an idiosyncratic shock for each exposure and compute the new DD for each exposure. One then simply counts the number of defaults for all exposures and this number forms a portfolio outcome. One then repeats this simulation many times recording the number of defaults for each simulation.¹⁰ From the distribution of the number of defaults across simulations, one can compute the median number of defaults, mean defaults, the 5th percentile of defaults, 95th percentile of defaults and so forth. The respective default rates are then derived by dividing the number of defaults by N . A recipe and sample SAS code are provided in Appendix A: that perform this calculation.¹¹

Table 1 below shows some results of this procedure for different numbers of exposures and assumptions of correlation, given a PD of 1%. The table presents the mean, the median, the 5th percentile and the 95th percentile of the default rate across 100,000 samples obtained from Monte Carlo simulations. We note three general observations from this set of results. First, in the absence of correlation the upper bound default rate remains close to the mean default rate for the exposures in the bucket. As gleaned from the case with 10,000 exposures and ρ equal to zero, the default rate in the 95th percentile is 1.17%. Second, upper bound default rates rise markedly with increasing correlation. Even where correlation is moderate (where ρ is equal to 0.2), the 95th percentile case becomes three times larger than the predicted PD of 1%. Finally, the median default rate is considerably lower than the prediction. The median being lower than the mean implies that most of the time a 'correct' model will over predict defaults.

⁹ Setting volatility to 1 is a scaling assumption.

¹⁰ For a sufficiently large number of observations, Vasicek has derived a closed form solution to this problem : See Appendix B.

¹¹ For this simple problem, one could also do numerical integration. For a one-dimensional integral, numerical integration is computationally more efficient—the same degree of accuracy can be achieved with fewer calculations. Nevertheless, simple extensions to the framework (e.g., a multi-period problem) will increase the dimensionality of the problem quickly. Higher dimensional problems are often easier to accommodate with Monte Carlo methods than numerical integration.

TABLE 1 Simulation Results Where The Actual PD is 1%

ρ	EXPOSURES	Mean Default Rate	Median Default Rate	P5	P95
0.0	100	1.0%	1.00%	0.00%	3.0%
0.0	1,000	1.0%	1.00%	0.50%	1.5%
0.0	10,000	1.0%	1.00%	0.84%	1.2%
0.2	100	1.0%	0.00%	0.00%	4.0%
0.2	1,000	1.0%	0.50%	0.00%	3.8%
0.2	10,000	1.0%	0.46%	0.03%	3.8%
0.4	100	1.0%	0.00%	0.00%	5.0%
0.4	1,000	1.0%	0.10%	0.00%	4.9%
0.4	10,000	1.0%	0.13%	0.00%	4.9%

Based on the results in Table 1, our response to the questions posed at the beginning of this section is as follows. Assuming that the PD of 1% is correct and there is no correlation between exposures, there is a 5 percent chance of observing 5 defaults or fewer (a default rate of 0.50% or smaller) and a 5 percent chance of observing 15 defaults or more (a default rate of 1.5% or bigger). With a modest amount of correlation (where ρ is assumed to be 0.2) there is a 5% chance of observing default rates that are 3.9% or higher. Therefore, observing 5, 15 or 25 defaults will not be unusual for a bucket with 1,000 exposures and a PD of 1%, given a reasonable amount of correlation (Appendix B.2 contains a discussion as to what constitutes a reasonable amount of correlation). In the language of classical statistical inference, one would not reject the null hypothesis of a PD=1% and $\rho = 0.2$.

To further illustrate the situation, Figure 1 presents the distribution of the default rate for 1,000 exposures with PDs of 1% and a ρ of 0.2. The grey lines mark the median, 95th and 99th percentiles of the distribution.

TABLE 2 Average Default Rate for Multi-Year Simulations

R^2	Exposures	Years	Mean Default Rate	Median Default Rate	P5	P95
0	1,000	4	1.0%	1.0%	0.75%	1.3%
0	1,000	8	1.0%	1.0%	0.83%	1.2%
0	1,000	12	1.0%	1.0%	0.85%	1.1%
0.2	1,000	4	1.0%	0.8%	0.20%	2.5%
0.2	1,000	8	1.0%	0.9%	0.34%	2.1%
0.2	1,000	12	1.0%	0.9%	0.43%	1.9%
0.4	1,000	4	1.0%	0.5%	0.03%	3.6%
0.4	1,000	8	1.0%	0.7%	0.11%	2.9%
0.4	1,000	12	1.0%	0.8%	0.19%	2.6%

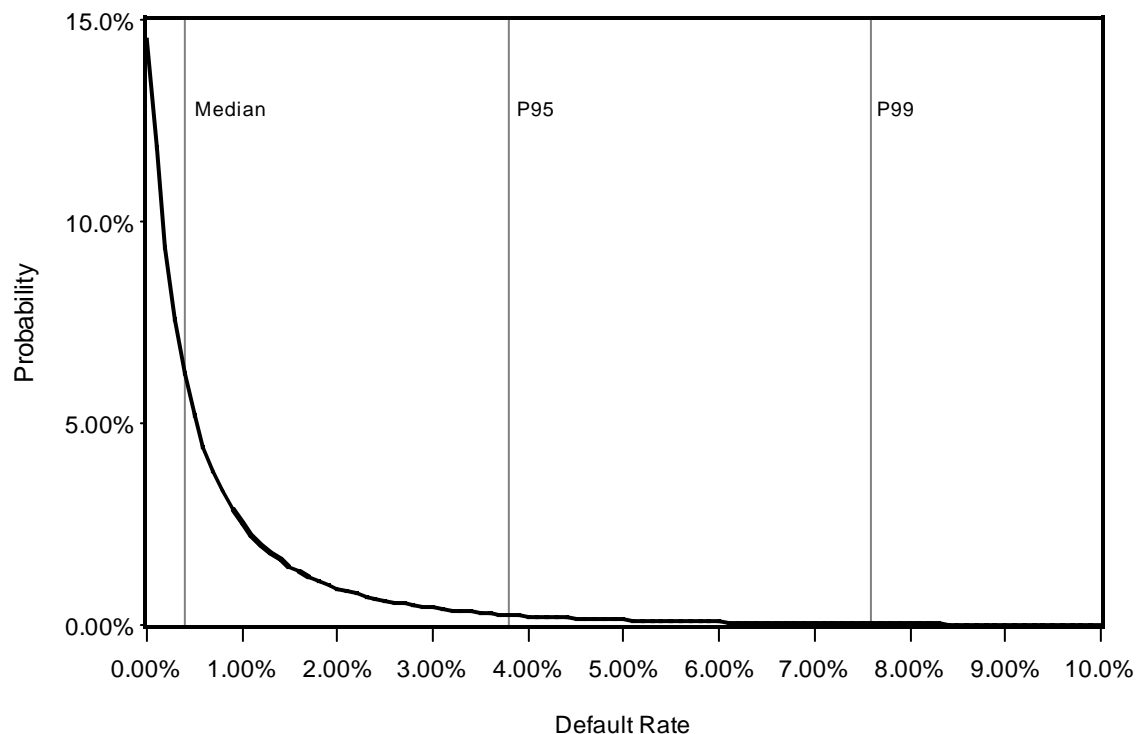


FIGURE 2 Distribution of Defaults

Using this framework, we extend the initial question to: What is the distribution of average default rates of a bucket over time, assuming that both the aggregate shocks and idiosyncratic shocks are both independent over time? Specifically, suppose that at the start of each year, we constructed a portfolio of 1,000 exposures all with a PD of 1%. What would be the average default rate of such a portfolio observed over several years? Table 2 presents the results from such an analysis with the time horizon varying from 4–12 years, for a portfolio of 1,000 exposures with a PD of 1%.¹²

The results are fairly striking. For a portfolio with moderate correlation (where ρ is assumed to be 0.2) over a 12-year period, the default rate can range from 0.43% to 1.86%. Nonetheless, the impact of correlation is reduced for portfolios with longer time horizons. For samples with a longer time horizon, the gap between the mean and median default rates shrink, indicating a more symmetrical and less skewed distribution.

Further, one may be interested in computing the expected number of defaults in a population with heterogeneous default probabilities. This scenario can be easily accommodated by computing the adjusted DD for each firm individually and modifying the computation in Step 4 of the recipe (Appendix A:) accordingly. One needs to specify the PD and ρ for each firm.

¹² We performed this analysis with a straightforward extension of the recipe provided in Appendix A: using 100,000 simulations.

3 ASSESSING PD CALIBRATION USING BAYESIAN APPROACHES

Now that we have established a framework for describing the distribution of defaults, we will show how the calibration of rating systems or PD models can be monitored using two different Bayesian approaches. Under Bayesian analysis, one defines a prior distribution characterized by one's assumption regarding the possible values for a parameter (or set of parameters) of interest. After observing an outcome, one updates the prior assumption, based on the actual outcome. In this section, we present two Bayesian approaches which can provide valuable insights in relation to two common modeling issues which occur in practice. The first approach helps determine the upper bound for a PD in relation to a low default portfolio or rating buckets with very few or zero observed defaults, while the second approach provides a means to monitoring the calibration of a rating system.

Under the first approach, we treat the PD as being unknown. We assume an uninformed prior—that we have no information regarding what the true PD would be. Consequently, our prior is that the PD is uniformly distributed between 0 and 1. We observe an outcome and update our prior on basis of the outcome. This is particularly useful for the case of observing 0 defaults.

