

Coppens, François; González, Fernando; Winkler, Gerhard

Research Report

The performance of credit rating systems in the assessment of collateral used in Eurosystem monetary policy operations

ECB Occasional Paper, No. 65

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Coppens, François; González, Fernando; Winkler, Gerhard (2007) : The performance of credit rating systems in the assessment of collateral used in Eurosystem monetary policy operations, ECB Occasional Paper, No. 65, European Central Bank (ECB), Frankfurt a. M.

This Version is available at:

<https://hdl.handle.net/10419/154518>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



EUROPEAN CENTRAL BANK

EUROSYSTEM

OCCASIONAL PAPER SERIES

NO 65 / JULY 2007

THE PERFORMANCE OF CREDIT RATING SYSTEMS IN THE ASSESSMENT OF COLLATERAL USED IN EUROSYSTEM MONETARY POLICY OPERATIONS

by François Coppens,
Fernando González
and Gerhard Winkler



EUROPEAN CENTRAL BANK

EUROSYSTEM



OCCASIONAL PAPER SERIES

NO 65 / JULY 2007

THE PERFORMANCE OF CREDIT RATING SYSTEMS IN THE ASSESSMENT OF COLLATERAL USED IN EUROSYSTEM MONETARY POLICY OPERATIONS¹

by François Coppens²,
Fernando González³
and Gerhard Winkler⁴

In 2007 all ECB
publications
feature a motif
taken from the
€20 banknote.

This paper can be downloaded without charge from
<http://www.ecb.int> or from the Social Science Research Network
electronic library at http://ssrn.com/abstract_id=977356.

¹ This paper contains background material produced by the authors for the Eurosystem Task Force on the Eurosystem Credit Assessment Framework (ECAF). The ECAF comprises the techniques and rules that establish the Eurosystem requirement of "high credit standards" for all eligible collateral in the Single List of collateral for Eurosystem monetary policy operations. One of its key aims is to maintain a minimum level of comparability between the different credit systems that participate in the credit assessment of collateral. The authors would like to thank the members of the ECAF Task Force and of the Working Group on Risk Assessment and an anonymous ECB referee for their helpful comments on earlier drafts of this paper.

² National Bank of Belgium, boulevard de Berlaimont 14, BE-1000 Brussels, Belgium.

³ European Central Bank, Kaiserstrasse 29, D-60311 Frankfurt am Main, Germany.

⁴ Oesterreichische Nationalbank, Otto Wagner Platz 3, A-1090 Vienna, Austria.

© European Central Bank, 2007

Address

Kaiserstrasse 29
60311 Frankfurt am Main
Germany

Postal address

Postfach 16 03 19
60066 Frankfurt am Main
Germany

Telephone

+49 69 1344 0

Website

<http://www.ecb.int>

Fax

+49 69 1344 6000

Telex

411 144 ecb d

All rights reserved. Any reproduction, publication or reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the author(s).

The views expressed in this paper do not necessarily reflect those of the European Central Bank.

ISSN 1607-1484 (print)
ISSN 1725-6534 (online)

CONTENTS

ABSTRACT	4
1 INTRODUCTION	5
2 A STATISTICAL FRAMEWORK – MODELLING DEFAULTS USING A BINOMIAL DISTRIBUTION	9
3 THE PROBABILITY OF DEFAULT ASSOCIATED WITH A SINGLE “A” RATING	12
4 CHECKING THE SIGNIFICANCE OF DEVIATIONS OF THE REALISED DEFAULT RATE FROM THE FORECAST PROBABILITY OF DEFAULT	15
4.1 Two possible backtesting strategies	21
4.2 The traffic light approach, a simplified backtesting mechanism	27
5 SUMMARY AND CONCLUSIONS	29
ANNEX HISTORICAL DATA ON MOODY’S A-GRADE	31
EUROPEAN CENTRAL BANK OCCASIONAL PAPER SERIES	36

CONTENTS

ABSTRACT

The aims of this paper are twofold: first, we attempt to express the threshold of a single “A” rating as issued by major international rating agencies in terms of annualised probabilities of default. We use data from Standard & Poor’s and Moody’s publicly available rating histories to construct confidence intervals for the level of probability of default to be associated with the single “A” rating. The focus on the single “A” rating level is not accidental, as this is the credit quality level at which the Eurosystem considers financial assets to be eligible collateral for its monetary policy operations. The second aim is to review various existing validation models for the probability of default which enable the analyst to check the ability of credit assessment systems to forecast future default events. Within this context the paper proposes a simple mechanism for the comparison of the performance of major rating agencies and that of other credit assessment systems, such as the internal ratings-based systems of commercial banks under the Basel II regime. This is done to provide a simple validation yardstick to help in the monitoring of the performance of the different credit assessment systems participating in the assessment of eligible collateral underlying Eurosystem monetary policy operations. Contrary to the widely used confidence interval approach, our proposal, based on an interpretation of p-values as frequencies, guarantees a convergence to an ex ante fixed probability of default (PD) value. Given the general characteristics of the problem considered, we consider this simple mechanism to also be applicable in other contexts.

Keywords: credit risk, rating, probability of default (PD), performance checking, backtesting.

JEL classification: G20, G28, C49.

I INTRODUCTION

To ensure the Eurosystem's requirement of high credit standards for all eligible collateral, the ECB's Governing Council has established the so-called Eurosystem Credit Assessment Framework (ECAF) (see European Central Bank 2007). The ECAF comprises the techniques and rules which establish and ensure the Eurosystem's requirement of high credit standards for all eligible collateral. Within this framework, the Eurosystem has specified its understanding of high credit standards as a minimum credit quality equivalent to a rating of "A",¹ as issued by the major international rating agencies.

In its assessment of the credit quality of collateral, the ECB has always taken into account, *inter alia*, available ratings by major international rating agencies. However, relying solely on rating agencies would not adequately cover all types of borrowers and collateral assets. Hence the ECAF makes use not only of ratings from (major) external credit assessment institutions, but also other credit quality assessment sources, including the in-house credit assessment systems of national central banks,² the internal ratings-based systems of counterparties and third-party rating tools (European Central Bank, 2007).

This paper focuses on two objectives. First, it analyses the assignation of probabilities of default to letter rating grades as employed by major international rating agencies and, second, it reviews various existing validation methods for the probability of default. This is done from the perspective of a central bank or system of central banks (e.g. the Eurosystem) in the special context of its conduct of monetary policy operations in which adequate collateral with "high credit standards" is required. In this context, "high credit standards" for eligible collateral are ensured by requiring a minimum rating or its quantitative equivalent in the form of an assigned annual probability of default. Once an annual probability of default at the required rating level has been assigned, it is

necessary to assess whether the estimated probability of default issued by the various credit assessment systems conform to the required level. The methods we review and propose throughout this paper for these purposes are deemed to be valid and applicable not only in our specific case but also in more general cases.

The first aim of the paper relates to the assignation of probabilities of default to certain rating grades of external rating agencies. Ratings issued by major international rating agencies often act as a benchmark for other credit assessment sources whose credit assessments are used for comparison. Commercial banks have a natural interest in the subject because probabilities of default are inputs in the pricing of all sorts of risk assets, such as bonds, loans and credit derivatives (see e.g. Cantor et al. (1997), Elton et al. (2004), and Hull et al. (2004)). Furthermore, it is of crucial importance for regulators as well. In the "standardised approach" of the New Basel Capital Accord, credit assessments from external credit assessment institutions can be used for the calculation of the required regulatory capital (Basel Committee on Banking Supervision (2005a)). Therefore, regulators must have a clear understanding of the default rates to be expected (i.e. probability of default) for specific rating grades (Blochwitz and Hohl (2001)). Finally, it is also essential for central banks to clarify what specific rating grades mean in terms of probabilities of default since most central banks also partly rely on ratings from external credit institutions for establishing eligible collateral in their monetary operations. Although it is well known that agency ratings may to some extent also be dependent on the expected severity of loss in the event of default

1 Note that we focus on the broad category "A" throughout this paper. The "A"-grade comprises three sub-categories (named A+, A, and A- in the case of Standard & Poor's, and A1, A2, and A3 in the case of Moody's). However, we do not differentiate between them or look at them separately, as the credit threshold of the Eurosystem was also defined using the broad category.

2 At the time of publication of this paper, only the national central banks of Austria, France, Germany and Spain possessed an in-house credit assessment system.

(e.g. Cantor and Falkenstein (2001)), a consistent and clear assignment of probabilities of default to rating grades should be theoretically possible because we infer from the rating agencies' own definitions of the meanings of their ratings that their prime purpose is to reflect default probability (Crouhy et al. (2001)). This especially holds for "issuer-specific credit ratings", which are the main concern of this paper. Hence a clear relation between probabilities of default and rating grades definitely exists, and it has been the subject of several studies (Cantor and Falkenstein (2001), Blochwitz and Hohl (2001), Tiomo (2004), Jafry and Schuermann (2004) and Christensen et al. (2004)). It thus seems justifiable for the purposes of this paper to follow the definition of a rating given by Krahnen et al. (2001) and regard agency ratings as "the mapping of the probability of default into a discrete number of quality classes, or rating categories" (Krahnen et al. (2001)).

We thus attempt to express the threshold of a single "A" rating by means of probabilities of default. We focus on the single "A" rating level because this is the level at which the ECB Governing Council has explicitly defined its understanding of "high credit standards" for eligible collateral in the ECB monetary policy operations. Hence, in the empirical application of our methods, which we regard as applicable to the general problem of assigning probabilities of default to any rating grades, we will restrict ourselves to a single illustrative case, the "A" rating grade. Drawing on the above-mentioned earlier works of Blochwitz and Hohl (2001), Tiomo (2004) and Jafry and Schuermann (2004), we analyse historical default rates published by the two rating agencies Standard & Poor's and Moody's. However, as default is a rare event, especially for entities rated "A" or better, the data on historically observed default frequencies shows a high degree of volatility, and probability of default estimates could be very imprecise. This may be due to country-specific and industry-specific idiosyncrasies which might affect rating migration dynamics (Nickel et al. (2000)). Furthermore,

macroeconomic shocks can generally also influence the volatility of default rates, as documented by Cantor and Falkenstein (2001). As discussed by Cantor (2001), Fons (2002) and Cantor and Mann (2003), however, agency ratings are said to be more stable in this respect because they aim to measure default risk over long investment horizons and apply a "through the cycle" rating philosophy (Crouhy et al. (2001) and Heitfield (2005)). Based on these insights we derive an ex ante benchmark for the single "A" rating level. We use data of Standard & Poor's and Moody's publicly available rating histories (Standard & Poor's (2005), Moody's (2005)) to construct confidence intervals for the level of probability of default to be associated with a single "A" rating grade. This results in one of the main contributions of our work, i.e. the statistical deduction of an ex ante benchmark of a single "A" rating grade in terms of probability of default.

The second aim of this paper is to explore validation mechanisms for the estimates of probability of default issued by the different rating sources. In doing so, it presents a simple testing procedure that verifies the quality of probability of default estimates. In a quantitative validation framework the comparison of performance could be based mainly on two criteria: the discriminatory power or the quality of calibration of the output of the different credit assessment systems under comparison. Whereas the "discriminatory power" refers to the ability of a rating model to differentiate between good and bad cases, calibration refers to the concrete assignment of default probabilities, more precisely to the degree to which the default probabilities predicted by the rating model match the default rates actually realised. Assessing the calibration of a rating model generally relies on backtesting procedures.³ In this paper we focus on the

³ To conduct a backtesting examination of a rating source the basic data required is the estimate of probability of default for a rating grade over a specified time horizon (generally 12 months), the number of rated entities assigned to the rating grade under consideration and the realised default status of those entities after the specified time horizon has elapsed (i.e. generally 12 months after the rating was assigned).

quality of the calibration of the rating source and not on its discriminatory power.⁴

Analysing the significance of deviations between the estimated default probability and the realised default rate in a backtesting exercise is not a trivial task. Realised default rates are subject to statistical fluctuations that could impede a straight forward assessment of how well a rating system estimates probabilities of default. This is mainly due to constraints on the number of observations available owing to the scarcity of default events and the fact that default events may not be independent but show some degree of correlation. Non-zero default correlations have the effect of amplifying variations in historically observed default rates which would normally prompt the analyst to widen the tolerance of deviations between the estimated average of the probabilities of default of all obligors in a certain pool and the realised default rate observed for that pool. In this sense, two approaches can be considered in the derivation of tests of deviation significance: tests assuming uncorrelated default events and tests assuming default correlation.

