

Task Force of the Market Operations Committee of the European System of Central Banks

Research Report

The use of portfolio credit risk models in Central Banks

ECB Occasional Paper, No. 64

Provided in Cooperation with:

European Central Bank (ECB)

Suggested Citation: Task Force of the Market Operations Committee of the European System of Central Banks (2007) : The use of portfolio credit risk models in Central Banks, ECB Occasional Paper, No. 64, European Central Bank (ECB), Frankfurt a. M.

This Version is available at:

<https://hdl.handle.net/10419/154517>

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 64 / JULY 2007

THE USE OF PORTFOLIO CREDIT RISK MODELS IN CENTRAL BANKS

Task Force
of the Market Operations Committee
of the European System of Central Banks



EUROPEAN CENTRAL BANK

EUROSYSTEM



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

OCCASIONAL PAPER SERIES

NO 64 / JULY 2007

THE USE OF PORTFOLIO CREDIT RISK MODELS IN CENTRAL BANKS

Task Force
of the Market Operations Committee
of the European System of Central Banks

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=977355.



© 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

CONTENTS

1 INTRODUCTION	5
2 CREDIT RISK IN CENTRAL BANK PORTFOLIOS	6
3 CREDIT RISK MODELS	9
3.1 Overview of credit risk modelling issues	9
3.2 Models and parameter assumptions used by task force members	10
3.2.1 Probabilities of default/migration	13
3.2.2 Correlation	16
3.2.3 Recovery rates	18
3.2.4 Yields/spreads	18
3.3 Output	20
4 SIMULATION EXERCISE	22
4.1 Introduction	22
4.2 Simulation results for Portfolio I using the common set of parameters	23
4.3 Simulation results for Portfolio II using the common set of parameters	27
4.4 Sensitivity analysis using individual sets of parameters	30
5 CONCLUSIONS AND LESSONS LEARNED	33
REFERENCES	36
EUROPEAN CENTRAL BANK OCCASIONAL PAPER SERIES	39

TASK FORCE OF THE MARKET OPERATIONS COMMITTEE OF THE EUROPEAN SYSTEM OF CENTRAL BANKS

This report was drafted by an ad hoc Task Force of the Market Operations Committee of the European System of Central Banks. The Task Force was chaired by Ulrich Bindseil. The coordination and editing of the report was carried out by the Secretary of the Task Force, Han van der Hoorn.

The full list of members of the Task Force is as follows:

Ulrich Bindseil	European Central Bank
Han van der Hoorn	
Ken Nyholm	
Henrik Schwartzlose	
Pierre Ledoyen	Nationale Bank van België/Banque Nationale de Belgique
Wolfgang Föttinger	Deutsche Bundesbank
Fernando Monar	Banco de España
Bérénice Boux	Banque de France
Gigliola Chiappa	Banca d'Italia
Noëlle Honings	De Nederlandsche Bank
Ricardo Amado	Banco de Portugal
Kai Sotamaa	Suomen Pankki – Finlands Bank
Dan Rosen	University of Toronto (external consultant)

I INTRODUCTION

In early 2006 nine Eurosystem central banks – the national central banks (NCBs) of Belgium, Germany, Spain, France, Italy, the Netherlands, Portugal and Finland, as well as the European Central Bank (ECB) – established a task force to analyse and discuss the use of portfolio credit risk methodologies by central banks.

The objectives of the task force were threefold. The first was to conduct a stock-taking exercise as regards current practices at NCBs and the ECB. The second followed directly from the first: to share views and know-how among participants. The third was to develop or agree on a “best practice” for central banks on certain central bank-specific modelling aspects and parameter choices. Two common portfolios were analysed by several task force members with different systems and the simulation results were compared.

This report summarises the findings of the task force. It is organised as follows. Section 2 starts with a discussion of the relevance of credit risk for central banks. It is followed by a short introduction to credit risk models, parameters and systems in Section 3, focusing on models used by members of the task force. Section 4 presents the results of the simulation exercise undertaken by the task force. The lessons from these simulations as well as other conclusions are discussed in Section 5.

2 CREDIT RISK IN CENTRAL BANK PORTFOLIOS

Credit risk may be defined as the risk of losses due to credit events, i.e. default (an obligor being unwilling or unable to repay its debt) or a change in the quality of the credit (rating change). Central banks may be exposed to at least two different sources of credit risk. The first is related to policy operations: central banks lend to commercial banks, with the aim of controlling the short-term interest rate. The amount may be very sizable: in 2006 the average amount lent to commercial banks outstanding in the euro area was more than €700 billion. The risk, on the other hand, is relatively small, since all policy-related lending is collateralised.¹ A central bank risks losing money only in the unlikely scenario of a “double default” on the part of the counterparty as well as issuer of the collateral, or in event of a default by the counterparty in combination with a large mark to market loss on the collateral. The latter risk is mitigated by applying haircuts to the collateral. The security from a collateral framework is not absolute – nor should it be: there is a trade-off between security and costs/efficiency of monetary policy implementation (Bindseil and Papadia, 2006) – but deemed sufficient for credit risk from policy operations to be disregarded in this report.

The second source of credit risk is investment operations. Traditionally, central banks have been very conservative investors, with little if any appetite for credit risk. Their investment portfolios have always been very risky on a mark to market basis, though, as a large proportion of assets has been denominated in foreign currency, and currency risk is typically not hedged (it is regarded as “unavoidable”). In addition, large gold holdings are subject to fluctuations in the price of gold. Compared with currency and commodity risks, however, other financial risks in the balance sheet – including credit and interest rate risk – are usually very small. Credit risk is only a minor component of overall financial risks, in particular at lower confidence levels of common

risk measures such as value at risk due to credit risk (CreditVaR). It becomes more relevant when the confidence level is increased, but remains much smaller than exchange rate and gold price risks.