Under the second approach, we treat the PD and asset correlation as being known and we update our expectation regarding the distribution of what the aggregate shock in the macroeconomic environment would have been in order to produce the observed sample default rate. One can then make a judgmental evaluation as to whether or not this distribution is consistent with one's knowledge of the general business conditions at the time.

The underlying analytics for both approaches are drawn from the foundation of the single-factor framework discussed above, and we will use the fact that the probability of default given the aggregate shock, r_m , is given by:

$$\Pi(PD, r_m) \equiv \Phi \left(\frac{\Phi^{-1}(PD) - \sqrt{\rho} r_m}{\sqrt{1-\rho}} \right)$$

3.1 Bayesian Updating with A Prior that the PD has a Uniform Distribution

High credit quality firms typically form a large proportion of the large corporate segment in a bank's portfolio. However, defaults are generally rare events for such investment grade exposures. Furthermore, high quality firms tend to be larger and more systematic (as reflected in the treatment of correlation as declining with increasing PDs in the Basel approach). Suppose that we have 10,000 of such exposures in a bucket where we have observed no defaults within 1 year. What can we say about the PD for this bucket? The maximum likelihood estimate of the PD is zero. However, economic intuition and longer-term data suggest that even the PDs for high credit quality exposures are unlikely to be stable over time and the credit cycle. Between 1970 and 2005, realized default rates for investment grade credits ranged from 0 to 0.485%, and in 22 of these years the group experienced zero defaults (Figure 1). As such, what degree of confidence can we place in this estimate of zero derived from data that is, like any single bank's, quite limited?

To address this issue, we propose to Bayesian update a prior distribution that reflects what the PD could be. The most uninformed prior as to what the PD could be is to say that it is uniformly distributed between [0,1]. Assuming zero correlation, the probability of D defaults and that the actual probability of default is less than or equal to PD is given by:

$$\text{Prob}(x \leq PD, D) = \int_0^{PD} P(D|x) f(x) dx$$

where

$$P(D|PD) = \binom{N}{D} PD^D (1-PD)^{N-D}$$

and $f(x) = 1$, i.e., the probability density function of a uniform random variable.

By Bayes' rule the posterior cumulative distribution of PD given D is the probability of the set $\{D, pd \in [0, PD]\}$ divided by the probability of the set $\{D, pd \in [0, 1]\}$. Therefore, the cumulative distribution of PD given D is given by:

$$K(PD|D) = \frac{\text{Prob}(x < PD, D)}{\text{Prob}(D)} = \frac{\int_0^{PD} P(D|x)f(x)dx}{\int_0^1 P(D|x)f(x)dx}$$

For zero defaults, this distribution has the solution:

$$K(PD | D = 0) = 1 - (1 - PD)^{N+1} .^{13}$$

Based on this formula, we present the posterior distribution of the PD in Table 3 for a bucket with one, five, ten and fifty thousand exposures, assuming zero defaults and no correlation. Where there is 1,000 exposures, one can be 95% certain that the true probability of default is less than 30 basis points. With 10,000 exposures, one can be 95 percent certain that the true probability of default is less than 3 basis points. For 50,000 exposures, the boundary becomes less than 1 basis point.

TABLE 3 Posterior Cumulative Distribution of PD given No Defaults and Assuming No Correlation

		Exposures			
		1,000	5,000	10,000	50,000
Probability of Default	0.01%	9.5%	39.3%	63.2%	99.3%
	0.02%	18.1%	63.2%	86.5%	100.0%
	0.03%	25.9%	77.7%	95.0%	100.0%
	0.04%	33.0%	86.5%	98.2%	100.0%
	0.05%	39.4%	91.8%	99.3%	100.0%
	0.06%	45.1%	95.0%	99.8%	100.0%
	0.07%	50.4%	97.0%	99.9%	100.0%
	0.08%	55.1%	98.2%	100.0%	100.0%
	0.09%	59.4%	98.9%	100.0%	100.0%
	0.10%	63.2%	99.3%	100.0%	100.0%
	0.20%	86.5%	100.0%	100.0%	100.0%
	0.30%	95.0%	100.0%	100.0%	100.0%

Allowing for correlation makes the problem somewhat more involved, because one needs to 'integrate out' the aggregate shock. Note that the probability that the pd is less than or equal to PD, the aggregate shock is less than or equal to r_m and that the number of defaults is equal to D is given by:

$$\text{Prob}(x \leq PD, y \leq r_m, D) = \binom{N}{D} \int_0^{PD} \int_{-\infty}^{r_m} \Pi(w, y)^D (1 - \Pi(w, y))^{N-D} n(y) f(w) dy dw$$

Where $n(\cdot)$ is a standard normal density function and recall that $\Pi(PD, r_m)$ is the probability of default given the aggregate shock. Therefore, the posterior distribution of PD is given by:

¹³ We can solve for the posterior distribution for different prior distributions. For example, suppose that the prior was uniform on $[0, \bar{P}]$ where \bar{P} is < 1 . Then one can show that the solution is $(1 - (1 - PD)^{N+1}) / (1 - (1 - \bar{P})^{N+1})$ on $[0, \bar{P}]$. As $(1 - (1 - \bar{P})^{N+1})$ is close to 1 for large N, this solution is very close to $(1 - (1 - PD)^{N+1})$ for large N. Therefore, the posterior distribution does not appear highly sensitive to the choice of prior for a large N. Appendix C: shows a similar result for another alternative prior distribution.

$$K(PD | D) = \frac{\binom{N}{D} \int_0^{PD} \int_{-\infty}^{\infty} \Pi(w, r_m)^D (1 - \Pi(w, r_m))^{N-D} n(r_m) f(w) dr_m dw}{\binom{N}{D} \int_0^1 \int_{-\infty}^{\infty} \Pi(w, r_m)^D (1 - \Pi(w, r_m))^{N-D} n(r_m) f(w) dr_m dw}$$

Since $\Pi(PD, r_m)$ is a nonlinear function of both PD and r_m it is unlikely that we will find an analytic solution to these double integrals. Nevertheless, they are amenable to numerical techniques.¹⁴

Table 4 presents a tabulation of the 95th percentile of this posterior distribution of the PD given zero defaults and a uniform prior for different numbers of exposures and correlations. As shown in Table 3, assuming zero correlation, for 10,000 exposures, 95th percentile PD is three basis points. Therefore, from a Bayesian perspective, if one observes 10,000 observations with no defaults one can be 95 percent certain that the actual PD is less than 0.03% under the assumption of no correlation. With a reasonable amount of correlation (0.2) and a large sample (100,000 exposures) a PD of 47 basis points can not be ruled out. For a high degree of correlation (0.5) a PD of 10.44% is possible even with zero defaults across 100,000 observations! This implies that even if the true probability of default for a bucket is relatively high, a large positive aggregate shock, r_m (i.e., where we are in the good part of the credit cycle), coupled with substantial correlation can result in an observed default rate that is zero or very close to zero in a single period.

TABLE 4 95th percentile of the Posterior Distribution of PD given Zero Defaults

		ρ							
		0	0.05	0.1	0.15	0.2	0.3	0.4	0.5
z	1,000	0.30%	0.74%	1.50%	2.65%	4.26%	8.92%	15.38%	23.27%
	5,000	0.06%	0.20%	0.50%	1.06%	1.97%	5.13%	10.36%	17.58%
	10,000	0.03%	0.11%	0.32%	0.72%	1.41%	4.05%	8.74%	15.59%
	50,000	0.006%	0.03%	0.11%	0.29%	0.65%	2.32%	5.88%	11.78%
	100,000	0.003%	0.02%	0.07%	0.19%	0.47%	1.83%	4.95%	10.44%

Based on this framework, institutions wishing to monitor the calibration of their rating systems on a zero default portfolio should be estimating suitable correlation values from which a 95th percentile PD can be calculated

3.2 Using Bayesian Updating to Size the Aggregate Shock

Suppose that we have a PD model that predicts a 1% default rate for a bucket with 1,000 exposures but we observed 30 defaults in a given year. According to Table 1 in Section 2, at a modest correlation of 0.2, we would expect such a default rate somewhat more often than 5% of the time. With zero correlation, such a default rate would be very unlikely. A possible inference is that the model may require re-calibration. Alternatively, one may infer that there must have been a large negative shock in the macro-economic environment in order to generate such an elevated default rate. However, how do we compare this inference to our knowledge of the general business conditions and credit cycle impacting these exposures during the relevant time period? When should we consider revisiting the calibration of the model? To address this issue, we begin by noting that the distribution of the sample default rate (SDR) given r_m can be approximated by a normal distribution with a mean of:

$$\Pi(r_M) = \Phi\left(\frac{\Phi^{-1}(PD) - \sqrt{\rho}r_M}{\sqrt{1-\rho}}\right)$$

¹⁴ As this function exhibits significant curvature when $D=0$ and N is large, we employ an adaptive quadrature method to evaluate the double integral.

and a variance of $\sigma_{r_M}^2 = (\Pi(r_M)(1 - \Pi(r_M))) / N$. Note that in contrast to Section 3.1, we have not included PD as a second argument to the function Π , because in this section we are treating PD as given. We have also approximated a binomial distribution by a normal distribution. Therefore, in this section we can easily work with density functions as both the random variables, r_M and SDR, are now continuous variables.