There is a growing literature on probability of default validation via backtesting (e.g. Cantor and Falkenstein (2001), Blochwitz et al. (2003), Tasche (2003), Rauhmeier (2006)). This work has been prompted mainly by the need of banking regulators to have validation frameworks in place to face the certification challenges of the new capital requirement rules under Basel II. Despite this extensive literature, there is also general acceptance of the principle that statistical tests alone would not be sufficient to adequately validate a rating system (Basel Committee on Banking Supervision (2005b)). As mentioned earlier, this is due to scarcity of data and the existence of a default correlation that can distort the results of a test. For example, a calibration test that assumes independence of default events would normally be very conservative in the presence of correlation in defaults. Such a test could send wrong messages for an otherwise well calibrated rating system. However, and given these caveats, validation

by means of backtesting is still considered valuable for detecting problems in rating systems.

We briefly review various existing statistical tests that assume either independence or correlation of defaults (cf. Brown et al. (2001), Cantor and Falkenstein (2001), Spiegelhalter (1986), Hosmer and Lemeshow (2000), Tasche (2003)). In doing so, we take a closer look at the binomial model of defaults that underpins a large number of tests proposed in the literature. Like any other model, the binomial model has its limitations. We pay attention to the discreteness of the binomial distribution and discuss the consequences of approximation, thereby accounting for recent developments in statistics literature regarding the construction of confidence intervals for binomially distributed random variables (for an overview see Vollset (1993), Agresti and Coull (1998), Agresti and Caffo (2000), Reiczigel (2004) and Cai (2005)).

We conclude the paper by presenting a simple hypothesis testing procedure to verify the quality of probability of default estimates that builds on the idea of a “traffic light approach” as discussed in, for example, Blochwitz and Hohl (2001) and Tiomo (2004). A binomial distribution of independent defaults is assumed in accordance with the literature on validation. Our model appears to be conservative and thus risk averse. Our hypothesis testing procedure focuses on the interpretation of p -values as frequencies, which, contrary to an approach based on confidence intervals, guarantees a long-run convergence to the probability of default of a specified or given level of probability of default that we call the benchmark level. The approach we propose is flexible and takes into account the number of objects rated by the specific rating system. We regard this approach as an early warning system that could identify problems of calibration in a rating

⁴ For an exposition of discriminatory power measures in the context of the assessment of performance of a rating source see, for example, Tasche (2006).

system, although we acknowledge that, given the fact that default correlation is not taken into account in the testing procedure, false alarms could be given for otherwise well-calibrated systems. Eventually, we are able to demonstrate that our proposed “traffic light approach” is compliant with the mapping procedure of external credit assessment institutions foreseen in the New Basel Accord (Basel Committee on Banking Supervision (2005a)).

The paper is organised as follows. In Section 2 the statistical framework forming the basis of a default generating process using binomial distribution is briefly reviewed. In Section 3 we derive the probability of default to be associated with a single “A” rating of a major rating agency. Section 4 discusses several approaches to checking whether the performance of a certain rating source is equivalent to a single “A” rating or its equivalent in terms of probability of default as determined in Section 3. This is done by means of their realised default frequencies. The section also contains our proposal for a simplified performance checking mechanism that is in line with the treatment of external credit assessment institutions in the New Basel Accord. Section 5 concludes the paper.

2 A STATISTICAL FRAMEWORK – MODELLING DEFAULTS USING A BINOMIAL DISTRIBUTION

The probability of default itself is unobservable because the default event is stochastic. The only quantity observable, and hence measurable, is the empirical default frequency. In search of the meaning of a single “A” rating in terms of a one year probability of default we will thus have to make use of a theoretical model that rests on certain assumptions about the rules governing default processes. As is common practice in credit risk modelling, we follow the “cohort method” (in contrast to the “duration approach”, see Lando and Skoedeborg (2002)) throughout this paper and, furthermore, assume that defaults can be modelled using a binomial distribution (Nickel et al. (2000), Blochwitz and Hohl (2001), Tiomo (2003), Jafry and Schuermann (2004)). The quality of each model’s results in terms of their empirical significance depends on the adequacy of the model’s underlying assumptions. As such, this section briefly discusses the binomial distribution and analyses the impact of a violation of the assumptions underlying the binomial model.⁵ It is argued that postulating a binomial model reflects a risk-averse point of view.⁶

We decided to follow the cohort method as the major rating agencies document the evolution of their rated entities over time on the basis of “static pools” (Standard & Poor’s 2005, Moody’s 2005). A static pool consists of N_Y rated entities with the same rating grade at the beginning of a year Y . In our case N_Y denotes the number of entities rated “A” at the beginning of year Y . The cohort method simply records the number of entities D_Y that have defaulted by the year end out of the initial N_Y rated entities (Nickel et al. (2000), Jafry and Schuermann (2004)).

It is assumed that D_Y , the number of defaults in the static pool of a particular year Y , is binomially distributed with a “success probability” p and a number of events N_Y (in notational form: $D_Y \approx B(N_Y; p)$). From this

assumption it follows that each individual (“A”-rated) entity has the same (one year) probability of default “ p ” under the assumed binomial distribution. Moreover the default of one company has no influence on the (one year) defaulting of the other companies, i.e. the (one year) default events are independent. The number of defaults D_Y can take on any value from the set $\{0, 1, 2, \dots, N_Y\}$. Each value of this set has a probability of occurrence determined by the probability density function of the binomial distribution which, under the assumptions of constant p and independent trials, can be shown to be equal to:

$$b(n_Y; N_Y; p) = P(D_Y = n_Y) = \binom{N_Y}{n_Y} p^{n_Y} (1-p)^{N_Y-n_Y} \quad (1)$$

The mean and the variance of the binomial distribution are given by

$$\begin{aligned} \mu_{D_Y} &= N_Y p \\ \sigma_{D_Y}^2 &= N_Y p(1-p) \end{aligned} \quad (2)$$

As indicated above, a clear distinction has to be made between the “probability of default” (PD) (i.e. the parameter p in formula (1)) and the “default frequency”. While the probability of default is the fixed (and unobservable) parameter “ p ” of the binomial distribution, the default frequency is the observed number of defaults in a binomial experiment, divided by the number of trials $\left(df_Y = \frac{n_Y}{N_Y}\right)$. This default

frequency varies from one experiment to another, even when the parameters p and N_Y stay the same. It can take on values from the set $df_Y \in \left\{\frac{0}{N_Y}, \frac{1}{N_Y}, \frac{2}{N_Y}, \dots, 1\right\}$. The value observed for

5 For a more detailed treatment of binomial distribution see e.g. Rohatgi (1984), and Moore and McCabe (1999).

6 An alternative distribution for default processes is the “Poisson distribution”. This distribution has some benefits, such as the fact that it can be defined by only one parameter and that it belongs to the exponential family of distributions which easily allow uniformly most powerful (UMP) one and two-sided tests to be conducted in accordance with the Neyman-Pearson theorem (see the Fisher-Behrens problem). However, in this paper we have opted to follow the mainstream literature on validation of credit systems which rely on binomial distribution to define the default generating process.

one particular experiment is the observed default frequency for that experiment.

The mean and variance of the default frequency can be derived from formula (1):

$$\begin{aligned}\mu_{df_y} &= p \\ \sigma_{df_y}^2 &= \frac{p(1-p)}{N_y}\end{aligned}\quad (2')$$

The probability density function can be derived from (1) by setting $f_y = \frac{n_y}{N_y}$:

$$P(df_y = f_y) = \binom{N_y}{f_y N_y} p^{f_y N_y} (1-p)^{(1-f_y)N_y} \quad (3)$$

As $f_y \in \left\{ \frac{0}{N_y}, \frac{1}{N_y}, \frac{2}{N_y}, \dots, 1 \right\}$ this distribution is discrete.

THE BINOMIAL DISTRIBUTION ASSUMPTIONS

It is of crucial importance to note that formula (1) is derived under two assumptions. First, the (one year) default probability should be the same for every “A”-rated company. Secondly, the “A”-rated companies should be independent with respect to the (one year) default event. This means that the default of one company in one year should not influence the default of another “A”-rated company within the same year.

THE CONSTANT “p”

It may be questioned whether the assumption of a homogeneous default probability for all “A”-rated companies is fulfilled in practice (e.g. Blochwitz and Hohl (2001), Tiomo (2004), Hui et al. (2005), Basel Committee on Banking Supervision (2005b)). The distribution of defaults would then not be strictly binomial. Based on assumptions about the distribution of probability of defaults within rating grades, Blochwitz and Hohl (2001) and Tiomo (2004) use Monte Carlo simulations to study the impact of heterogeneous probabilities of default on confidence intervals.

The impact of a violation of the assumption of a uniform probability of default across all

entities with the same rating may, however, also be modelled using “mixed binomial distribution”, of which “Lexian distribution” is a special case. Lexian distribution considers a mixture of “binomial subsets”, each subset having its own PD. The PDs can be different between subsets. The mean and variance of the Lexian variable x , which is the number of defaults among n companies, are given by⁷

$$\begin{aligned}\mu_x &= n\bar{p}, \\ \sigma_x^2 &= n\bar{p}(1-\bar{p}) + n(n-1)\text{var}(p)\end{aligned}\quad (4)$$

Where $\bar{p}$ is the average value of all the (distinct) PDs and $\text{var}(p)$ is the variance of these PDs.

Consequently, if a mixed binomial variable is treated as a pure binomial variable, its mean, the average probability of default would still be correct, whereas the variance would be underestimated when the “binomial estimator” $np(1-p)$ is used (see the additional term in (4)). The mean and the variance will be used to construct confidence intervals. An underestimated variance will lead to narrower confidence intervals for the (average) probability of default and thus to lower thresholds. Within the context of this paper, lower thresholds imply a risk-averse approach.

INDEPENDENT TRIALS

Several methods for modelling default correlation have been proposed in literature (e.g. Gordy (1998), Nagpal and Bahar (2001), Servigny and Renault (2002), Blochwitz, When and Hohl (2003, 2005) and Hamerle, Liebig and Rösch (2003)). They all point to the difficulties of measuring correlation.

Although default correlations are low for sufficiently high levels of credit quality such as a single “A” rating, they could be an important factor in performance testing for lower rating grades. Over the period 1981-2005 Standard and Poor’s historical default experience (see Table 1) shows that, with the exception of 2001,

⁷ See e.g. Johnson (1969)

not more than one company defaulted per year, a fact which indicates that correlation cannot be very high. Secondly, even if we assumed that two firms were highly correlated and one defaulted, the other one will most likely not default in the same year, but only after a certain lag! Given that the primary interest is in an annual testing framework, the possibility of intertemporal default patterns beyond the one year period is of no interest. Finally, from a risk management point of view, providing that the credit quality of the pool of obligors is high (e.g. single “A” rating or above), it could be seen as adequate to assume that there is no default correlation, because not accounting for correlation leads to confidence intervals that are more conservative.⁸ Empirical evidence for these arguments is provided by Nickel et al. (2000). Later on we will relax this assumption when presenting for demonstration purposes a calibration test accounting for default correlation.

8 As in the case of heterogeneous PDs, this is due to the increased variance when correlation is positive. Consider, for example, the case where the static pool can be divided into two subsets. Within each subset issuers are independent, but between subsets they are positively correlated. The number of defaults in the whole pool is then a sum of two (correlated) binomials. The total variance is given by $\frac{N}{2}p(1-p) + \frac{N}{2}p(1-p) + 2\sigma_{12}$, which is again higher than the “binomial variance”.



3 THE PROBABILITY OF DEFAULT ASSOCIATED WITH A SINGLE “A” RATING

In this section we derive a probability of default that could be assigned to a single “A” rating. We are interested in this rating level because this is the minimum level at which the Eurosystem has decided to accept financial assets as eligible collateral for its monetary policy operations. The derivation could easily be followed to compute the probability of default of other rating levels.

Table 1 shows data on defaults for issuers rated “A” by Standard & Poor’s (the corresponding table for Moody’s is given in Annex 1). The first column lists the year, the second shows the number of “A” rated issuers for that year. The column “Default frequency” is the observed one-year default frequency among these issuers. The last column gives the average default

frequency over the “available years” (e.g. the average over the period 1981-1984 was 0.04%).

The average one-year default frequency over the whole observation period spanning from 1981 to 2004 was 0.04%, the standard deviation of the annual default rates was 0.07%.

The maximum likelihood estimator for the parameter p of a binomial distribution is the observed frequency of success. Table 1 thus gives for each year between 1981 and 2004 a maximum likelihood estimate for the probability of default of companies rated “A” by S&P, i.e. 24 (different) estimates.

One way to combine the information contained in these 24 estimates is to apply the central limit theorem to the arithmetic average of the default frequency over the period 1981-2004

Table 1 One-year default frequency within Standard and Poor’s A-rated class

Year	Number of issuers	Default frequency (%)	Average (1981-YYYY) (%)
1981	494	0.00	0.00
1982	487	0.21	0.11
1983	466	0.00	0.07
1984	471	0.00	0.05
1985	510	0.00	0.04
1986	559	0.18	0.07
1987	514	0.00	0.06
1988	507	0.00	0.05
1989	561	0.00	0.04
1990	571	0.00	0.04
1991	583	0.00	0.04
1992	651	0.00	0.03
1993	719	0.00	0.03
1994	775	0.13	0.04
1995	933	0.00	0.03
1996	1,027	0.00	0.03
1997	1,106	0.00	0.03
1998	1,116	0.00	0.03
1999	1,131	0.09	0.03
2000	1,118	0.09	0.04
2001	1,145	0.17	0.04
2002	1,176	0.09	0.04
2003	1,180	0.00	0.04
2004	1,209	0.00	0.04
Average 1981-2004		0.04	
Standard deviation 1981-2004		0.07	

Source: Standard & Poor’s, “Annual Global Corporate Default Study: Corporate defaults poised to rise in 2005”.

which is 0.04% according to Table 1. As such, it is possible to construct confidence intervals for the true mean $\mu_{\bar{x}}$ of the population around this arithmetic average. The central limit theorem states that the arithmetic average $\bar{x}$ of n independent random variables x_i , each having mean μ_i and variance σ_i^2 , is approximately normally distributed with parameters $\mu_{\bar{x}} = \frac{\sum_{i=1}^n \mu_i}{n}$ and $\sigma_{\bar{x}}^2 = \frac{\sum_{i=1}^n \sigma_i^2}{n^2}$ (see e.g. DeGroot (1989), and Billingsley (1995)). Applying this theorem to S&P's default frequencies, random variables with $\mu_i = p$ and $\sigma_i^2 = p(1-p)/N_i$, yields the result that the arithmetic average of S&P's default frequencies is approximately normal with mean $\mu_{\bar{x}} = \frac{\sum_{i=1}^n p}{n} = p$ and variance $\sigma_{\bar{x}}^2 = \frac{\sum_{i=1}^n p(1-p)}{n^2}$. If the probability of default “ p ” is not constant over the years then a confidence interval for the average probability of default is obtained. In that case the estimated benchmark would be based on the average probability of default. After estimating p and $\sigma_{\bar{x}}^2$ from S&P data ($\hat{p} = 0.04\%$, $\hat{\sigma}_{\bar{x}} = 0.0155\%$, for “A” and $\hat{p} = 0.27\%$, $\hat{\sigma}_{\bar{x}} = 0.0496\%$ for “BBB”), confidence intervals for the mean, i.e. the default probability p , can be constructed. These confidence intervals are given in Table 2 for S&P's rating grades “A” and “BBB”. Similar estimates can be derived for Moody's data using the same approach. The confidence intervals for a single “A” rating from Moody's have lower limits than those shown for S&P in Table 2. This is due to the lower mean realised default frequency recorded in Moody's ratings. However, in the next paragraph it will be shown that Moody's performance does not differ significantly from that of S&P for the single “A” rating grade.

A similar result is obtained when the observations for the 24 years are “pooled”. Pooling is based on the fact that the sum of independent binomial variables with the same p is again binomial with parameters $\sum D_i \approx B(\sum N_i; p)$ (see e.g. DeGroot (1989)). Applying this theorem to the 24 years of S&P data (and assuming independence) it can be seen that eight defaults are observed among 19,009 issuers (i.e. the sum of all issuers rated

Table 2 Confidence intervals for the $\mu_{\bar{x}}$ of S&P's “A” compared to “BBB”

(percentages)		
Confidence level	Lower	Upper
S&P A		
95.0	0.01	0.07
99.0	0.00	0.08
99.5	0.00	0.09
99.9	0.00	0.10
S&P BBB		
95.0	0.17	0.38
99.0	0.13	0.41
99.5	0.12	0.43
99.9	0.09	0.46

single “A” over the 1981-2004 period). This yields an estimate for p of 0.04% and a binomial variance of 0.015%, similar to the estimates based on the central limit theorem.

The necessary condition for the application of the central limit theorem or for pooling is the independence of the annual binomial variables. This is hard to verify. Nevertheless, several arguments in favour of the above method can be brought forward. First, a quick analysis of the data in Table 1 shows that there are no visible signs of dependence among the default frequencies. Second, and probably the most convincing argument, the data in Table 1 confirms the findings for the confidence intervals that are found in Table 2. Indeed, the last column in Table 1 shows the average over 2, 3, ..., 24 years. As can be seen, with a few exceptions, these averages lie within the confidence intervals (see Table 2). For the exceptions it can be argued (1) that not all values have to be within the limits of the confidence intervals (in fact, for a 99% confidence interval one exception is allowed every 100 years, and for a 95% interval it is even possible to exceed the limits every 20 years) and (2) that we did not always compute 24-year averages although the central limit theorem was applied to a 24-year average. When random samples of size 23 are drawn from these 24 years of data, the arithmetic average seems to be within the limits given in Table 2. The third argument in support of our

findings is a theoretical one. In fact, a violation of the independence assumption would change nothing in the findings about the mean $\mu_{\bar{x}}$. However, the variance would no longer be correct as the covariances should be taken into account. Furthermore, dependence among the variables would no longer guarantee a normal distribution. The sum of dependent and (right) skewed distributions would no longer be symmetric (like the normal distribution) but also skewed (to the right). Assuming positive covariances would yield wider confidence intervals. Furthermore, as the resulting distribution will be skewed to the right, and as values lower than zero would not be possible, using the normal distribution as an approximation would lead to smaller confidence intervals. As such, a violation of the independence assumption implies a risk-averse result.

An additional argument can be brought forward which supports our findings: First, in the definition of the “A” grade we are actually also interested in the minimum credit quality that “A-grade” stands for. We want to know the highest value the probability of default can take to be still accepted as equivalent to “A”. Therefore we could also apply the central limit theorem to the data for Standard & Poor’s BBB. Table 2 shows that in that case the PD of a BBB rating is probably higher than 0.1%.

We can thus conclude that there is strong evidence to suggest that the *probability of default* for the binomial process that models the observed default frequencies of Standard & Poor’s “A” rating grade is between 0.00% and 0.1% (see Table 2). The average point estimate is 0.04%. For reasons mentioned above, these limits are conservative, justifying the use of values above 0.04% (but not higher than 0.1%). An additional argument for the use of a somewhat higher value for the average point estimate than 0.04% is the fact that the average observed default frequency for the last five years of Table 1 equals 0.07%.

TESTING FOR EQUALITY IN DEFAULT FREQUENCIES OF TWO RATING SOURCES AT THE SAME RATING LEVEL

The PD of a rating source is unobservable. As a consequence, a performance checking mechanism cannot be based on the PD alone. In this section it is shown that the central limit theorem could also be used to design a mechanism that is based on an average observed default frequency.⁹

Earlier on, using the central limit theorem, we found that the 24-year average of S&P’s default frequencies is normally distributed:

$$\bar{x}^{S\&P} \approx N(\mu_{\bar{x}^{S\&P}}; \sigma_{\bar{x}^{S\&P}}) \quad (5)$$

with $\mu_{\bar{x}^{S\&P}}$ and $\sigma_{\bar{x}^{S\&P}}$ estimated at 0.04% and 0.0155% respectively.

In a similar way, the average default frequency of any rating source is normally distributed:

$$\bar{x}^{rs} \approx N(\mu_{\bar{x}^{rs}}; \sigma_{\bar{x}^{rs}}) \quad (6)$$

The formulae (5) and (6) can be used to test whether the average default frequency of the rating source is at least as good as the average of the benchmark by testing the statistical hypothesis

$$H_0 : \mu_{\bar{x}^{rs}} < \mu_{\bar{x}^{S\&P}} \text{ against } H_1 : \mu_{\bar{x}^{rs}} \geq \mu_{\bar{x}^{S\&P}} \quad (7)$$

Although seemingly simple, such a performance checking mechanism has several disadvantages. First, assuming, for example, 24 years of data for the rating source, the null hypothesis cannot be rejected if the annual default frequency is 0.00% on 23 occasions and 0.96% once ($\bar{x}^{rs} = \frac{23 \times 0.00\% + 1 \times 0.96\%}{24} = 0.04\%$, p -value is 50%).

In other words, extreme values for the observed default frequencies are allowed (0.96%). Second, the performance rule is independent of the static pool size. A default frequency of 0.96% on a sample size of 10,000 represents

⁹ This is only possible when historical data are available, i.e. when a n -year average can be computed.

96 defaults, while it is only 2 defaults for a sample of 200. Third, requiring 24 years of data to compute a 24-year average is impractical. Other periods could be used (e.g. a 10-year average), but that is still impractical as 10 years of data must be available before the rating source can be backtested. Taking into account these drawbacks, two alternative performance checking mechanisms will be presented in Section 4.1.

This rule can, however, be used to test whether the average default frequencies of S&P and Moody's are significantly different. Under the null hypothesis

$$H_0 : \mu_{\bar{x}^{S\&P}} = \mu_{\bar{x}^{Moody's}} \quad (8)$$

the difference of the observed averages is normally distributed, i.e. (assuming independence)

$$\bar{x}^{S\&P} - \bar{x}^{Moody's} \approx N(0; \sqrt{\sigma_{\bar{x}^{S\&P}}^2 + \sigma_{\bar{x}^{Moody's}}^2}) \quad (9)$$

Using an estimate of the variance, the variable $\frac{\bar{x}^{S\&P} - \bar{x}^{Moody's}}{\sqrt{s_{\bar{x}^{S\&P}}^2 + s_{\bar{x}^{Moody's}}^2}}$ has a t -distribution with 46 degrees of freedom and can be used to check the hypothesis (8) against the alternative hypothesis $H_1 : \mu_{\bar{x}^{S\&P}} \neq \mu_{\bar{x}^{Moody's}}$.

Using the figures from S&P and Moody's ($\hat{p} = 0.04\%$, $\hat{\sigma}_{\bar{x}} = 0.0155\%$, for S&P's "A" and $\hat{p} = 0.02\%$, $\hat{\sigma}_{\bar{x}} = 0.0120\%$ for Moody's "A"), a value of 0.81 is observed for this t -variable. This t -statistic has an implied p -value (2-sided) of 42% so the hypothesis of equal PDs for Moody's & S&P's "A" grade cannot be rejected. In formula (9) S&P and Moody's "A" class were considered independent. Positive correlation would thus imply an even lower t -value.

PERFORMANCE CHECKING: THE DERIVATION OF A BENCHMARK FOR BACKTESTING

To allow performance checking, the assignment of PDs to rating grades alone is not enough. In fact, as can be seen from S&P data in Table 1, the observed annual default frequencies often

exceed 0.1%. This is because the PD and the (observed) default frequencies are different concepts. A performance checking mechanism should, however, be based on "observable" quantities, i.e. on the observed default frequencies of the rating source.

In order to construct such a mechanism it is assumed that the annually observed default rates of the benchmark may be modelled using a binomial distribution. The mean of this distribution, the probability of default of the benchmark, is estimated at $\hat{p} \in [0.0\%, 0.1\%]$ (with an average of 0.04%). The other binomial parameter is the number of trials N . To define the benchmark N is taken to be the average size of S&P's static pool or $N = 792$ (see Table 1). This choice may appear somewhat arbitrary because the average size over the period 2000-2004 is higher (i.e. 1,166), but so is the average observed default frequency over that period (0.07%). If the binomial parameters were based on this period, then the mean and the variance of this binomial benchmark would be higher, and so confidence limits would also be higher.