Based on this formula, Table 5 presents the probability of observing 30 or more defaults for different values of the aggregate shock, r_m , assuming a correlation of 0.2. The first column presents the aggregate shock which ranges from -2 to 0.¹⁵ The second column presents the probability that the aggregate shock is less than or equal to this value. The third and fourth columns present the mean and standard deviation of the default rate given the aggregate shock and the fifth column presents the probability that the default rate exceeds 3% given the aggregate shock. If the aggregate shock were greater than -1, then the probability of observing the given default rate is essentially zero. Therefore, one can conclude that if the model is correct then there was a negative shock that was in excess of -1. Such an event has a p-value of 16% [$\Phi(-1)$]. If the model is correct, then given that we observed a 3% default rate during the year, we can infer that this particular year was approximately a one-in-six year event for companies in this portfolio. If this inference does not seem plausible according to our observations of the credit cycle over the past year, then we should consider revisiting the calibration of the model.

TABLE 5

Aggregate Shock Size	Probability of Aggregate Shock or Worse	Default Rate Mean $\Pi(r_m)$	Default Rate Std Dev $\sigma_{r_m}^2$	P-Value that Default Rate Exceeds 3%
-2.00	2.3%	5.5%	0.7%	1.00
-1.80	3.6%	4.4%	0.7%	0.99
-1.60	5.5%	3.6%	0.6%	0.84
-1.40	8.1%	2.9%	0.5%	0.40
-1.20	11.5%	2.3%	0.5%	0.06
-1.00	15.9%	1.8%	0.4%	0.00
-0.80	21.2%	1.4%	0.4%	0.00
-0.60	27.4%	1.1%	0.3%	0.00
-0.40	34.5%	0.8%	0.3%	0.00
-0.20	42.1%	0.6%	0.2%	0.00
0.00	50.0%	0.5%	0.2%	0.00

One can take this analysis one step further, by deriving an expression for how certain we are that the aggregate shock must have been less than -1 through Bayesian methods. Specifically, we derive a posterior distribution for the aggregate shock given our prior that the distribution is a standard normal combined with the observed sample default rate (SDR) and taking as given the PD and the correlation assumption.

We can specify the probability density function of the sample default rate and the aggregate shock by:

$$f(SDR, r_M) = h(SDR | r_M)g(r_M)$$

Where g is the prior density function for the aggregate shock (a standard normal distribution), and $h(SDR|r_m)$ is the conditional distribution of SDR given r_m (a normal distribution with a mean of $\Pi(r_m)$ and a variance $\sigma_{r_m}^2$). The posterior density function of r_m is then given by:

¹⁵ As these are draws from a normal distribution with a mean of 0 and a standard deviation of 1, the units are standard deviations.

$$k(r_m | SDR) = \frac{f(SDR, r_m)}{\int_{-\infty}^{\infty} h(SDR | r_m) g(r_m) dr_m}.$$

The posterior cumulative distribution function is simply:

$$K(r_m | SDR) = \int_{-\infty}^{r_m} k(x | SDR) dx.$$

While both $h(SDR | r_m)$ and $g(r_m)$ have normal distributions, the mean and variance of $h(SDR | r_m)$ are non-linear functions of r_m . Consequently, deriving an analytic expression for this integral does not appear to be straightforward. This integral can easily be approximated using the trapezoidal rule. The denominator can be calculated as:

$$\int_{-\infty}^{\infty} h(SDR | r_m) g(r_m) dr_m \approx \frac{1}{\sqrt{2\pi}(N+1)} \sum_{i=0}^N \exp\left(-\frac{1}{2}\left(zscore(r_i)^2 + r_i^2\right)\right)$$

Where $zscore(r_m) = \frac{SDR - \Pi(r_m)}{\sqrt{\sigma_{r_m}^2}}$, and $r_i = -5 + \frac{i}{10N}$, i.e., $r_m(i)$ runs from -5 to 5.

Similarly, the numerator is calculated in the same way, except that the summation runs from 0 to $N(r_m)$ rather than to N , where $N(r_m) = \max\{r_m(i)\}$ such that $r_m(i) < r_m$.

Based on this approach, we present two examples in Figure 3 of how we might monitor the calibration of the model. The left-hand panel of Figure 3 presents the prior and posterior CDF for r_m in our example. The dotted lines cross the CDF at the 95th percentile value. Given our prior, we are 95 percent certain that the aggregate shock is less than -1.16. The probability of the aggregate shock being less than or equal to -1.16 is 12 percent $\{\Phi(-1.16)\}$. Therefore, observing a default rate of 3% when the model predicts a default rate of 1% is only consistent with the default experience being the worst of 1 out of 8 years, i.e., at or close to the bottom of the credit cycle.

The right-hand panel of Figure 3 presents the result of how we can interpret a particularly low default rate. If one takes as given that PD=10% and $r=0.2$ and one encounters a sample default rate of 1%, then one could be 95 percent certain that the aggregate shock was larger than 1.3. Such an aggregate shock would occur approximately one out of 10 years $\{1 - \Phi(1.3)\}$. Therefore, unless the macroeconomic environment was particularly favorable, such a low sample default rate would suggest that the model was overly conservative.

Note that both of these examples are fairly extreme. In the first, the sample default rate was three times the predicted and in the second the sample default rate was one-tenth. This does reflect the nature of credit as a unique asset class. Most outcomes in any one year for any one bucket will not provide enough information to conclude that the model produces PDs that are either too high or are too low with any degree of certainty (cf., Stein, 2006 and Tasche, 2005). Nevertheless, a large departure from the predicted value is informative regarding the aggregate shock that must have occurred. If this information is inconsistent with one's knowledge of the credit environment, one can reject a model calibration even though the observed default rate was within the 95 percent confidence level.

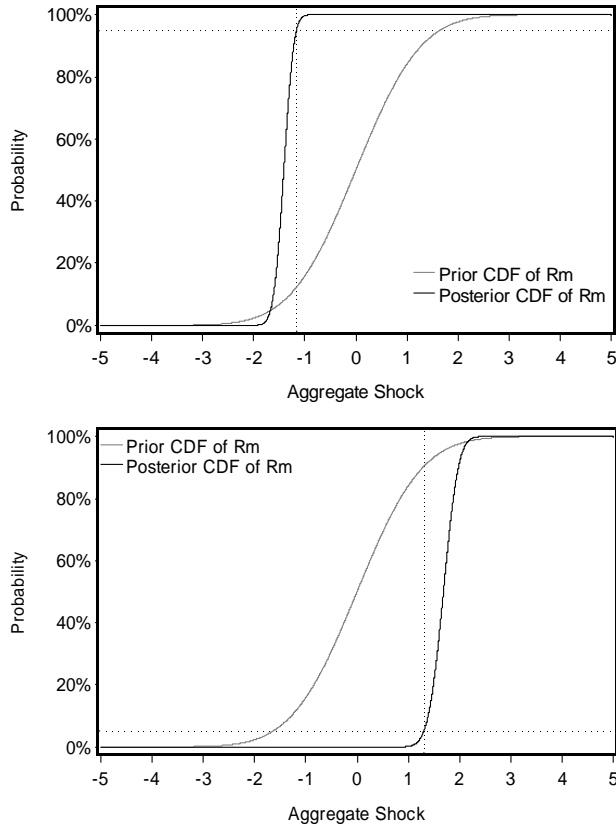


FIGURE 3 Prior and Posterior CDF of the Aggregate Shock

The left-hand panel presents the prior and posterior distribution of the aggregate shock given a sample default rate of 3% from a bucket of 1,000 exposures with a PD of 1% and a correlation of 0.2. The right-hand panel presents the prior and posterior distribution of the aggregate shock given a sample default rate of 1% from a bucket of 1,000 exposures with a PD of 10% and a correlation of 0.2.

This analysis can easily be extended to the case of several buckets, because given the aggregate shock the distribution of defaults in each bucket is a binomial distribution that can be approximated by a normal distribution for a large sample and the distribution of the sum of normal distributions is a normal distribution.

Institutions that seek to calibrate their rating models correctly need to incorporate a methodology that accounts for the credit cycle. As macroeconomic shocks cause spikes in default rates unevenly across the portfolios of financial institutions, these institutions need to be able to quantify whether their rating systems are accurately capturing changes in credit quality with appropriate distributional shifts of defaults across rating buckets.

This method will allow one to determine that a model calibration needs to be revisited where the standard approach of computing a confidence level for a given correlation assumption (cf., Table 1, or Tasche 2005) will not. Consider the previously mentioned example of observing 30 defaults in a bucket of a 1,000 exposures with a PD of 1%. Under a standard approach, one would not reject the hypothesis of a PD of 1% given 20% correlation. If 30 defaults were observed in an expansionary period in the economy, however, one would conclude that the model calibration needed to be revisited.

4 DATA CONSIDERATIONS

The above methodology is developed under the assumption that the data is clean – that a default is a clearly defined and meaningful concept, all defaults are captured in the data and each default can be properly linked to the relevant information regarding the firm (e.g., financial statements). Data samples used to develop, validate and calibrate private firm default risk models typically do not have all of these characteristics. Consequently, the interpretation of level validation studies must acknowledge and adjust for the data deficiencies. These issues include sample selection biases, data collection issues and differences in the definition of default. There are similar issues associated with level validation of ratings (cf., Hamilton and Varma, 2006) and EDF measures for publicly traded firms (cf., Kurbat and Korablev, 2002). Arguably, these issues are larger and more difficult to manage for the validation of private firm models. This section provides an overview of the data considerations encountered in validating private firm models.