In Section 4.1 below two alternatives for the benchmark will be used:

1. A fixed upper limit of 0.1% for the benchmark probability of default.
2. A stochastic benchmark, i.e. a Binomial distribution with parameters p equal to 0.1% and N equal to 792.

4 CHECKING THE SIGNIFICANCE OF DEVIATIONS OF THE REALISED DEFAULT RATE FROM THE FORECAST PROBABILITY OF DEFAULT

As realised default rates are subject to statistical fluctuations it is necessary to develop mechanisms to show how well the rating source estimates the probability of default. This is generally done using statistical tests to check the significance of the deviation of the realised default rate from the forecast probability of default. The statistical tests would normally check the null hypothesis that “the forecast probability of default in a rating grade is correct” against the alternative hypothesis that “the forecast default probability is incorrect”.

As shown in Table 1, the stochastic nature of the default process allows for observed default frequencies that are far above the probability of default. The goal of this section is to find upper limits for the observed default frequency that are still consistent with a PD of 0.1%.

We will first briefly describe some statistical tests that can be used for this purpose. The first one is to test a realised default frequency for a rating source against a fixed upper limit for the PD, this is the “Wald test” for single proportions. The second test will assess the significance of the difference between two proportions or, in other words, two default rates that come from two different rating sources. We will then proceed to a test that considers the significance of deviations between forecast probabilities of default and realised default rates of several rating grades, the “Hosmer-Lemeshow test”. In some instances, the probability of default associated with a rating grade is considered not to be constant for all obligors in that rating grade. The “Spiegelhalter test” will assess the significance of deviations when the probability of default is assumed to vary for different obligors within the rating grade. Both the Hosmer-Lemeshow and the derived Spiegelhalter test can be seen as extensions of the Wald test. Finally, we introduce a test that accounts for correlation and show how the critical values for assessing significance in

deviations can be dramatically altered in the presence of default correlation.

THE WALD TEST FOR SINGLE PROPORTIONS

For hypothesis testing purposes, the binomial density function is often approximated by a normal density function with parameters given by (2) or (2') in Section 2 (see e.g. Cantor and Falkenstein (2001), Nickel et al. (2000)).

$$df_Y \approx N\left(p; \sqrt{\frac{p(1-p)}{N_Y}}\right) \quad (10)$$


When testing the null hypothesis H_0 : “the realised default is consistent with a specified probability of default value lower than p_0 or benchmark” against H_1 : “the realised default is higher than p_0 ”, a Z-statistic

$$Z = \frac{df - p_0}{\sqrt{\frac{df(1-df)}{N_Y}}} \quad (11)$$

can be used, which is compared to the quantiles of the standard normal distribution.

The quality of the approximation depends on the values of the parameters N_Y , the number of rated entities with the same rating grade at the beginning of a year Y , and p , the forecast probability of default (see e.g. Brown et al. 2001). A higher N_Y results in better approximations. For the purpose of this paper, N_Y is considered to be sufficiently high. The low PD values for “A” rated companies (lower than 0.1%) might be problematic since the quality of the approximation degrades when p is far away from 50%. In fact, the two parameters interact, the higher N_Y is, the further away from 50% p can be. Low values of p imply a highly skewed (to the right) binomial distribution, and since the normal distribution is symmetric the approximation becomes poor. The literature on the subject is extensive (for an overview see Vollset (1993), Agresti and Coull (1998), Newcombe (1998), Agresti and Caffo (2000), Brown et al. (2001), Reiczigel (2004), and Cai (2005)). Without going into more details, the problem is briefly explained in a graphical way.

Chart 1 The performance of the Wald interval for different values of N , and for $p = 0.1\%$ (left) and $p = 0.05\%$ (right)



In Chart 1 the performance of the Wald interval is shown for several values of N , once for $p = 0.05\%$ and once for $p = 0.10\%$. Formula (10) can then be used to compute the upper limit (df_U) of the 90% one-sided confidence interval. As the normal distribution is only an approximation for the binomial distribution, the cumulative binomial distribution for this upper limit will seldom be exactly equal to 90%, i.e. $B(df_U \times N_Y; N_Y; p) = P(D_Y \leq df_U \times N_Y) \neq 90\%$.

The zigzag line shows, for different values of N , the values for the cumulative binomial distribution in the upper limit of the Wald interval. For $p = 0.1\%$ and $N = 500$ this value seems to be close to 90%. However for $p = 0.05\%$ and $N = 500$ the coverage is far below 90%. This shows that for $p = 0.05\%$ the 90% Wald confidence interval is in fact not a 90% but only a 78% confidence interval, meaning that the Wald confidence interval is too small and that a test based on this approximation (for $p = 0.05\%$ and $N = 500$) is (too) conservative. The error is due to the approximation of the binomial distribution (discrete and asymmetric) by a normal (continuous and symmetric) one. Thus it is to be noted that, the higher the value of N , the better the approximation becomes, and that in most cases the test is conservative.¹⁰

Our final traffic light approach will be based on a statistical test for differences of proportions.

This test is also based on an approximation of the binomial distribution by a normal one. In this case, however, the approximation performs better as is argued in the next section.

THE WALD TEST FOR DIFFERENCES OF PROPORTIONS

To check the significance of deviations between the realised default rates of two different rating systems, as opposed to just testing the significance of deviations of one single default rate against a specified value p_0 , a Z-statistic can also be used.

If we define the realised default rate and the number of rated entities of one rating system (1) as df^1 and N_1 respectively and of another rating system (2) as df^2 and N_2 respectively, we can test the null hypothesis $H_0: df^1 = df^2$ (or $df^1 - df^2 = 0$) against $H_1: df^1 \neq df^2$. To derive such a test of difference in default rates we need to pool the default rates of the two rating systems and compute a pooled standard deviation of the difference in default rates in the following way,

¹⁰ The authors are well aware of the fact that the Poisson distribution (discrete and skewed, just like the binomial) is a better approximation than the normal distribution. However the normal approximation is more convenient for differences of proportions (because the difference of independent normal variables is again a normal variable, a property that is not valid for Poisson distributed variables).

$$df^{pooled} = \frac{N_1 df^1 + N_2 df^2}{N_1 + N_2} \quad (12)$$

$$\sigma_{df^1 - df^2} = \sqrt{df^{pooled} (1 - df^{pooled}) \left(\frac{1}{N_1} + \frac{1}{N_2} \right)} \quad (13)$$

Assuming that the two default rates are independent, the corresponding Z-statistic is given by

$$Z = \frac{df^1 - df^2}{\sqrt{df^{pooled} (1 - df^{pooled}) \left(\frac{1}{N_1} + \frac{1}{N_2} \right)}} \quad (14)$$

The value for the Z-statistic may be compared with the percentiles of a standard normal distribution.


Since the binomial distributions considered have success probabilities that are low ($< 0.1\%$) they are all highly skewed to the right. Taking the difference of two right skewed binomial distributions, however, compensates for the asymmetry problem to a large extent.

Chart 2 illustrates the performance of the Wald approximation applied to differences of proportions. For several binomial distributions (i.e. $(N, p) = (500, 0.20\%), (1,000, 0.20\%), (5,000, 0.18\%)$ and $(10,000, 0.16\%)$) the 80% confidence threshold for their difference with respect to the binomial distribution with parameters $(0.07\%, 792)$ is computed using the Wald interval. Then the exact confidence level of this “Wald threshold” is computed.¹¹

The figure shows that for the difference between the binomials with parameters $(792, 0.07\%)$ and $(500, 0.20\%)$ the 80% confidence threshold resulting from the Wald approximation is in fact an 83.60% confidence interval. For the difference between the binomials with parameters $(792, 0.07\%)$ and $(1,000, 0.20\%)$ the 80% confidence threshold resulting from the Wald approximation is in fact a 79.50% confidence interval, and so on.

It can be seen that the Wald approximations for differences in proportions perform better than

Chart 2 Performance of the Wald interval for differences of proportions



the approximations in Chart 1 for single proportions (i.e. the coverage is close to the required 80%). From this it may be concluded that hypothesis tests for differences of proportions, using the normal approximation, work well, as is demonstrated by Chart 2. Thus they seem to be more suitable for our purposes in this context.

THE HOSMER-LEMESHOW TEST (1980, 2000)

The binomial test (or its above mentioned normal/Wald test extensions) is mainly suited to testing a single rating grade, but not several or all rating grades simultaneously. The Hosmer-Lemeshow test is in essence a joint test for several rating grades.

Assume that there are k rating grades with probabilities of default $p_1, \dots, p_k$. Let n_i be the number of obligors with a rating grade i and d_i be the number of defaulted obligors in grade i . The statistic proposed by Hosmer-Lemeshow (HLS) is the sum of the squared differences of forecast and observed numbers of default, weighted by the inverses of the theoretical variances of the number of defaults.

$$HLS = \sum_{i=1}^k \frac{(n_i p_i - d_i)^2}{n_i p_i (1 - p_i)} \quad (15)$$

¹¹ The values that were chosen for the parameters will become clear in Section 4.1.2.

The Hosmer-Lemeshow statistic is χ^2 distributed with k degrees of freedom under the hypothesis that all the probability of default forecasts match the true PDs and that the usual assumptions regarding the adequacy of the normal distribution (large sample size and independence) are justifiable.¹² It can be shown that, in the extreme case, when there is just one rating grade, the HSLS statistic and the (squared) binomial test statistic are identical.

THE SPIEGELHALTER TEST (1986)

Whereas the Hosmer-Lemeshow test, like the binomial test, requires all obligors assigned to a rating grade to have the same probability of default, the Spiegelhalter test allows for variation in PDs within the same rating grade. The test also assumes independence of default events. The starting point is the mean square error (MSE) also known as the Brier score (see Brier 1950)

$$MSE = \frac{1}{N} \sum_i (y_i - p_i)^2 \quad (16)$$

where there are $1, \dots, N$ obligors with individual probability of default estimates p_i . y_i denotes the default indicator, $y = 1$ (default) or $y = 0$ (no default).

The MSE statistic is small if the forecast PD assigned to defaults is high and the forecast PD assigned to non-defaults is low. In general, a low MSE indicates a good rating system.

The null hypothesis for the test is that “all probability of default forecasts, p_i , match exactly the true (but unknown) probability of default” for all i . Then under the null hypothesis, the MSE has an expected value of

$$E[MSE] = \frac{1}{N} \sum_{i=1}^N p_i(1-p_i) \quad (17)$$

and

$$\text{var}[MSE] = \frac{1}{N^2} \sum_{i=1}^N p_i(1-p_i)(1-2p_i)^2 \quad (18)$$

Under the assumption of independence and using the central limit theorem, it can be shown that under the null hypothesis the test statistic

$$Z = \frac{MSE - E[MSE]}{\sqrt{\text{var}[MSE]}} \quad (19)$$

follows approximately a standard normal distribution which allows a standard test decision (see Rauhmeier and Scheule (2005) for practical examples).

CHECKING DEVIATION SIGNIFICANCE IN THE PRESENCE OF DEFAULT CORRELATION

Whereas all the tests presented above assume independence of defaults, it is also important to discuss tests that take into account default correlation. The existence of default correlation within a pool of obligors has the effect of reinforcing the fluctuations in default rate of that pool. The tolerance thresholds for the deviation of realised default rates from estimated values of default may be substantially larger when default correlation is taken into account than when defaults are considered independent. From a conservative risk management point of view, assuming independence of defaults is acceptable, as this approach will overestimate the significance of deviations in the realised default rate from the forecast rate. However, even in that case, it is necessary to determine at least the approximate extent to which default correlation influences probability of default estimates and their associated default realisations.

Most of the relevant literature models correlations on the basis of the dependence of default events on a common systematic random factor (cf. Tasche (2003) and Rauhmeier (2006)). This follows from the Basel II approach underlying risk weight functions which utilise a one factor model.¹³ If D_N is the realised number of defaults in the specified period of time for a 1 to N obligor sample:

¹² If we use the HSLS statistic as a measure of goodness of fit when building the rating model using “in-sample” data then the degrees of freedom of the χ^2 distribution are $k-2$. In the context of this paper, we use the HSLS statistic as backtesting tool on “out of sample” data which has not been used in the estimation of the model.