Sample Selection Issues

When banks originate a new loan, they will ask for the two most recent financial statements in the loan application and enter both into their systems.¹⁶ In such cases, it will be impossible for a firm to default within 15 months of the first financial statement, because the originating bank does not yet have a loan outstanding for the obligor to default on. Additionally, the possibility of a firm defaulting within 15 months of the second financial statement will often be limited. For example, suppose one applied for a loan at the end of September of 2003 with financial statements dated December 31st for both 2001 and 2002. In order to default within 15 months of the second financial statement, the borrower would have to default before April of 2004, which is within 6 months of applying for the loan. The process of a loan becoming approved, the funds being disbursed, the ceasing of payments and finally becoming classified as a defaulter can easily use up more than these six months. We often find significantly lower default rates on the first and second statements for a firm in our own modeling on the data from bank databases (Dwyer, 2005). When part of the sample of historical financial statements lacks any chance of defaulting, the result is a sample default rate lower than the population default rate.

The window for default chosen for a study may also omit many defaults. For example, suppose one defined a 1 year default as a default within 3-15 months of a financial statement. This definition omits defaults that did not occur within 15 months of the last financial statement. We find that many private firm defaulters do not deliver their last financial statement. For example, suppose a firm's fiscal year ends in December and that they defaulted in May of 2005. They probably would have delivered their December 2003 financial statement in February or March of 2004. There is a good chance, however, that they would not have delivered their December 2004 financial statement in February or March of 2005, because this would have been about the time that they stopped making payments on their debt. In the Moody's KMV Credit Research Database (CRD), 25% of the defaults do not have financial statements within 15 months of their default date in North America. This type of sample selection bias will also result in the sample default rate being lower than the population default rate.

A third form of sample selection bias can result from including observations for which the default window extends past the end of the data collection period. For example, consider a financial statement dated December 31 of 2004. For such a statement, a 15 month default window would end on March 31, 2006. Suppose that the data was gathered in September of 2005. Then firms that defaulted between October of 2005 and March of 2006 would not yet have been captured as having defaulted in the database. The technical term for this issue is *right-hand censored* data (cf., Greene, 2000, chapter 20.3). This type of sample selection bias will also result in the sample default rate being lower than the population default rate.

Data Collection Issues

All financial institutions are in the process of implementing systems to capture defaults events. Default events are typically captured in a different system than financial statements, and there are often problems in linking defaults to

¹⁶Of course, the extent to which this practice is implemented will vary both across financial institutions and between business units within each financial institution.

financial statements. Active, performing obligors are more likely to have historical data maintained in bank systems than inactive, defaulted borrowers. Post-merger systems integration will often lead to database problems. The net effect is often an under-representation of defaults (c.f., Dwyer and Stein, 2006).

Default Definition Issues

Basel II provides a single default definition that is intended to be the world-wide standard for default. This standard can be characterized as 90 days past due and everything else that could be considered a default or a bankruptcy. Achieving uniform application of this standard is taking some time. Many countries have only recently begun collecting 90 days past due information. Consequently, most data sets that collect defaults over a long-time period do not consistently apply the Basel definition of default throughout the data collection process. This will tend to bias downward the default rates observed in a sample relative to the Basel definition.

Implications of Data Issues for Model Building and Validation

Private firm default risk models are often not calibrated to exactly match the average default rate in the development data set. In fact, many modelers chose to set the average default rate in the model above the actual default rate in the model development sample. This difference is intended to adjust for the extent to which the data set does not capture all default events and the extent to which the data set does not reflect a full credit cycle. Further, adjustments may be made to account for differences in how the definition of default is defined in the database and the targeted definition of default. Often the target definition is either the expectation of a real credit loss or the Basel Definition of default or both. The chosen average default rate for a model, often referred to as the central default tendency, is validated relative to a wide variety of available sources for overall reasonableness. Basel II explicitly recognizes these issues (cf. paragraphs 417, 448 and 462 of *A Revised Framework*).

In validating a model, one needs to take into account the same issues to the extent possible when comparing the realized default rates to the actual default rate: is the coverage of default events consistent and complete? What is the definition of default? What is the stage of the credit cycle?

5 CONCLUSION: IMPLICATIONS FOR MODEL BUILDING, CALIBRATING AND VALIDATING

A correct model for single-firm default risk will over-predict defaults most of the time. There are three reasons for this phenomenon. First, low PDs have skewed distributions in small samples. Second, correlation of negative events will skew the distribution of defaults as well. Third, data issues will generally understate the realized default rate on many samples.

These observations have a number of implications for building and calibrating models. For example, the fact that a correct model should over-predict defaults in most years needs to be taken into account during the model building process. By maximizing the log-likelihood of a model on a particular sample, one fits the average PD of the model to that of the development sample.¹⁷ If one does so, on a sample that does not contain a recession, then the average PD produced by the model will be too low. Table 2 can be interpreted to give a sense of how large this effect could be. If one were to use 4,000 observations across 4 years to build a model with a correlation of 0.2, then the median default rate on such a sample will be 20% below the true default rate and the 5th percentile is 80 percent smaller than the true default rate.¹⁸ Therefore, 50% of the time one would understate default risk by more than 20% and 5% of the time one would understate default risk by 80%.

If one observes a model that consistently under-predicts defaults, then one can conclude that the PDs are too low with some degree of confidence. Except for times during a recession, a standalone PD model will over-predict defaults most of the time. Under certain circumstances one can show that the PDs produced by a model are too high. For example, a model that consistently over predicts defaults during recessions is likely to be too conservative. Because of the skewness of the default distribution, however, it is easier to show that the PDs are too low than too high. By bringing knowledge of the economic conditions facing the portfolio, this approach allows one to determine that a calibration needs to be revisited even when the default rate is within the 95% confidence level computed under the standard approach that allows for correlation.

When applying the framework to actual data, one must be cognizant of the data issues involved, and consider the following questions: How complete is the data set? Does it capture all defaults? Has the default definition changed over time? What is the stage of the credit cycle? Without a large data set that spans more than one credit cycle, that relies on a consistent definition of default, contains uniform accounting practices and utilizes consistent and complete data collection mechanisms, there will be significant intrinsic uncertainty regarding the actual level of risk associated with borrowers. Consequently, level validation of private firm default risk models will continue to be a challenging and interesting exercise for some time to come.

Appendix D: presents an application of the technique using two long time series of default rates (speculative grade and investment grade issuers as rated by Moody's Investors Service). With a long time series, this framework can be extended to actually calibrate the correlation assumption to be consistent with the variability of default rates observed in the time series. For example, if one observes a relatively constant default rate across a large number of homogenous exposures over time then such data is consistent with a low degree of correlation. In contrast, if the default rate changes markedly over time, this is evidence of substantial correlation.¹⁹ In Appendix D:, we estimate that speculative grade bonds exhibit a correlation of 11% while investment grade bonds exhibit a correlation of 10%. Further, we show that this technique could be used to reject the hypothesis of the PD of speculative grade defaults being 1.4% using only 2004 or 2005 data even though the p-value of the sample default rate observed in these years given this hypothesis is less than 90%. The

¹⁷ For example, one can show that in a logistic model an implication of the first order conditions is that the average default rate will equal the average of the predicted probability of default (cf., Maddala, 1983, page 26).

¹⁸ This interpretation should be thought of as follows. Suppose one had many samples that are statistically independent in terms of both the aggregate and idiosyncratic shocks. If one were to build many different models across these samples by maximizing a likelihood function that does not allow for correlation, then 50% of the time the average PD produced by the model would be 20% below the true value and 5% of the time the average PD produced by the model would be 80% below the true value.

¹⁹ Moody's KMV has used this approach to estimate correlations for retail exposures. Other examples include Cantor and Falkenstein (2001) and Duffie, et. al., (2006). While neither of these two explicitly uses a Gaussian one-factor framework, both estimate parameters associated with an underlying latent variable. Examples that do use a Gaussian one-factor model include Jakubík (2006) and Rösch (2003).

reason is that the observed default rates in 2004 & 2005 would have implied a negative aggregate shock which is inconsistent with economic conditions during these years.

APPENDIX A: RECIPE FOR COMPUTING THE DISTRIBUTION OF DEFAULTS

Given N , PD , ρ , $NumSims$

1. Draw r_m from a standard normal distribution
2. Draw e_j from a standard normal distribution for each exposure
3. Compute $\sqrt{\rho} r_m + \sqrt{1-\rho} e_j - \Phi^{-1}(PD)$ for each exposure and flag the negative values as defaulters
4. Sum up the number of defaulters and record
5. Repeat $NumSims$ times
6. Divide the number of defaults for each simulation by N to convert to a default rate (if desired)
7. Compute the p5, p10, median, and mean of the number of defaults or the default rate as desired across the simulations

SAS Code:

```
%let rho = 0.2;
%let pd = .01;
%let Exposures = 1000;
%let NumSims = 10000;

data test;
  do Sims =1 to &NumSims;
    default=0;
    rm = normal(0);
    do n = 1 to &exposures;
      ei = normal(0);
      dd= &rho**.5*rm + (1-&rho)**.5*ei;
      default = default + (dd<probit(&PD))/&exposures;
    end;
  output;
end;

proc univariate data=test ; var default ; run;
```

APPENDIX B: RESULTS FOR A LARGE NUMBER OF EXPOSURES AND THE RELATIONSHIP TO REGULATORY CAPITAL REQUIREMENTS

There are some well-known closed-form results for the distribution of defaults as the number of exposures in the portfolio goes to infinity (cf., Vasicek 1987 and 1991; and Gordy, 2002). These results form the basis for the capital allocation rules used in Basel II. We will show these results assuming that a portfolio can be divided up into a finite number of buckets with homogeneous risk.²⁰ All exposures in the portfolio are assumed to be subject to the same aggregate shock. Further, we assume that all the exposures within a bucket have the same correlation with the aggregate shock. We do allow, however, for correlation to differ across buckets. In order to derive the results, some additional notation is required:

Notation:

- PD_i = Probability of default over the horizon for bucket i
- B = Number of buckets
- ρ_i = Asset correlation for bucket i
- X_{ij} = An indicator variable for whether or not exposure j in Bucket i defaulted
- N_i = # of exposures for bucket i
- w_i = The proportion of exposures in bucket i

The default rate for the portfolio is given by

$$D_N = \sum_{i=1}^B \left(\frac{w_i}{N_i} \sum_{j=1}^{N_i} X_{ij} \right).$$

It will be convenient to define a quantile variable, u , for the aggregate shock, defined as $u = \Phi^{-1}(r_m)$. Therefore, the probability that the aggregate shock r_m exceeds $\Phi^{-1}(r_m)$ is u .

The probability limit of the default rate ($N_i \rightarrow \infty \forall i$) conditional on r_m is given by:

$$L(u) \equiv \text{plim}_{N_i \rightarrow \infty \forall i} (D_N | r_m) = \sum_{i=1}^B \left(w_i \Phi \left(\frac{\Phi^{-1}(PD_i) - \sqrt{\rho_i} \Phi^{-1}(u)}{\sqrt{1 - \rho_i}} \right) \right).$$

Note that the probability limit concept requires that the number of exposures in every bucket goes to infinity. Further, $L(u)$ is strictly decreasing in u . Therefore, the probability that the large sample default rate is less than $L(u)$ is equal to u . We will refer to $L(u)$ as the loss function. We can also define a recovery function and a cumulative distribution for defaults:

$$R(u) = 1 - L(u)$$

We define the cumulative distribution of defaults as the inverse of the recovery function.

$$F(L) = R^{-1}$$

A cumulative distribution function takes as an input a value for a random variable (in this case the large sample default rate, L) and returns the probability that the random variable will be less than this value (in this case the probability that the default rate will be less than L). So the cumulative distribution for the default rate returns the probability that the

²⁰ We will assume that each exposure is of equal size. Gordy (2002) demonstrates how to relax this assumption.

large sample default rate is less than L . The probability density function for the default rate is the derivative of the cumulative distribution, which is given by:

$$f(\mathcal{D}) = \frac{\partial R^{-1}}{\partial L} = \frac{-\partial u}{\partial L}$$

If there is only one bucket, this $L(u)$ can be inverted by allowing u to be expressed as a function of L :

$$u = \Phi \left(\frac{\Phi^{-1}(PD) - \sqrt{1-\rho}\Phi^{-1}(L)}{\sqrt{\rho}} \right)$$

which yields the density function through differentiation:

$$f(L) = -\frac{\partial u}{\partial L} = n \left(\frac{\Phi^{-1}(PD) - \sqrt{1-\rho}\Phi^{-1}(L)}{\sqrt{\rho}} \right) \left(\sqrt{\frac{1-\rho}{\rho}} \right) \left(\frac{1}{n(\Phi^{-1}(L))} \right)$$

where n denotes the probability density function for a standard normal.²¹ This expression can be shown to be equivalent to the density expression provided by Vasicek (1991).²²

For more than one bucket, explicitly solving for u may not be possible. Nevertheless, the cumulative distribution function is implicitly defined from the loss function.

B.1 Relation to Regulatory Capital

Under Paragraph 272 of the Revised Framework, the capital requirement of an individual exposure is given by:

$$K = LGD \times N[(1-R)^{-.5}G(PD) + (R/(1-R))^{.5}G(.999)] - PD \times LGD \times (1-1.5*b)^{-1} \times (1 + (M-2.5) \times b)$$

Where LGD is loss given default, N is the normal distribution, R is asset correlation, G is the inverse of a normal distribution, M is maturity and b is a maturity adjustment. Note that the notation actually used in the framework was followed.

When the maturity of a loan equals one, the term $(1-1.5*b)^{-1} \times (1 + (M-2.5) \times b)$ drops out. Replacing N with Φ , and G with Φ^{-1} and R with ρ_i yields:

$$K = LGD \left(\Phi \left(\frac{\Phi^{-1}(PD_i) + \sqrt{\rho_i}\Phi^{-1}(.999)}{\sqrt{1-\rho_i}} \right) - PD \right)$$

or

²¹ Use the Leibniz integral rule to differentiate through the cumulative normal, i.e., if $y = \int_{-\infty}^x n(w)dw$ then $\frac{dy}{dx} = n(x)$.

²² Simply write out the expression using $\frac{1}{\sqrt{2\pi}} \exp\left(\frac{-x^2}{2}\right)$ for the normal distribution, make the two necessary changes to notation (change L to θ and PD to ρ) and simplify.

$$K = LGD \left(\Phi \left(\frac{\Phi^{-1}(PD_i) - \sqrt{\rho_i} \Phi^{-1}(.001)}{\sqrt{1 - \rho_i}} \right) - PD \right)$$

Since $\Phi^{-1}(.999) = -\Phi^{-1}(.001)$. Therefore, regulatory capital for a given exposure is based on the 1 in 1,000 worse case scenario as measured by the Vasicek distribution. It is multiplied by LGD to account for recovery. Expected losses (PD*LGD) is subtracted out under the view that the bank would be earning enough spread to cover the expected losses.

The capital requirements formula in Basel II is derived as the limiting case of one factor model. The level of capital requirements ensures sufficient capital to cover the 1 in 1,000 worst case scenario. Therefore, the capital requirements are a reasonable approximation for the 1 in 1,000 value at risk provided: that (1) the largest exposure in a bank's portfolio is relatively small and that (2) a single factor model is a reasonable approximation for the systematic risk that a bank's portfolio is subject to. There are at least three other issues that one should consider: (1) credit migration risk associated with exposures that are downgraded but do not default; (2) the tendency of the normal distribution to understate the risk of correlated extreme events and (3) the validity of the correlation parameter.

B.2 What Represents a Reasonable Value for Correlation?

In the Moody's KMV Portfolio Manager Product, correlations range from 0.1 to 0.65. For private firms, we provide a model to predict the correlation of private firms on the basis of their country, industry and size. Nevertheless, these correlations are designed to work in the context of a multifactor model. The relationship between the correlation that is appropriate for a multi-factor model and that of a single-factor model is unclear. We do find that correlation increases in the size of the firm.

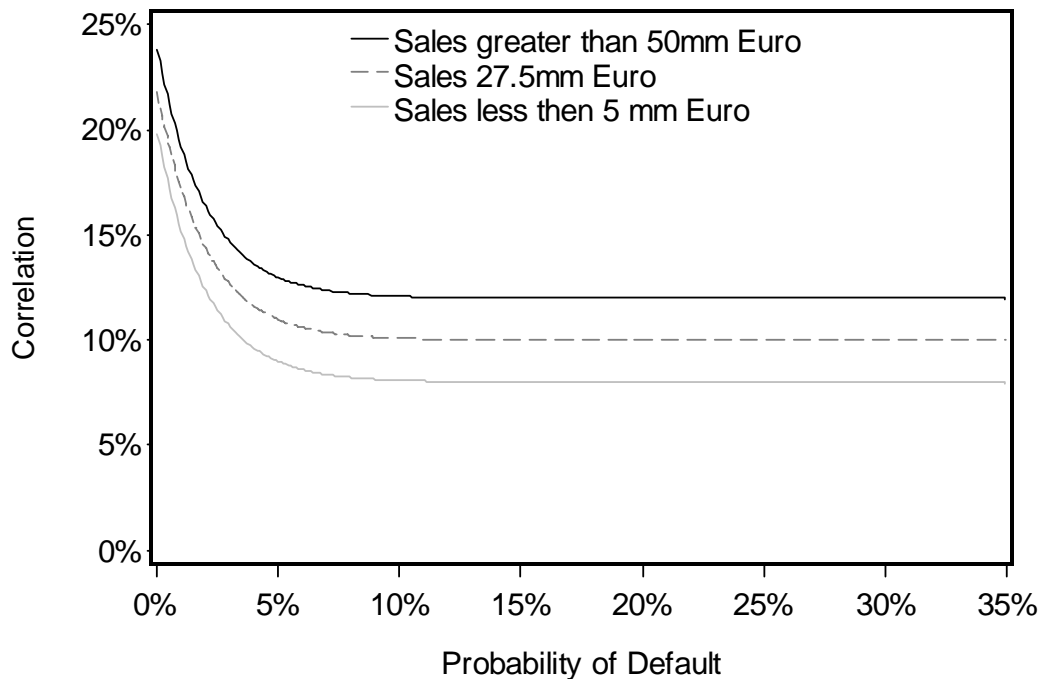


FIGURE 4 Correlation and Probability of Default from Basel II

Figure 4 presents the relationship between correlation and probability of default as prescribed by paragraphs 272&273 of 'A Revised Framework' for different firm sizes.