¹³ See Finger (2001) for an exposition.

$$D_N = \sum_{i=1}^N 1[\sqrt{\rho}X + \sqrt{1-\rho}\varepsilon_i \leq \theta] \quad (20)$$

The default of an obligor i is modelled using a latent variable $AV_i = \sqrt{\rho}X + \sqrt{1-\rho}\varepsilon_i$ representing the asset value of the obligor. The (random) factor X is the same for all the obligors and represents systemic risk. The (random) factor ε_i depends on the obligor and is called the idiosyncratic risk. The common factor X implies the existence of (asset) correlation among the N obligors.

If the asset value AV_i falls below a particular value θ (i.e. the default threshold) then the obligor defaults. The default threshold should be chosen in such a way that $E[D_N] = Np$. This is the case if $\theta = \Phi^{-1}(p)$ where Φ^{-1} denotes the inverse of the cumulative standard normal distribution function and p the probability of default (see e.g. Tasche (2003)). The indicator function $1[\cdot]$ has the value 1 if its argument is true (i.e. the asset value is below θ and the obligor defaults) and the value 0 otherwise (i.e. no default). The variables X and ε_i are normally distributed random variables with a mean of zero and a standard deviation of one (and as a consequence AV_i is also standard normal). It is further assumed that idiosyncratic risk is independent for two different borrowers and that idiosyncratic and systematic risk are independent. In this way, the variable X introduces the dependency between two borrowers through the factor ρ , which is the asset correlation (i.e. the correlation between the asset values of two borrowers). Asset correlation can be transformed into default correlations as shown, for example, in Basel Committee on Banking Supervision (2005b).

Tasche 2003 shows that on a confidence level α we can reject the assumption that the actual default rate is less than or equal to the estimated probability of default whenever the number of defaults D is greater than or equal to the critical value given by

$$D_\alpha = q + \frac{2q-1}{2N} - \frac{q(1-q)}{\varphi\left(\frac{\sqrt{\rho}\phi^{-1}(1-\alpha) - \phi^{-1}(PD)}{\sqrt{1-\rho}}\right)} \times \frac{(1-2\rho)\phi^{-1}(1-\alpha) - \sqrt{\rho}\phi^{-1}(PD)}{2N\sqrt{\rho(1-\rho)}} \quad (21)$$

where

$$q = \phi^{-1}\left(\frac{\sqrt{\rho}\phi^{-1}(\alpha) + \phi^{-1}(PD)}{\sqrt{1-\rho}}\right), \varphi(x) = \frac{d\Phi(x)}{dx}, \quad (22)$$

and Φ^{-1} denotes the inverse of the cumulative standard normal distribution function and ρ the asset correlation. However, the above test, which includes dependencies and a granularity adjustment, as in the Basel II framework, shows a strong sensitivity to the level of correlation.¹⁴

It is interesting to see how the binomial test and the correlation test as specified above behave under different assumptions. As can be seen in Tables 3 and 4, the critical number of defaults that can be allowed before we could reject the null hypothesis that the estimated probability of default is in line with the realised number of defaults, goes up as we increase the level of asset correlation among obligors for every level of sample size from 0.05 to 0.15.¹⁵ The binomial test produces consistently lower critical values of default than the correlation test for all sample sizes. However, the test taking into account correlation suffers from dramatic changes in the critical values, especially for larger sample sizes (i.e. over 1,000).

¹⁴ Tasche (2003) also discusses an alternative test to determine default-critical values assuming a Beta distribution, with the parameters of such a distribution being estimated by a method of matching the mean and variance of the distribution. This approach will generally lead to results that are less reliable than the test based on the granularity adjustment.

¹⁵ The $\rho = 0.05$ may be justified by applying the non-parametric approach proposed by Gordy (2002) to data on the historical default experiences of all the rating grades of Standard & Poor's, which yields an asset correlation of ~5%. Furthermore, Tasche (2003) also points out that " $\rho = 0.05$ appears to be appropriate for Germany". 24% is the highest asset correlation according to Basle II (see Basel Committee on Banking Supervision (2005a)).

Table 3 95% critical values for a benchmark probability of default of 0.10% under different calibration tests

N	Binomial	Correlation = 0.05	Correlation = 0.15
100	1	2	2
500	2	3	3
1,000	3	4	5
5,000	9	15	21

Table 4 99.9% critical values for a benchmark probability of default of 0.10% under different calibration tests

N	Binomial	Correlation = 0.05	Correlation = 0.15
100	2	4	4
500	2	6	12
1,000	5	10	22
5,000	13	37	102

As can be inferred from the above tables, the derivation of critical values of default, taking into account default correlation, is not a straight forward exercise. First, we need to have a good estimate of asset correlation. In practice, this number could vary depending on the portfolio considered. A well-diversified portfolio of retail loans across an extensive region will present very different correlation characteristics than that of a sector concentrated portfolio of corporate names. In practice, default correlations could be seen in the range of 0-5%.¹⁶ Second, the validation analyst should take into account that there should be a consistency between the modelling of correlation for risk measurement in the credit assessment system that is going to be validated and the validation test to derive consistent confidence intervals for such credit system. This consistency is in practice difficult to achieve because the correlation dynamics in the validation test may not be in line with those assumed in the rating system.

The binomial test, although conservative, is seen as a good realistic proxy for deriving critical values. It is considered a good early warning tool, free of all the estimation problems seen in tests that incorporate correlation estimates. Therefore, in what remains of this paper we will focus on the binomial distribution paradigm and its extension in the normal distribution as the general statistical framework to derive a simple mechanism for performance checking based on backtesting.

4.1 TWO POSSIBLE BACKTESTING STRATEGIES

In what follows we will concentrate on elaborating two backtesting strategies that focus on a rating level of a single “A” as defined by the main international rating agencies. This is the credit quality level set by the Eurosystem for determining eligible collateral for its monetary policy operations. The single “A” rating is thus considered the “benchmark”. The previous section defined this benchmark in terms of a probability of default. As the true probability of default is unobservable, this section presents two alternative backtesting strategies based on the (observable) default frequency:

1. A backtesting strategy that uses a fixed, absolute upper limit for the probability of default as a benchmark.
2. A backtesting strategy that uses a stochastic benchmark. This assumes that the benchmark is not constant as in the first strategy

These alternatives will be summarised in a simplified rule which will result in a traffic light approach for backtesting, much in the same vein as in Tasche (2003), Blochwitz and Hohl (2001) or Tiomo (2004).

¹⁶ Hushens and Stahl (2005) show evidence that, for a well diversified German retail portfolio, asset correlations are in the range between 0% and 5%, which implies even smaller default correlations.

4.1.1 A BACKTESTING STRATEGY RELYING ON A FIXED BENCHMARK

Using the central limit theorem we found in Section 3 that the probability of default of the benchmark (p^{bm}) is at most 0.1%. A rating source is thus in line with the benchmark if its default probability for the single “A” rating is at most 0.1%.

Assuming that the rating source’s default events are distributed in accordance with a binomial distribution with parameters PD^{rs} and N_y^{rs} , the backtesting should check whether

$$PD^{rs} \leq 0.1\% \quad (23)$$

Since PD^{rs} is an unobservable variable, (23) can not be used for validation purposes. A quantity that can be observed is the number of defaults in a rating source’s static pool within one particular year, i.e. df_y^{rs} .

The performance checking mechanism should thus check whether observing a value df_y^{rs} for a random variable which is (approximately) normally distributed $df_y^{rs} \approx N\left(PD^{rs}; \sqrt{\frac{PD^{rs}(1-PD^{rs})}{N_y^{rs}}}\right)$ is consistent with (23).

This can be done using a statistical hypothesis test. The null hypothesis that $H_0: p^{rs} \leq 0.1\%$ must be tested against the alternative hypothesis $H_1: p^{rs} > 0.1\%$.

Assuming that the null hypothesis of this statistical test, H_0 , is true, the probability of observing the value df_y^{rs} can be computed. This is the p -value of the hypothesis test or the probability of obtaining a value of df_y^{rs} or higher, assuming that H_0 is true. This p -value is given by

$$1 - \Phi \left(\frac{df_y^{rs} - 0.1\%}{\sqrt{\frac{0.1\%(1-0.1\%)}{N_y^{rs}}}} \right) \quad (24)$$

where Φ is the cumulative probability function for the standard normal distribution. Table 5

Table 5 Test of credit quality assessment source against the limit of 0.1% for a sample size of 10,000. N denotes the number of defaults

(percentage)			
N	df ^(rs)	p-value	Probability of “N” if H_0 is true
0	0.00		
1	0.01	99.78	0.35
2	0.02	99.43	0.77
3	0.03	98.66	1.54
4	0.04	97.12	2.80
5	0.05	94.32	4.60
6	0.06	89.72	6.84
7	0.07	82.87	9.22
8	0.08	73.66	11.24
9	0.09	62.41	12.41
10	0.10	50.00	12.41
11	0.11	37.59	11.24
12	0.12	26.34	9.22
13	0.13	17.13	6.84
14	0.14	10.28	4.60
15	0.15	5.68	2.80
16	0.16	2.88	1.54
17	0.17	1.34	0.77
18	0.18	0.57	0.35
19	0.19	0.22	0.14
20	0.20	0.08	0.05
21	0.21	0.03	0.02
22	0.22	0.01	0.01
23	0.23	0.00	0.00
24	0.24	0.00	0.00
25	0.25	0.00	0.00

gives an example for an eligible set of $N_y^{rs} = 10,000$ companies.

The first column of the table gives different possibilities for the number of defaults observed in year “Y”. The observed default frequency is derived by dividing the number of defaults by the sample size. This is shown in the second column of the table. The third column shows the p -values computed using formula (24). So the p -value for observing at least 15 defaults out of 10,000, assuming that H_0 is true, equals 5.68%. In the same way it follows from the table that if H_0 is true, then the probability of observing at least 18 defaults in 10,000 is 0.57%, or “almost impossible”. Or, to put it another way, if we observe 18 defaults or more then it is almost impossible for H_0 to be true.

The last column “probability” computes the theoretical probability for observing a particular number of defaults if H_0 is true. It is the difference between two successive p -values. For example, if H_0 is true, then the probability of observing at least one default out of 10,000 equals 99.78%, and the probability of observing at least two defaults is 99.43%. As a consequence, if H_0 is true, the probability of having exactly one default is 0.35%. The column “probability” can thus be used as an exact behavioural rule, i.e. if H_0 is true then one can have

- exactly one default in 10,000 every 0.35 years out of 100 years
- exactly two defaults in 10,000 every 0.77 years out of 100 years
- exactly three defaults in 10,000 every 1.54 years out of 100 years
-

Averaging this rule over a 100 year period shows that in the long run the average default frequency will converge to 0.1%. However, such a rule is, of course, too complex to be practical. It is simplified below.

Table 5 can be used as backtesting for a sample of 10,000 obligor names with an ex ante probability of default of 0.10% after fixing a confidence level (i.e. a minimum p -value, e.g. 1%): if the size of the static pool is 10,000

then the rating source is in line with the benchmark only if at most 17 defaults are observed (confidence level of 1%) i.e. $df_{Y}^{rs} \leq 0.17\%$.

This technique has the disadvantage of first having to decide on a confidence level. Moreover, fixing only one limit (0.17% in the case above) does not guarantee a convergence over time to an average of 0.1% or below.

A p -value, being a probability, can be interpreted in terms of “number of occurrences”. From Table 5 we infer that, if the null hypothesis is true, the observed default frequency must be lower than 0.12% in 80% of cases. In other words, a value above 0.12% should be observable only once every 5 years (i.e. if the realised default frequency should be lower than 0.12% in 80 out of 100 years, then the realised default frequency could be higher than 0.12% in 20 out of 100 years, or once every 5 years), otherwise the rating source is not in line with the benchmark.

This gives a second performance checking rule: a rating source with a static pool of size 10,000 is in line with the benchmark if at most once every five years a default frequency above 0.12% is observed. A default frequency above 0.17% should “never” be observed.

Table 6 Backtesting strategy based on a fixed benchmark for different static pool sizes

(percentage)	All time	Once in 5y	Never	Average DF
500	0.00-0.00	0.20-0.40	>0.40	0.06
1,000	0.00-0.10	0.20-0.40	>0.40	0.10
2,000	0.00-0.10	0.15-0.25	>0.25	0.08
3,000	0.00-0.10	0.13-0.23	>0.23	0.08
4,000	0.00-0.13	0.15-0.25	>0.25	0.09
5,000	0.00-0.12	0.14-0.20	>0.20	0.08
6,000	0.00-0.12	0.13-0.20	>0.20	0.08
7,000	0.00-0.11	0.13-0.19	>0.19	0.08
8,000	0.00-0.11	0.13-0.19	>0.19	0.08
9,000	0.00-0.11	0.12-0.18	>0.18	0.07
10,000	0.00-0.11	0.12-0.17	>0.17	0.07
50,000	0.00-0.112	0.114-0.134	>0.134	0.07

The intervals for other sizes of the static pool are shown in Table 6. The lower value of the “Once in 5y” interval is derived from the 80% confidence limit, the absolute upper limit “Never” is derived from a 99% confidence interval.

The column “average DF” is an estimated average using 4 in 5 occurrences at the midpoint of the first interval and 1 in 5 occurrences at the midpoint of the second. These averages are clearly below the benchmark limit of 0.1%.

Notice, however, that the validation strategies proposed above make use of hypothesis tests for one proportion. As illustrated earlier in Section 4, the Wald approximation performs worse for one proportion than for differences of proportions. Hence an alternative test based on differences of proportions will be developed in the following section.

4.1.2 A BACKTESTING STRATEGY BASED ON A STOCHASTIC BENCHMARK

In the preceding section a performance checking mechanism using a fixed upper limit for the benchmark was derived. That fixed upper limit followed from the central limit theorem and was found to be 0.1%.

An examination of Table 1 could also prompt the idea that the benchmark is not fixed but stochastic. Thus we will develop an alternative backtesting strategy in this section, based on a stochastic benchmark. In fact, in Section 3 we concluded that the benchmark can be defined as

$$df^{bm} \approx N \left(PD^{bm}; \sqrt{\frac{PD^{bm}(1-PD^{bm})}{N^{bm}}} \right) \quad (25)$$

where PD^{bm} was estimated at 0.04% and N^{bm} was estimated at 792.

On the other hand, the rating source’s default frequency is also normally distributed,

$$df_Y^{rs} \approx N \left(PD^{rs}; \sqrt{\frac{PD^{rs}(1-PD^{rs})}{N_Y^{rs}}} \right) \quad (26)$$

If one assumes a stochastic benchmark, there is no longer an upper limit for the PD of the rating source. The condition on which to base the performance-checking mechanism should be that “the rating source should do at least as well as the benchmark”. In terms of a probability of default, this means that the rating source’s PD should be lower than or equal to that of the benchmark. The hypothesis to be tested is thus $H_0: PD^{rs} \leq PD^{bm}$ against $H_1: PD^{rs} > PD^{bm}$ where PD^{bm} was estimated at 0.04% and N^{bm} was estimated at 792.

The test is completely different from the one in the preceding section. Indeed we *cannot* replace PD^{bm} by 0.04% because this is only an estimate of the benchmark’s PD. The true PD of the benchmark is unknown. We should therefore combine the variance of the ex-ante estimated probability of default of the rating source and that of the benchmark in one measure in order to conduct the backtesting.

The difference of two normally distributed variables also has a normal distribution thus, assuming that both are independent:¹⁷

$$df_Y^{rs} - df^{bm} \approx N \left(PD^{rs} - PD^{bm}; \sqrt{\frac{PD^{rs}(1-PD^{rs})}{N_Y^{rs}} + \frac{PD^{bm}(1-PD^{bm})}{N^{bm}}} \right) \quad (27)$$

PD^{rs} and PD^{bm} are unknown, but if the null hypothesis is true then their difference should be $PD^{rs} - PD^{bm} \leq 0$. An estimate of the combined variance $\frac{PD^{rs}(1-PD^{rs})}{N_Y^{rs}} + \frac{PD^{bm}(1-PD^{bm})}{N^{bm}}$ is needed. A standard hypothesis test, testing the equality of two proportions, would use a “pooled variance” as estimator. This pooled variance itself being derived from a “pooled proportion” estimator (see e.g. Moore and McCabe (1999), and Cantor and Falkenstein (2001)). The reasoning is that as we test the

17 If the rating source’s eligible class and the benchmark are dependant then the variance of the combined normal distribution should include the covariance term.

hypothesis of equal proportions, all observations can be pooled so that there are a total of $N_Y^{rs} + N^{bm}$ observations, among which there are $N^{bm} df^{bm} + N_Y^{rs} df_Y^{rs}$. The pooled proportion is thus

$$df^{pooled} = \frac{792 \times 0.04\% + N_Y^{rs} df_Y^{rs}}{792 + N_Y^{rs}} \quad (28)$$

and the two variances are then

$$\sigma_{bm}^2 = \frac{df^{pooled} (1 - df^{pooled})}{792},$$

$$\sigma_{rs}^2 = \frac{df^{pooled} (1 - df^{pooled})}{N_Y^{rs}} \quad (29)$$

and (9) becomes

$$df_Y^{rs} - df^{bm} \approx N \left(0; \sqrt{df^{pooled} (1 - df^{pooled}) \left(\frac{1}{N_Y^{rs}} + \frac{1}{792} \right)} \right) \quad (30)$$

However, as we have an estimate of the benchmark that is based on only 24 past observations, we decided not to touch the variance estimate of the benchmark. So, taking a risk-averse position, we assigned the empirical variance of the benchmark recorded in these 24 observations (1981-2004) untouched, and the hypothesis test uses the distribution given in (31).

$$df_Y^{rs} - df^{bm} \approx N \left(0; \sqrt{\frac{df^{pooled} (1 - df^{pooled})}{N_Y^{rs}} + 0.07\%^2} \right) \quad (31)$$

Using the observed default frequency (df^{rs}) as an estimate for the rating source, and using the estimated benchmark values, the p -values of the test are given by:

$$1 - \Phi \left(\frac{df^{rs} - 0.04\%}{\sqrt{0.07\%^2 + \frac{df^{pooled} (1 - df^{pooled})}{N_Y^{rs}}}} \right) \quad (32)$$

Table 7 Test of credit quality assessment source against the limit of stochastic benchmark for a sample size of 10,000, using $df^{bm} = 0.04\%$

(percentages)			
BM	Mean 0.04	Stand dev 0.07	df [*] (bm) 0.04
Size	10,000	mean	
N	df [*] (rs)	p-value	Probability of "N" if H ₀ is true
0	0.0000	71.27	5.10
1	0.0100	66.17	5.31
2	0.0200	60.86	5.43
3	0.0300	55.43	5.43
4	0.0400	50.00	5.34
5	0.0500	44.66	5.16
6	0.0600	39.49	4.91
7	0.0700	34.59	4.59
8	0.0800	30.00	4.23
9	0.0900	25.77	3.84
10	0.1000	21.94	3.44
11	0.1100	18.50	3.04
12	0.1200	15.46	2.65
13	0.1300	12.81	2.29
14	0.1400	10.52	1.95
15	0.1500	8.57	1.65
16	0.1600	6.92	1.37
17	0.1700	5.55	1.14
18	0.1800	4.41	0.93
19	0.1900	3.48	0.75
20	0.2000	2.73	0.61
21	0.2100	2.12	0.48
22	0.2200	1.64	0.38
23	0.2300	1.26	0.30
24	0.2400	0.96	0.23
25	0.2500	0.73	0.18

The results for an estimated benchmark PD of 0.04% and a static pool of 10,000 companies are shown in Table 7.

Table 7 can be used as backtesting for a sample of 10,000 obligor names with an stochastic benchmark, after fixing a confidence level (i.e. a minimum p -value, e.g. 1%). A rating source with a static pool of size 10,000 is in line with the benchmark if at most once every five years a default frequency above 0.1% is observed. A default frequency above 0.23% should "never" be observed.

The intervals for other sizes of the static pool are shown in Table 8. The average default

Table 8 Backtesting strategy based on a stochastic benchmark for different static pool sizes

(percentages)						
	p (bm) = 0.04%			p (bm) = 0.07%		
	All time	Once in 5y	Average DF	All time	Once in 5y	Average DF
500	0-0	0.2-0.6	0.08	0-0	0.2-0.8	0.10
1,000	0-0	0.2-0.5	0.07	0-0.1	0.2-0.5	0.11
2,000	0-0.1	0.15-0.35	0.09	0-0.1	0.2-0.45	0.11
3,000	0-0.1	0.13-0.3	0.08	0-0.1	0.13-0.37	0.09
4,000	0-0.1	0.125-0.275	0.08	0-0.125	0.15-0.37	0.10
5,000	0-0.1	0.12-0.28	0.08	0-0.16	0.18-0.34	0.12
6,000	0-0.1	0.12-0.25	0.08	0-0.15	0.16-0.35	0.11
7,000	0-0.1	0.11-0.24	0.08	0-0.15	0.17-0.34	0.11
8,000	0-0.1	0.11-0.237	0.07	0-0.15	0.16-0.32	0.11
9,000	0-0.1	0.11-0.23	0.07	0-0.15	0.16-0.32	0.11
10,000	0-0.1	0.11-0.23	0.07	0-0.15	0.16-0.32	0.11
50,000	0-0.1	0.11-0.21	0.07	0-0.15	0.152-0.29	0.10

frequency seems to be lower than 0.1% for all sizes. As argued in Section 3, a higher average than 0.04% could be justified. Table 8 also shows the results when an estimate of 0.07% is used for the benchmark PD. As in Table 6, the lower value of the “Once in 5y” interval is derived from the 80% confidence limit, the absolute upper limit “Never” is derived from a 99% confidence interval.

It appears that with these assumptions (i.e. an ex ante estimate for the probability of default of 0.07% and benchmark sample size of 792 obligors) the backtesting strategy based on the stochastic benchmark is less conservative than that based on a fixed benchmark (cf. Table 6 and 8). For every sample size the confidence intervals for the realised default rates are wider for the stochastic benchmark test. It can also be seen that the confidence intervals are relatively wider for large sample sizes when using the stochastic benchmark.

4.1.3 THE BACKTESTING MECHANISMS AND BASEL II

Under the new rules of Basel II supervisors will be responsible for assigning an eligible External Credit Assessment Institution’s (ECAI) credit risk assessment to the risk weights available under the standardised approach. Annex 2 of

the revised Basel II framework (2005a) contains the Committee’s proposal for a consistent mapping of credit risk assessments into the available risk weights. The mapping mechanism uses two quantifiable parameters:

- a ten-year average of a three-year cumulative default frequency, and
- the two most recent three-year cumulative default frequencies

These measures have to be compared to benchmark values:

- For the two most recent three-year cumulative default frequencies the Basel II documents give two benchmark levels:
 - a monitoring level derived from Monte Carlo simulations and fixed at the 99% quantile which takes a value of 1.0% for the “A” grade, and
 - a trigger level representing the 99.9% quantile, which takes a value of 1.3% for the “A” grade.

Using the 10 year average of the 3-year default frequencies observed over the years 1993-2002

Table 9 Confidence intervals based on a stochastic benchmark rule applied to the 3 year default rates

(percentages)	Conf 99%	Conf 99.9%
500	1.00	1.40
600	1.00	1.30
700	1.00	1.28
800	0.90	1.25
900	0.90	1.22
1,000	0.90	1.20
11,00	0.90	1.20
12,00	0.90	1.15
Average	0.94	1.25
p.m. Basel II	1.00	1.30

by Standard&Poor's (0.25%) and the average number of issuers over the same period (1,024), the 99% and the 99.9% confidence intervals resulting from tests with a stochastic benchmark are shown for different static pool sizes in Table 9.

This table shows that our stochastic benchmark tests for confidence levels of 99% and 99.9% respectively yield results similar to the threshold values given in Annex 2 of the revised Basel II framework. It appears that our proposed test is slightly more conservative. If, instead of the stochastic benchmark, we were to use the fixed benchmark strategy, the confidence intervals would be even more conservative than those provided by the Basel II rules.

4.2 THE TRAFFIC LIGHT APPROACH, A SIMPLIFIED BACKTESTING MECHANISM

The testing procedures outlined in the previous sections allow

- confidence intervals to be defined for the annual default frequency (i.e. the annual interpretation of the rule) and
- the specification of how often a value should fall within a specific interval (i.e. the multi-period interpretation of the rule) in order to converge to a long-run average of 0.10%.

Moreover, the intervals for the annual default frequency depend on the size of a credit quality assessment source's eligible set (i.e. the static pool).