Basel II requires a formula for correlation in computing regulatory capital that is based on industry practices. This formula could be employed as a reasonable starting point. Paragraphs 272 and 273 of *A Revised Framework* provide the following formula:

If $S > 50$ then:

$$\text{Correlation (R)} = .12 * \text{weight} + .24 * (1 - \text{weight})$$

If $S \in [5, 50]$

$$\text{Correlation (R)} = .12 * \text{weight} + .24 * (1 - \text{weight}) - .04 * (1 - (S - 5) / 45)$$

If $S \leq 5$

$$\text{Correlation (R)} = .12 * \text{weight} + .24 * (1 - \text{weight}) - .04$$

Where $\text{weight} = (1 - \exp(-50 * PD)) / (1 - \exp(-50))$ and S is total annual sales in millions of Euros.

This formula implies that a large firm's correlation is bounded between 0.12 and 0.24 and that the smallest firm is bounded between 0.08 and 0.20 (See Figure 4).

APPENDIX C: AN ALTERNATIVE TO THE UNIFORM DISTRIBUTION ASSUMPTION IN THE CASE OF NO DEFAULTS WITH ZERO CORRELATION

In this section, we show for a prior distribution on the PD that is very different from the uniform, the posterior distribution is very similar for a large sample size given no correlation.

One method for establishing an upper bound for a PD in the case of zero defaults and no correlation is based on an answer to the following question: what is the smallest PD such that the probability of observing zero defaults is 5%?²³

The answer to this question is solving $0.05 = (1 - PD)^N$ for N .²⁴ Note the similarity between this solution and the solution to the 95th percentile of the posterior distribution given a uniform distribution: $0.05 = (1 - PD)^{N+1}$. We can give the answer to this question a Bayesian interpretation, by finding a prior distribution that would yield a posterior distribution consistent with this answer.

Let the prior distribution of PD be given by:

$$f(PD) = \frac{\phi}{(1 - PD)} \quad \text{on } [0, \varphi)$$

$$0 \text{ otherwise}$$

Where $\varphi < 1$ and ϕ is $(-\log(1 - \varphi))^{-1}$. This distribution places a larger weight on the larger PDs. In fact, the distribution is only well defined if φ is less than 1, as otherwise the density would approach infinity as PD approaches 1. The constant, ϕ , simply ensures that the density function integrates to one.

For this prior, the posterior distribution of PD can be derived as:

$$k(PD | SDR = 0) = \frac{\int_0^{PD} (1 - w)^{N-1} dw}{\int_0^{\varphi} (1 - w)^{N-1} dw} = \frac{1 - (1 - PD)^N}{1 - (1 - \varphi)^N}$$

Clearly, the limit of the posterior as φ approaches 1 is $1 - (1 - PD)^N$ and the 95th percentile of this distribution is determined by the solution to: $0.05 = (1 - PD)^N$. Therefore, a rather significant change in the prior distribution has only a modest change in the posterior for a large N .

The Bayesian approach also allows one to state that we are 95 percent certain that the true value is in the range approximated by $[0, 3/N]$ given the prior. Under the classical approach, one can state that any null hypothesis that the PD was not in this region would have been rejected with 95% certainty (cf., footnote 1 of Pluto and Tasche 2005). Nevertheless, this region is not a confidence interval in the sense of classical statistics. Under classical statistics, the confidence interval is random while the true value is fixed. A 95% confidence interval has the interpretation that the probability that the confidence interval contains the true value is 0.95. The difficulty of constructing such an interval in the case of a binomial distribution results from the variance of the distribution changing with the mean of the distribution (cf., Section 6.3, Hogg and Craig, 1978).

²³ Pluto and Tasche (2005) take such an approach

²⁴ Solving this equation yields one form of the so-called *rule of threes*. Taking logs of both sides and using the facts $\log(.05) \approx -3$ and $\log(1 - PD) \approx -PD$ yields $PD \approx 3/N$, i.e., that the minimum PD is three divided by the sample size. Put differently, suppose a portfolio has $3/PD$ exposures and that each have a probability of default of PD . Then one can be 95 percent certain that the portfolio will have at least one default under the assumption of independence.

APPENDIX D: AN APPLICATION TO INVESTMENT AND SPECULATIVE GRADE DEFAULT RATES

The default studies of Moody's Investors Service provide a long time series that one can use to test the model. For simplicity, we will break the series into two categories: investment and speculative grade issuers. The long time series enables us to set the PD equal to the average of the default rate time series and to calibrate the correlation assumption to the observed standard deviation of the default rate time series.

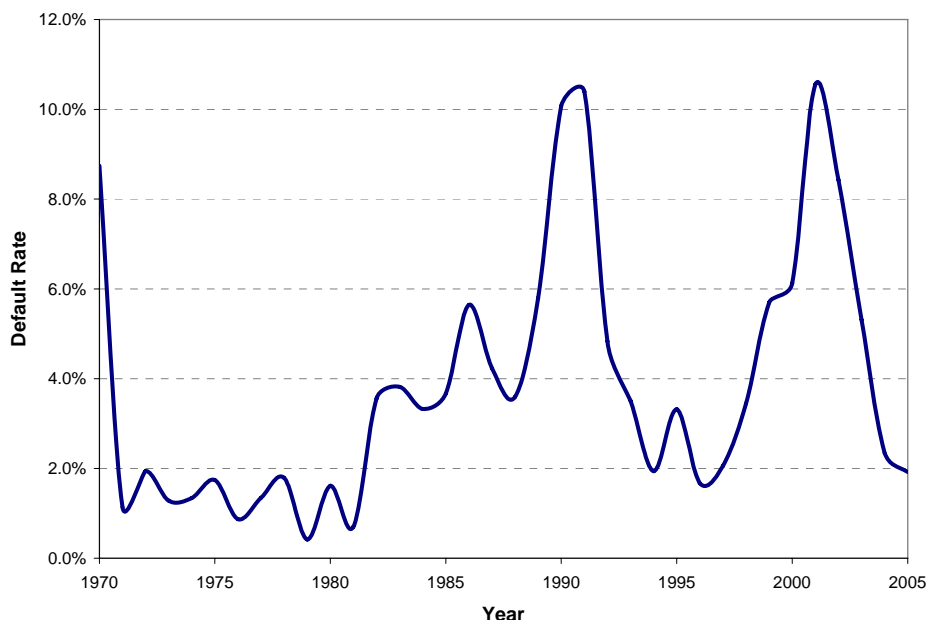


FIGURE 5 The Default Rate of Speculative Grade Issuers 1970–2005

Source: Exhibit 30 of Hamilton and Varma 2006.

Figure 1 and Figure 5 present the time series of default rates for Investment and Speculative Grade issuers as rated by Moody's Investors Service for a long time series (1970–2005).²⁵ The second and fifth column of Table 6 provides statistics and quantiles regarding these two time series. Given a PD and a correlation assumption one can simulate this time series by drawing an aggregate shock for each year and an idiosyncratic shock for each issuer in each year. One then computes the time series of the default rate as well as the average, standard deviation, 25th percentile, median, 75th percentile, 90th percentile and maximum for this time series and records these values. One can repeat this exercise 10,000 times and take the average of each statistic and quantile across the simulations.

One can also take the standard deviation of these values to determine the extent of sampling variability that is in the data. For these simulations, we chose the PD to match the time series average of the realized default rate and chose the correlation parameter to match the standard deviation of the realized default rate. Our algorithm is to perform the simulations for each value of $\rho \in \{0, 0.01, 0.02, \dots, 0.2\}$ and to then chose the ρ that minimizes the square of the difference between the average of the standard deviation of the time series default rate and the standard deviation of the realized time series default rate. We present the results of such an analysis for a PD and correlation assumption of $\{0.068\%, 0.11\}$ and $\{3.839\%, 0.10\}$ for investment and speculative grade respectively.

²⁵ For both of these time series the number of rated issuers is growing. Investment grade issues grew from 738 to 3,390 and speculative grade issues grew from 286 to 1,559. The counts for the speculative grade and investment issuers are computed from Exhibit 27 and Exhibit 30 of Hamilton (2005).

This very simple model fits the general properties of the time series rather well.²⁶ The second and fifth columns present the realized statistics and quantiles for the actual the time series. The third and sixth columns present the average of these values across the 10,000 simulations. The fourth and seventh columns present the standard deviation of the values across the simulations. The realized 25th, median, 75th, 90th and 95th percentiles are within the sampling variability that one would expect given the model. The largest departure is the realized 90th percentile for investment grade which is approximately one standard deviation away from the average of the simulations. The worst default rate observations in the realized 36-year time period are lower than the average value across the simulations. The finding that the tail risk appears worse in the simulations than in the actual time series is perhaps surprising.

The Gaussian framework has been criticized for under predicting correlated extreme events. Under the risk neutral measure, it has been argued that the Gaussian framework does not price CDO tranches effectively for this reason. The implied correlation for the equity tranche tends to be lower than the implied correlation for the more senior tranches when implied correlations are quoted using the so-called *base correlation* approach (cf., Hull and White, 2006). Therefore, it is perhaps surprising that a Gaussian framework can fit the tail risk over a 36-year period reasonably well, when working in the *physical world* of actual default events rather than the *risk neutral* world of prices.

TABLE 6 Distribution of Observed versus Modeled Default Rates for Investment and Speculative Grade Issuers

	Investment Grade			Speculative Grade		
	Realized	Avg across Sims	STDev across Sims	Realized	Avg across Sims	STDev across Sims
Average	0.068%	0.068%	0.021%	3.8%	3.8%	0.49%
Std Dev	0.118%	0.118%	0.049%	2.8%	2.9%	0.61%
p25	0.000%	0.000%	0.000%	1.7%	1.8%	0.36%
Median	0.000%	0.011%	0.019%	3.3%	3.1%	0.51%
p75 th	0.102%	0.097%	0.028%	5.4%	5.2%	0.80%
p90 th	0.271%	0.204%	0.068%	8.7%	7.8%	1.41%
p95 th	0.318%	0.331%	0.134%	10.4%	10.1%	2.19%
Max	0.485%	0.526%	0.287%	10.5%	12.6%	3.40%

This table compares the distribution of the observed time series of default rates from 1970–2005 to the distribution of default rates observed across 10,000 simulations.