To apply the tests discussed in 4.1 in practice, a traffic light approach is proposed. Instead of defining exactly how often every possible default frequency may be observed for a certain credit quality assessment source, Tables 6 and 8 can be simplified to a restriction on how often a realised default frequency should fall within one of only three intervals (the three zones (green, orange, and red) of the "traffic light approach").

1. Depending on the size of a rating source's static pool, two threshold levels are defined that separate these three zones: (1) a monitoring level, and (2) a trigger level.
2. If the annually observed default frequency is (strictly) below the monitoring level, then the rating source is in line with the benchmark and is in the green zone.
3. If the observed default frequency is above (or equal to) the monitoring level and (strictly) below the trigger level, then the rating source is in the orange zone. The rating source is allowed to be in the orange zone only once in five years (on average).
4. If the observed default frequency is above (or equal to) the trigger level, then the rating source is in the red zone.

A practical example of a traffic light approach as defined above, with the monitoring and the trigger levels derived from Table 8, is given in Table 10 below. A similar example of a traffic light approach could also be constructed on the basis of Table 6 (test based on a fixed benchmark rule). The monitoring and trigger levels would be somewhat more conservative than those shown in Table 8.

One way of applying the traffic light approach in practice, assuming that a rating source has a

Table 10 Example of traffic light monitoring and trigger levels based on a stochastic benchmark rule

(percentages)		
Size of eligible set	Monitoring level (orange)	Trigger level (red)
Up to 500	0.20	1
Up to 1,000	0.20	0.80
Up to 5,000	0.18	0.34
Up to 50,000	0.16	0.28

static pool in the vicinity of 500 obligors, could be as follows: if the rating source records an annual realised default rate that is in the red zone (i.e. a realised default rate above 1%) or a default rate that repeatedly falls in the orange zone (i.e. more than once over a period of 5 years), then the analyst may wish to consult with the relevant rating provider to understand why its default experience is considerably worse than the historical default experience of the benchmark rating agencies. The credit quality assessment system provider will be asked to provide additional information to justify its violation of the rules.

If no convincing arguments were brought forward, then the conclusion would be that the rating source's model estimates probabilities of default which are too low and the model must be re-calibrated.¹⁸

Finally, please note that default frequencies are discrete variables, so the upper limit of the green zone can be far below the lower limit of the orange zone. E.g. for a static pool size below 500 the green zone means no defaults at all, because the lowest non zero is 0.2% (1/500) (see Table 8).

¹⁸ If forecast PDs and ex-post default information are available for every individual borrower in the static pool of the rating source, then the Brier score/Spiegelhalter test, for example, could be used to check the forecasting performance of the rating source's model.

5 SUMMARY AND CONCLUSIONS

In this paper we concentrate on two main goals. First, we are interested in translating the single “A” rating as published by major rating agencies for debt issuers into a quantitative metric, the annual probability of default. In particular we look at the single “A” rating, because that is the minimum level at which the Eurosystem sets its requirement of high credit standards for collateral that can be used in its monetary policy operations. This translation method could be useful for mapping credit assessments of other rating sources to those of the major rating agencies by means of this probability of default metric. Although, the information that is contained in a rating goes beyond the probability of default of an obligor, we present arguments in support of the translation from a rating to a PD. The example presented with the single “A” rating could also be extended to other rating grades. We demonstrate that the probability of default for a single “A” issued by the main rating agencies is at most 0.1%.

Second, we are interested in assessing the quality of the estimated probability of default for a rating grade, in this case the single “A” rating grade, by means of the realised default rate, also called backtesting of the probability of default. We review briefly the main statistical tests that have appeared in the literature focusing on the binomial test and its normal extension, analysing in particular its main underlying assumptions, independence of default events and constant probability of default. We show that the existence of default correlation would imply wider confidence intervals than those derived with the binomial test, in which independence of default events is assumed. However, it is also argued that in practice default correlations would be low, in particular for high credit quality debtors, and that, from a risk management perspective, it is preferable to rely on a more conservative test, such as the binomial test to derive critical values.

Assuming that the default generating process follows a binomial distribution, the paper proposes two generic backtesting strategies for testing the quality of forecast probabilities of default: first, a backtesting strategy that uses a fixed, absolute upper limit for the probability of default, which in the case of a single “A” rating is derived at 0.10%; and second, a backtesting strategy that relies on a stochastic benchmark, a benchmark probability of default that is not constant, unlike in the first strategy. The second strategy could be justified in cases where there is uncertainty about the level of the benchmark or if the benchmark is expected to move over time. We show that a backtesting strategy based on a stochastic benchmark would produce wider confidence intervals than those obtained using a fixed benchmark. The two strategies are based on one and five-year multi-period tests. The use of a multi-period test is intended to provide a more informative statement about the performance of a rating source as reliance only on annual tests may be misleading due to, for example, problems in the measurement of default with scarce data, situations of unforeseen non-normal stress that increase default rates, or the existence of default correlation.

The backtesting strategies presented are implemented through a traffic light approach in the same vein as in Tasche (2003). Depending on the size of a rating source’s static pool, two threshold levels are defined that separate three zones, green, orange and red: (1) a monitoring level, and (2) a trigger level. If the annually observed default frequency is (strictly) below the monitoring level, then the rating source is in the green zone and is in line with the benchmark. If the observed default frequency is above (or equal to) the monitoring level and (strictly) below the trigger level, then the rating source is in the orange zone, which implies that the realised default rate is not compatible with the PD forecast but still in the range of usual statistical deviations. We implement the multi-period rule by allowing the default frequency to be in the orange zone only once in five years (on average). If the observed default frequency

is above (or equal to) the trigger level, then the rating source is in the red zone, indicating that the realised default frequency is unequivocally not in line with the PD forecast.

We see the backtesting techniques and strategies described in this paper as early warning tools for identifying performance problems in credit assessment systems. This could be useful in the context of the Eurosystem Credit Assessment Framework, in which various credit assessment sources can be employed to assess the credit quality standards of eligible collateral. In such a setting it is important to guarantee that the different participating credit systems fulfil their rating mandates correctly and comparably. In this sense, this paper puts emphasis on risk management and therefore there is a general preference for backtesting strategies that are more conservative. However, the techniques presented in this paper are by no means the only mechanism that an analyst has at his disposal for validating the functioning of a credit system, but they are an important one. They could be considered as a first step (i.e. early warning) in a more comprehensive process that should take into account also more qualitative elements (see Basel Committee (2005b)). The drawbacks shown in this paper as regards problems of measurement, existence of correlation, or the existence of non-normal stress situations should be weighed carefully when assessing credit assessment systems based solely on the results of backtesting tools of the type presented.

ANNEX

HISTORICAL DATA ON MOODY'S A-GRADE

Historical data on Moody's A-grade		
Year	Issuers	1Y-Def.-Freq. (%)
1981	376	0.00
1982	387	0.26
1983	432	0.00
1984	472	0.00
1985	524	0.00
1986	579	0.00
1987	555	0.00
1988	553	0.00
1989	587	0.00
1990	614	0.00
1991	609	0.00
1992	694	0.00
1993	740	0.00
1994	880	0.00
1995	968	0.00
1996	1,071	0.00
1997	1,133	0.00
1998	1,154	0.00
1999	1,173	0.00
2000	1,237	0.00
2001	1,287	0.16
2002	1,301	0.16
2003	1,279	0.00
2004	1,244	0.00
Mean		0.02
Standard deviation		0.07
Source: Moody's (2005), Moody's Default Report 2005. Annual Default Study.		

REFERENCES

- Agresti, A. and Coull, B. (1998), Approximate is better than “exact” for interval estimation of binomial proportions, *The American Statistician*, May.
- Agresti, A. and Caffo, B. (2000), Simple and effective confidence intervals for proportions and differences of proportions result from adding two successes and two failures, *The American Statistician*, November.
- Basel Committee on Banking Supervision (2005a), International Convergence of Capital Measurement and Capital Standards. A Revised Framework.
- Basel Committee on Banking Supervision (2005b), Studies on the validation of internal rating systems (revised), Working Paper, Bank for International Settlements.
- Billingsley, P. (1995), Probability and Measure, 3rd edition, John Wiley & Sons.
- Blochwitz, S. and Hohl, S. (2001), The worst case or what default rates do we have to expect from the rating agencies? Working Paper, Deutsche Bundesbank, Frankfurt.
- Blochwitz, S., When, C. and Hohl, S. (2003), Reconsidering ratings, Working Paper. Deutsche Bundesbank, Frankfurt.
- Brier, G. (1950), Verification of forecasts expressed in terms of probability, *Monthly Weather Review*, vol. 78, No. 1, pp. 1-3.
- Brown, L. D., Cai, T. and DasGupta, A. (2001), Interval estimation for a binomial proportion, *Statistical Science*, Vol. 16, No. 2.
- Cai, T. (2005), One-sided confidence intervals in discrete distributions, *J. Statistical Planning and Inference*, Vol. 131.
- Cantor, R. (2001), Moody’s Investors Service’s response to the consultative paper issued by the Basel Committee on Banking Supervision “A new capital adequacy framework”, *Journal of Banking and Finance*, Vol. 25.
- Cantor, R. and Falkenstein, E. (2001), Testing for rating consistency in annual default rates, *Journal of Fixed Income*.
- Cantor, R. and Mann, C. (2003), Are corporate bond ratings procyclical? *Moody’s Special Comment*, October.
- Cantor, R., Packer, F. and Cole, K. (1997), Split ratings and the pricing of credit risk, *Journal of Fixed Income*, December.
- Christensen, J., Hansen, E. and Lando, D. (2004), Confidence sets for continuous-time rating transition probabilities, *Journal of Banking and Finance*, Vol. 28.

- Crouhy, M., Galai, D. and Mark, R. (2001), Prototype risk rating system, *Journal of Banking and Finance*, Vol. 25.
- DeGroot, M. (1989), Probability and statistics, Second edition, Addison-Wesley Publishing Company.
- Elton, E., Martin, J., Gruber, J., Agrawal, D. and Mann, C. (2004), Factors affecting the valuation of corporate bonds, in Cantor, R. (ed.), Recent Research on Credit Ratings, *Journal of Banking and Finance*, Vol. 28, special issue.
- European Central Bank (2007), The implementation of monetary policy in the Euro area. General documentation on Eurosystem monetary policy instruments and procedures.
- Finger, C. (2001), The one-factor Creditmetrics model in the new Basel Capital Accord, *Riskmetrics journal*, Vol. 2(1), pp. 9-18.
- Fons, J. (2002), Understanding Moody's corporate bond ratings and rating process, *Moody's Special Comment*, May.
- Gordy, M. (1998), A Comparative Anatomy of Credit Risk Models, Working Paper, Federal Reserve System.
- Hamerle, A., Liebig, T., Röscher, D. (2003), Benchmarking Asset Correlations, *RISK*, Nov., pp. 77-81.
- Heitfield, H. (2005), Studies on the Validation of Internal Rating Systems, Working Paper 14, Basel Committee on Banking Supervision.
- Hosmer, D., Lemeshow, S. and Klar J. (1988), Goodness of fit testing for multiple logistics regression analysis when the estimated probabilities are small, *Biometrical Journal*, Vol. 30, pp. 991-924.
- Hosmer, D. and Lemeshow, S. (1980), A goodness of test for the multiple logistic regression, *Communication in Statistics*, A10, 1043-1069.
- Hosmer, D. and Lemeshow, S. (2000), Applied logistic regression, Wiley series in Probability and Statistics.
- Huschens, S. and Stahl, G. (2005), A general framework for IRBS backtesting, *Bankarchiv, Zeitschrift für das gesamte Bank und Börsenwesen*, 53, pp. 241-248.
- Hui, C.-H., Wong, T.-C., Lo, C.-F. and M.-X. Huang (2005), Benchmarking Model of Default Probabilities of Listed Companies, *Journal of Fixed Income*, September.
- Hull, J., Predescu, M. and White, A. (2004), The relationship between credit default swap spreads, bond yields, and credit rating announcements, in Cantor, R. (ed.), Recent Research on Credit Ratings, *Journal of Banking and Finance*, Vol. 28, special issue.