Figure 6 and Figure 7 present that time series of the default rates and the corresponding p-value of the observed default rates for these series using the same PD and correlation assumptions that were used in Table 6 and the methodology in Section 2. For the investment grade series the p-values exceed the 95th percentile in 1986, 1989, and 2002. The years 1989 and 2002 correspond to recessionary periods in the economy. Similarly, for the speculative grade series the p-value exceeds the 95th percentile in 1990, 1991, and 2001, which also correspond to recessionary periods.

Figure 8 and Figure 9 present the 10th, 50th, and 90th percentile for the posterior distribution of the aggregate shock for the investment grade and speculative grade default rates computed using the methodology in Section 3.2. For the speculative grade series there were large negative shocks in 1970, 1990, 1991, 2000, and 2001. These time periods correspond to recessions with the exception of 1970. During 1970, Penn Central Railroad and 25 of its affiliates defaulted, which explains the spike. For the investment grade series, the negative aggregate shocks are likely to have occurred in 1970, 1973, 1982, 1986, 1989, 2000, 2001, and 2002. All of these years correspond to recessionary periods with the exception of 1970 and 1986, with 1970 being explained by Penn Central Railroad.

²⁶ For purposes of this simple analysis, we assume that the aggregate shocks are independent between time periods. This assumption implies zero correlation between the current and prior year default rate. In the case of the investment grade default rate series we can not reject the hypothesis of zero correlation. For the speculative default rate series, however, there is substantial positive correlation (0.65). The model could be extended to reflect this feature of the data by incorporating some autocorrelation into the aggregate shock (e.g., an AR1 process).

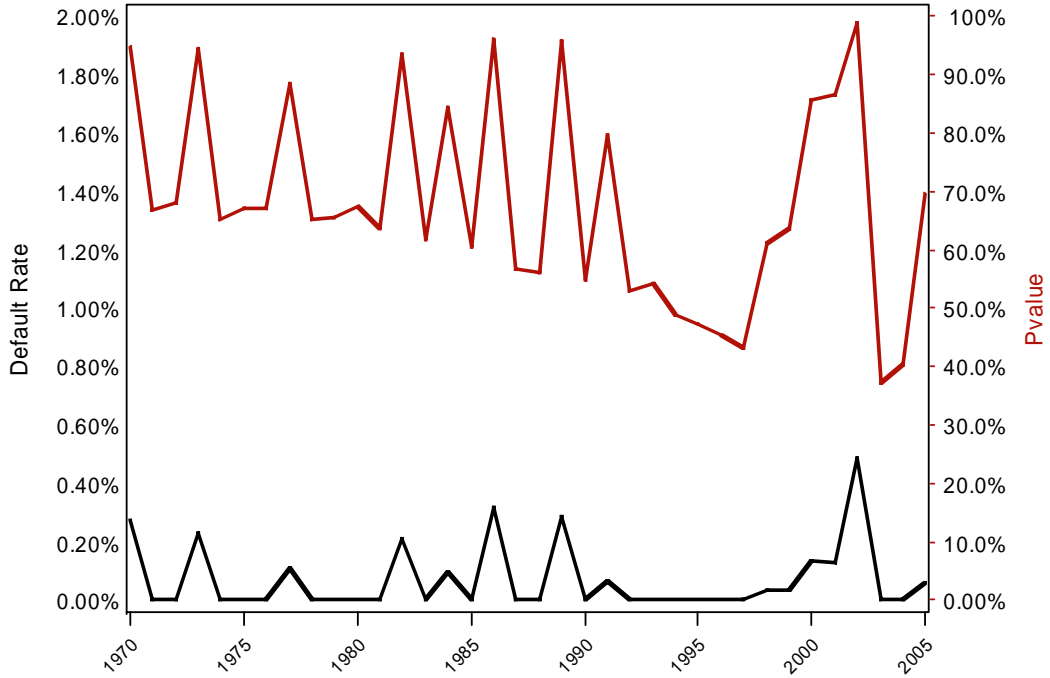


FIGURE 6 Investment Grade Default Rate and Corresponding P-Value

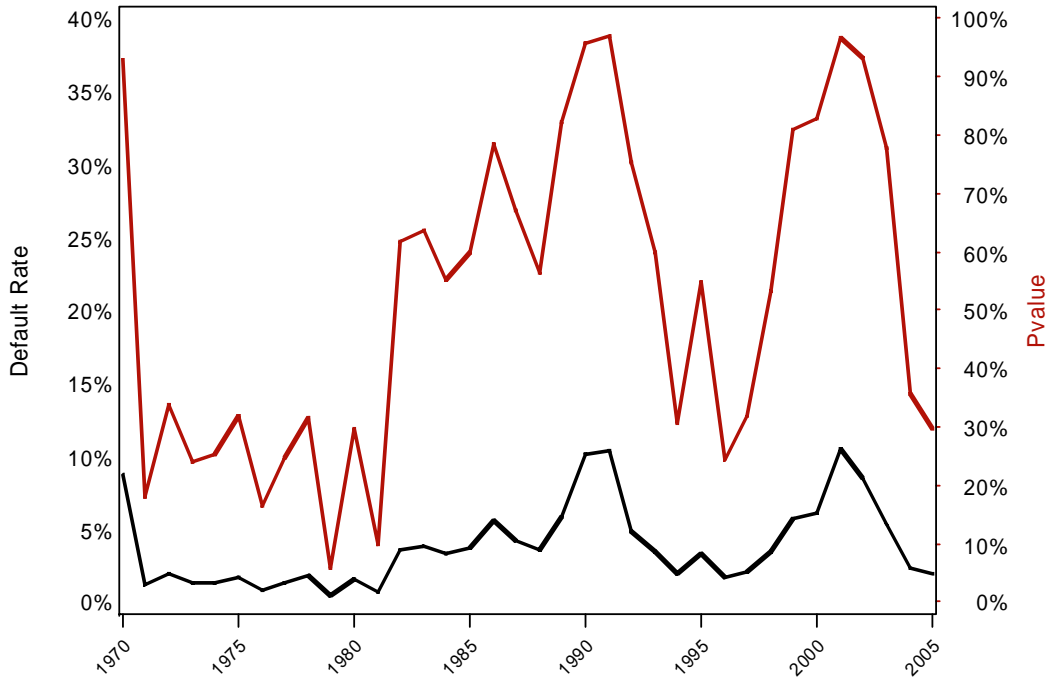


FIGURE 7 Speculative Grade Default Rate and Corresponding P-Value

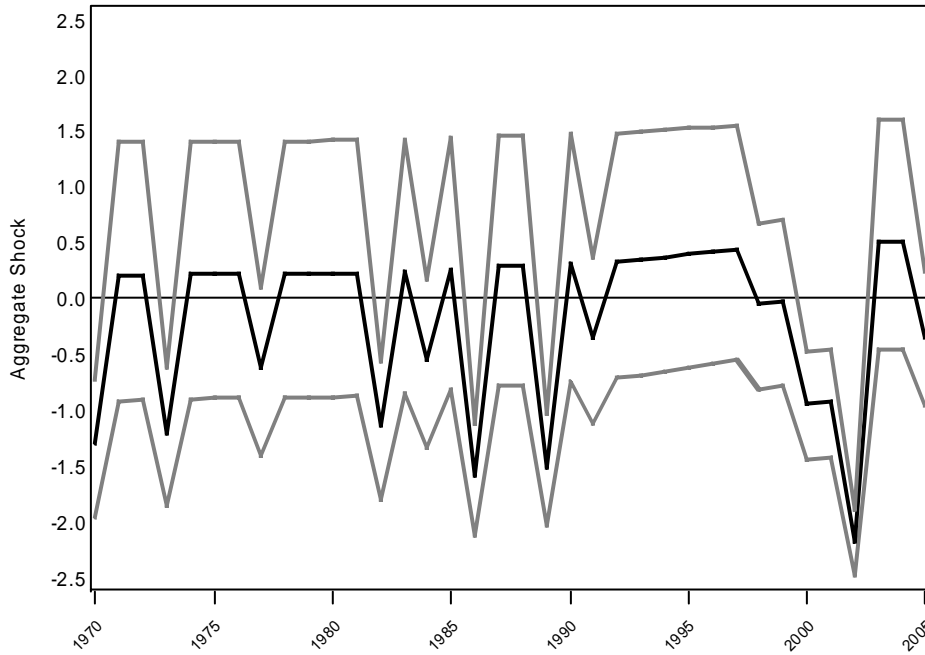


FIGURE 8 Posterior Distribution of the Aggregate Shock for Investment Grade Credits

This figure presents 10th, 50th, and 90th percentiles of the posterior distribution for the aggregate shock for investment grade issuers given the observed default rate and the assumptions that the PD and correlation are 0.068% and 11%, respectively.

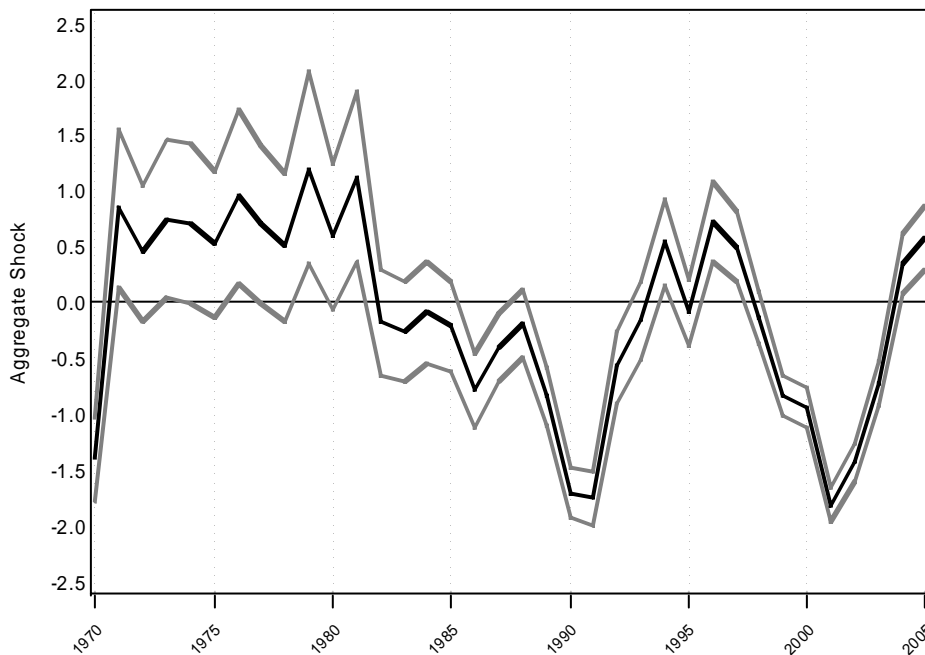


FIGURE 9 Posterior Distribution of Aggregate Shock for Speculative Grade Credits

This figure presents the 10th, 50th and 90th percentiles of the posterior distribution for the aggregate shock for speculative grade issuers given the observed default rate and the assumptions that the PD and correlation are 3.8% and 10%, respectively.

With the benefit of a 36-year time series, one can calibrate the model well and the data appears consistent with the implications of this simple model. Often banks will have a very limited time series of default rates to evaluate their model on due to changing business practices. Can this model show that a calibration is wrong? For example, suppose that one claimed that the default rate for speculative grade bonds was 1.4% and one only had access to 2004 and 2005 data to test this hypothesis. In these years the actual default rate was 2.3% and 1.9%, respectively. Figure 10 presents the p-value and the 10th, 50th and 90th percentiles for the distribution of the aggregate shock for the entire 36 year time series given a PD of 1.4% and a correlation of 10%.

Looking only at 2004 and 2005, the p-value for a PD of 1.4% is 84% and 79%, respectively. Therefore, the standard methodology would not allow you to reject the hypothesis of a 1.4% PD. Nevertheless, if the PD really were 1.4% then there would have been a negative aggregate shock in both of these years—the 90th percentile for the distribution of the aggregate shock is less than 0 in both cases. This implication is counterfactual; these years have been characterized as expansionary in the world economies. Therefore, one could conclude that the calibration needed to be revisited even though the p-value of the realized default rate was less than 95% in both years according to the traditional approach.

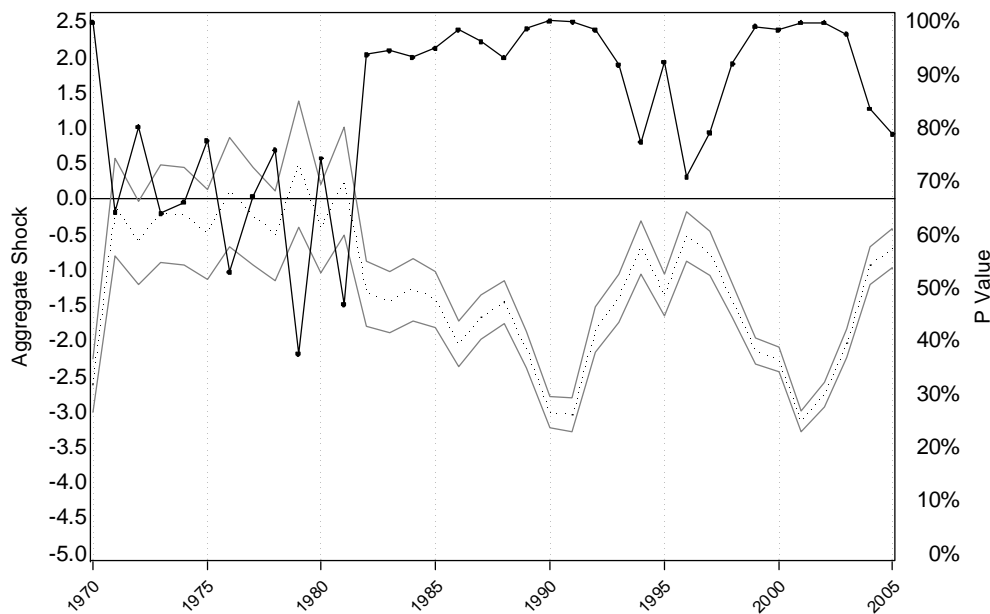


FIGURE 10 P-Value and Posterior Distribution for Speculative Grade Default Rates Given PD = 1.4%

The grey lines in this figure presents the 10th, 50th and 90th percentiles of the posterior distribution for the aggregate shock for speculative grade issuers given the observed default rate and the assumptions that the PD and correlation are 1.4% and 10%, respectively. The black line at the top, which is read off the right axis, provides the p-value of the observed default rate in each year.

REFERENCES

1. Basel Committee on Banking Supervision (2004). “International Convergence of Capital Measurement and Capital Standards (*A Revised Framework*),” Bank for International Settlements.
2. Basel Committee on Banking Supervision (2005). “An Explanatory Note on the Basel II IRB Risk Weight Functions,” Bank for International Settlements.
3. Cantor, Richard, and Eric Falkenstein (2001). “Testing for Rating Consistency in Annual Default Rates,” Moody’s Investors Service.
4. Christensen, Jens. H. E., Ernst Hanson, and David Lando (2004). “Confidence Sets for Continuous-time Rating Transition Probabilities.” *Journal of Banking and Finance* 28: 2575–2602.
5. Duffie, Darrell, Andreas Eckner, Guillaume Horel and Leandro Saita (2006). “Frailty Correlated Default,” Stanford University.
6. Dwyer, Douglas, and Roger Stein (2005). *Technical Document on EDF RiskCalc v3.1 Methodology*. Moody’s KMV.
7. Dwyer, Douglas (2005). “Examples of Overfitting Encountered When Building Private Firm Default Prediction Models.” Moody’s KMV.
8. Dwyer, Douglas and Roger M. Stein (2006). “Inferring the Default Rate In a Population By Comparing Two Incomplete Default Databases.” *Journal of Banking and Finance* 30: 797–810.
9. Embrechts, P., Alexander McNeil, and Daniel Straumann (1999). “Correlation and Dependence in Risk Management: Properties and Pitfalls.”
10. Gelman, A., John B. Carlin, Hal S. Stern and Donald B. Rubin (1995). “Bayesian Data Analysis.” New York, Chapman & Hall/CRC.
11. Gordy, Michael. (2002). “A Risk-Factor Model Foundation for Ratings-Based Bank Capital Rules.” Washington DC, Board of Governors of the Federal Reserve Board.
12. Greene, William H., *Econometric Analysis*. Fourth Edition, Upper Saddle River, NJ: Prentice Hall, 2000.
13. Hamilton, David and Praveen Varma. (2006). “Default and Recovery Rates of Corporate Bond Issuers,” 1920–2005, Moody’s Investors Service.
14. Robert V. Hogg and Allen T. Craig. (1978). *Introduction to Mathematical Statistics*, Fourth Edition, Macmillan.
15. Hull, John and Alan White. (2006). “Valuing Credit Derivatives Using an Implied Copula Approach.”
16. Jakubik, P. (2006). “Does Credit Risk Vary with Economic Cycles?” The Case of Finland, Charles University in Prague.
17. Kurbat, Matthew. and Irina Korablev (2002). “Methodology for Testing the Level of the EDFTM Credit Measure,” Moody’s KMV.
18. Maddala, G.S. *Limited Dependent and Qualitative Variables in Econometrics*, Econometric Society Monographs: Cambridge University Press, 1983.
19. Pluto, Katja. and Dirk. Tasche. (2005). “Estimating Probabilities for Low Default Portfolios,” Deutsche Bundesbank.
20. Rösch, D. (2003). “Correlations and Business Cycles of Credit Risk: Evidence from Bankruptcies in Germany,” University of Regensburg.
21. Stein, R. (2006). “Are the Probabilities Right? A First Approximation to the Lower Bound on the Number of Observations Required to Test for Default Rate Accuracy.” *Journal of Investment Management*, Vol 4., No 2.
22. Tasche, D. (2003). “A Traffic Lights Approach to PD Validation,” Deutsche Bundesbank.
23. Tasche, Dirk (2005), “Rating and Probability of Default Validation,” in *Studies on the Validation of Internal Rating Systems*, Working Paper No. 14, Basel Committee on Banking Supervision, Bank for International Settlements.
24. Vasicek, Oldrich. (1987). “Probability of Loss on Loan Portfolio.” Moody’s KMV.
25. Vasicek, Oldrich. (1991). “Limiting Loan Loss Probability Distribution.” Moody’s KMV.