- Jafry, Y. and Schuermann, T. (2004), Measurement, Estimation and Comparison of Credit Migration Matrices, *Journal of Banking and Finance*.
- Johnson, N. L. (1969), Discrete distributions, Houghton Mifflin Company, Boston.
- Krahnhen, J. P. and Weber, M. (2001), Generally accepted rating principles: a primer, *Journal of Banking and Finance*, Vol. 25.
- Lando, D. and Skødeberg, T. M. (2002), Analyzing rating transitions and rating drift with continuous observations, *Journal of Banking and Finance*, Vol. 26.
- Moody's (2005), Moody's Default Report 2005. Annual Default Study.
- Moore, D. S. and McCabe, G. P. (1999), Introduction to the practice of statistics, W. H. Freeman and Company, New York.
- Newcombe, R. G. (1998), Two-sided confidence intervals for the single proportion: comparison of seven methods, *Statistics in Medicine*, Vol. 17.
- Nickel, P., Perraudin, W. and Varotto, S. (2000), Stability of Rating Transitions, *Journal of Banking and Finance*, Vol. 24.
- Rauhmeier, R. (2006), PD-validation, experience from banking practice, in Engelman, B. and Rauhmeier, R. (eds.), The Basel II risk parameter, estimation, validation and stress testing, Springer, Berlin, pp. 307-343.
- Rauhmeier, R. and Scheule, H. (2005), Rating properties and their implications for Basel II capital, *Risk*, 18(3), pp. 78-81.
- Reiczigel, J. (2004), Confidence intervals for the binomial parameter: some new considerations, Working Paper, Szent István University, Budapest.
- Rohatgi, V. K. (1984), Statistical Inference, Wiley Series in Probability and mathematical statistics, John Wiley & Sons.
- Spanos, A. (1986), Statistical foundations of econometric modelling, Cambridge University Press.
- Spiegelhalter, D. (1986), Probabilistic prediction in patient management and clinical trials, *Statistics in Medicine*, Vol. 5, pp. 421-433.
- Standard & Poor's (2005), Annual global corporate default study: corporate defaults poised to rise in 2005. Global fixed income research.
- Tasche, D. (2003), A traffic lights approach to PD validation, Working Paper, Deutsche Bundesbank, Frankfurt.
- Tasche, D. (2006), Validation of internal rating systems and PD estimates, Working paper, Deutsche Bundesbank, Frankfurt.

REFERENCES

- Tiomo, A. (2004), Credit risk and variability of default rates: an empirical analysis using simulations, Working Paper, Banque de France, Paris.
- Vollset, S. E. (1993), Confidence intervals for a binomial proportion, *Statistics in Medicine*, Vol. 12.

**EUROPEAN CENTRAL BANK
OCCASIONAL PAPER SERIES**

- 1 “The impact of the euro on money and bond markets” by J. Santillán, M. Bayle and C. Thygesen, July 2000.
- 2 “The effective exchange rates of the euro” by L. Buldorini, S. Makrydakis and C. Thimann, February 2002.
- 3 “Estimating the trend of M3 income velocity underlying the reference value for monetary growth” by C. Brand, D. Gerdesmeier and B. Roffia, May 2002.
- 4 “Labour force developments in the euro area since the 1980s” by V. Genre and R. Gómez-Salvador, July 2002.
- 5 “The evolution of clearing and central counterparty services for exchange-traded derivatives in the United States and Europe: a comparison” by D. Russo, T. L. Hart and A. Schönenberger, September 2002.
- 6 “Banking integration in the euro area” by I. Cabral, F. Dierick and J. Vesala, December 2002.
- 7 “Economic relations with regions neighbouring the euro area in the ‘Euro Time Zone’” by F. Mazzaferro, A. Mehl, M. Sturm, C. Thimann and A. Winkler, December 2002.
- 8 “An introduction to the ECB’s survey of professional forecasters” by J. A. Garcia, September 2003.
- 9 “Fiscal adjustment in 1991-2002: stylised facts and policy implications” by M. G. Briotti, February 2004.
- 10 “The acceding countries’ strategies towards ERM II and the adoption of the euro: an analytical review” by a staff team led by P. Backé and C. Thimann and including O. Arratibel, O. Calvo-Gonzalez, A. Mehl and C. Nerlich, February 2004.
- 11 “Official dollarisation/euroisation: motives, features and policy implications of current cases” by A. Winkler, F. Mazzaferro, C. Nerlich and C. Thimann, February 2004.
- 12 “Understanding the impact of the external dimension on the euro area: trade, capital flows and other international macroeconomic linkages” by R. Anderton, F. di Mauro and F. Moneta, March 2004.
- 13 “Fair value accounting and financial stability” by a staff team led by A. Enria and including L. Cappiello, F. Dierick, S. Grittini, A. Maddaloni, P. Molitor, F. Pires and P. Poloni, April 2004.
- 14 “Measuring financial integration in the euro area” by L. Baele, A. Ferrando, P. Hördahl, E. Krylova, C. Monnet, April 2004.

- 15 “Quality adjustment of European price statistics and the role for hedonics” by H. Ahnert and G. Kenny, May 2004.
- 16 “Market dynamics associated with credit ratings: a literature review” by F. Gonzalez, F. Haas, R. Johannes, M. Persson, L. Toledo, R. Violi, M. Wieland and C. Zins, June 2004.
- 17 “Corporate ‘excesses’ and financial market dynamics” by A. Maddaloni and D. Pain, July 2004.
- 18 “The international role of the euro: evidence from bonds issued by non-euro area residents” by A. Geis, A. Mehl and S. Wredenburg, July 2004.
- 19 “Sectoral specialisation in the EU: a macroeconomic perspective” by MPC task force of the ESCB, July 2004.
- 20 “The supervision of mixed financial services groups in Europe” by F. Dierick, August 2004.
- 21 “Governance of securities clearing and settlement systems” by D. Russo, T. Hart, M. C. Malaguti and C. Papathanassiou, October 2004.
- 22 “Assessing potential output growth in the euro area: a growth accounting perspective” by A. Musso and T. Westermann, January 2005.
- 23 “The bank lending survey for the euro area” by J. Berg, A. van Rixtel, A. Ferrando, G. de Bondt and S. Scopel, February 2005.
- 24 “Wage diversity in the euro area: an overview of labour cost differentials across industries” by V. Genre, D. Momferatou and G. Mourre, February 2005.
- 25 “Government debt management in the euro area: recent theoretical developments and changes in practices” by G. Wolswijk and J. de Haan, March 2005.
- 26 “The analysis of banking sector health using macro-prudential indicators” by L. Mörttinen, P. Poloni, P. Sandars and J. Vesala, March 2005.
- 27 “The EU budget – how much scope for institutional reform?” by H. Enderlein, J. Lindner, O. Calvo-Gonzalez, R. Ritter, April 2005.
- 28 “Reforms in selected EU network industries” by R. Martin, M. Roma, I. Vansteenkiste, April 2005.
- 29 “Wealth and asset price effects on economic activity”, by F. Altissimo, E. Georgiou, T. Sastre, M. T. Valderrama, G. Sterne, M. Stocker, M. Weth, K. Whelan, A. Willman, June 2005.
- 30 “Competitiveness and the export performance of the euro area”, by a Task Force of the Monetary Policy Committee of the European System of Central Banks, June 2005.
- 31 “Regional monetary integration in the member states of the Gulf Cooperation Council (GCC)” by M. Sturm and N. Siegfried, June 2005.

- 32 “Managing financial crises in emerging market economies: experience with the involvement of private sector creditors” by an International Relations Committee task force, July 2005.
- 33 “Integration of securities market infrastructures in the euro area” by H. Schmiedel, A. Schönenberger, July 2005.
- 34 “Hedge funds and their implications for financial stability” by T. Garbaravicius and F. Dierick, August 2005.
- 35 “The institutional framework for financial market policy in the USA seen from an EU perspective” by R. Petschnigg, September 2005.
- 36 “Economic and monetary integration of the new Member States: helping to chart the route” by J. Angeloni, M. Flad and F. P. Mongelli, September 2005.
- 37 “Financing conditions in the euro area” by L. Bê Duc, G. de Bondt, A. Calza, D. Marqués Ibáñez, A. van Rixtel and S. Scopel, September 2005.
- 38 “Economic reactions to public finance consolidation: a survey of the literature” by M. G. Briotti, October 2005.
- 39 “Labour productivity in the Nordic EU countries: a comparative overview and explanatory factors – 1998-2004” by A. Annenkov and C. Madaschi, October 2005.
- 40 “What does European institutional integration tell us about trade integration?” by F. P. Mongelli, E. Dorrucci and I. Agur, December 2005.
- 41 “Trends and patterns in working time across euro area countries 1970-2004: causes and consequences” by N. Leiner-Killinger, C. Madaschi and M. Ward-Warmedinger, December 2005.
- 42 “The New Basel Capital Framework and its implementation in the European Union” by F. Dierick, F. Pires, M. Scheicher and K. G. Spitzer, December 2005.
- 43 “The accumulation of foreign reserves” by an International Relations Committee Task Force, February 2006.
- 44 “Competition, productivity and prices in the euro area services sector” by a Task Force of the Monetary Policy Committee of the European System of Central banks, April 2006.
- 45 “Output growth differentials across the euro area countries: Some stylised facts” by N. Benalal, J. L. Diaz del Hoyo, B. Pierluigi and N. Vidalis, May 2006.
- 46 “Inflation persistence and price-setting behaviour in the euro area – a summary of the IPN evidence”, by F. Altissimo, M. Ehrmann and F. Smets, June 2006.
- 47 “The reform and implementation of the stability and growth pact” by R. Morris, H. Ongena and L. Schuknecht, June 2006.

- 48 “Macroeconomic and financial stability challenges for acceding and candidate countries” by the International Relations Committee Task Force on Enlargement, July 2006.
- 49 “Credit risk mitigation in central bank operations and its effects on financial markets: the case of the Eurosystem” by U. Bindseil and F. Papadia, August 2006.
- 50 “Implications for liquidity from innovation and transparency in the European corporate bond market” by M. Laganá, M. Peřina, I. von Köppen-Mertes and A. Persaud, August 2006.
- 51 “Macroeconomic implications of demographic developments in the euro area” by A. Maddaloni, A. Musso, P. Rother, M. Ward-Warmedinger and T. Westermann, August 2006.
- 52 “Cross-border labour mobility within an enlarged EU” by F. F. Heinz and M. Ward-Warmedinger, October 2006.
- 53 “Labour productivity developments in the euro area” by R. Gomez-Salvador, A. Musso, M. Stocker and J. Turunen, October 2006.
- 54 “Quantitative quality indicators for statistics – an application to euro area balance of payment statistics” by V. Damia and C. Picón Aguilar, November 2006
- 55 “Globalisation and euro area trade: Interactions and challenges” by U. Baumann and F. di Mauro, February 2007.
- 56 “Assessing fiscal soundness: Theory and practice” by N. Giammarioli, C. Nickel, P. Rother, J.-P. Vidal, March 2007.
- 57 “Understanding price developments and consumer price indices in south-eastern Europe” by S. Herrmann and E. K. Polgar, March 2007.
- 58 “Long-Term Growth Prospects for the Russian Economy” by R. Beck, A. Kamps and E. Mileva, March 2007.
- 59 “The ECB Survey of Professional Forecasters (SPF) a review after eight years’ experience”, by C. Bowles, R. Friz, V. Genre, G. Kenny, A. Meyler and T. Rautanen, April 2007.
- 60 “Commodity price fluctuations and their impact on monetary and fiscal policies in Western and Central Africa” by U. Böwer, A. Geis and A. Winkler, April 2007.
- 61 “Determinants of growth in the central and eastern European EU Member States – A production function approach” by O. Arratibel, F. Heinz, R. Martin, M. Przybyla, L. Rawdanowicz, R. Serafini and T. Zumer, April 2007.
- 62 “Inflation-linked bonds from a Central Bank perspective” by J. A. Garcia and A. van Rixtel, June 2007.
- 63 “Corporate finance in the euro area – including background material”, Task Force of the Monetary Policy Committee of the European System of Central Banks, June 2007.

- 64 “The use of portfolio credit risk models in central banks”, Task Force of the Market Operations Committee of the European System of Central Banks, July 2007.
- 65 “The performance of credit rating systems in the assessment of collateral used in Eurosystem monetary policy operations” by F. Coppens, F. González and G. Winkler, July 2007.

ISSN 1607148-4



9 771607 148006