This relatively limited (perceived) relevance of credit risk is changing gradually, for a number of reasons.² First, central bank reserves have been growing rapidly in recent years, in particular in Asia. Some of these reserves may not be directly needed to fulfil public duties (e.g. to fund interventions). At the same time, central banks are feeling increasing pressure to ensure that, within the constraints imposed by their public duties and in an environment of generally decreased interest rates and lower expected returns, an adequate return is nonetheless made on these public assets. Moreover, as demonstrated in Section 4 of this report, even a high credit quality portfolio may show a considerable amount of credit risk once the confidence level of CreditVaR or other tail measures approaches 100%. These observations may be used as arguments for transferring a proportion of central bank reserves into “non-traditional” assets, which offer higher expected returns than more traditional central bank assets, such as sovereign and supranational debt, as well as possibly bonds issued by government sponsored enterprises, at little additional risk. Some of these newer asset classes include asset-backed securities (ABS), mortgage-backed securities (MBS), corporate bonds and, to a lesser extent, equities. A recent description of these trends in central bank reserves management can be found, for instance, in Wooldridge (2006).

1 Article 18.1 of the Statute of the European System of Central Banks and of the European Central Bank requires that Eurosystem lending to banks be based on adequate collateral.

2 In one of their annual surveys of reserve management trends, Pringle and Carter (2005) observe that “The single most important risk facing central banks in 2005 is seen as *market risk* (reflecting expectations of volatility in securities markets and exchange rates). However, large central banks view *credit risk* as likely to be equally if not more important for them as diversification of asset classes increases their exposure to a wider range of borrowers/investments”.

The case for corporate bonds in central bank portfolios has been put by, among others, de Beaufort et al. (2002) and Grava (2004), who focus on the attractive risk-return trade-off of corporate bonds vis-à-vis government debt. Several studies have even argued not only that the expected return on corporate bonds is higher than the expected return on similar government bonds, but that the risk is also lower, as a result of negative correlations between spreads and the level of interest rates (see, for instance, Loeys, 1999). In general, one can argue that in most cases *adding* a small position to an existing portfolio should not change the overall risk level substantially, and that *substituting* existing assets with newer assets that have lower correlations with the rest of the portfolio might even reduce the portfolio risk.

Most central banks within the euro area are already exposed to credit risk through uncollateralised deposits with commercial banks, but only a few central banks invest in corporate bonds. Several others are, however, exploring the possibilities. As credit risk exposure grows, central banks must necessarily invest time and resources in credit risk measurement tools. Value at risk (VaR) models for market risk are now common in most, if not all, central banks. The introduction of portfolio credit risk models is a logical next step, also as a precondition for making credit and market risks more comparable and for making progress towards a more integrated risk management approach. In addition, central banks study credit risk models for reasons unrelated to their investments, notably in their capacity as bank supervisors or for market surveillance.

Only a few central banks have practical experience with credit risk modelling, but many others are testing or implementing systems. Of those represented in the task force, three central banks have an operational system. Their models measure credit risk in all investment portfolios, i.e. foreign reserves as well as domestic fixed income portfolios. Given the portfolio compositions, the scope of the models is restricted to fairly “plain vanilla” instruments

such as bonds, covered bonds, deposits, repos and over-the-counter derivative instruments such as forwards and swaps (but not yet credit default swaps (CDSs)). Government bonds or other bonds that are perceived as credit risk-free are sometimes excluded from the calculations.

These models are used for a variety of purposes, starting with reporting, typically done on a monthly basis. Indirectly, portfolio credit risk models are also used for limit setting, for instance, if the limit structure is designed in such a way that a certain CreditVaR for the whole portfolio is not exceeded. Individual limits, however, are not derived from a CreditVaR. Other applications are limited or still at an early stage. Strategic asset allocation decisions, for example, are not (yet) based on a trade-off between credit and market risk. Risk-return considerations do play a role, however, when assessing the desired allocation to credit. One central bank’s decision to invest in corporate bonds was motivated by the wish to increase portfolio returns by reducing the allocation to Treasuries and, hence, avoiding paying the liquidity premium embedded in Treasury yields. Credit spreads were decomposed into compensations for default risk and for other risks, in order to identify assets with the largest compensation for risks other than default (mainly liquidity risk). At the time, this compensation was found to be in the AA-A range, which is still the bulk of this central bank’s portfolio.

The motivation for implementing a portfolio credit risk model in those NCBs that do not have a model already, is primarily to be able to identify and quantify sources of risk and to be able to reduce them whenever considered necessary. CreditVaR is also expected to facilitate the decision-making process surrounding benchmarks, investment universe and limit system. Another envisaged application of a portfolio credit risk model would be in stress testing. A precondition is that models are transparent and, wherever possible, simple, in

order to be able to communicate output to decision makers.

Ultimately, the aim of some of the banks which have advanced further in this field, as well as of academic research, is to develop a framework for integrated risk management, which would include market as well as credit risk, and possibly also other risks such as liquidity and operational risk. The calculation of tail measures of credit risk is clearly a first key step in this direction, as it provides the same types of risk measure as those used typically for market risks. In the practice of most task force members, there have so far been few concrete attempts to integrate market and credit risk models. One model permits market and credit risk to be combined, using stochastic yield curves. Nevertheless, one of the main (and well-known) complications of integration is the difference in horizon for credit and market risk. Clearly, this is an area that is still underdeveloped, in theory as well as in practice.

3 CREDIT RISK MODELS

3.1 OVERVIEW OF CREDIT RISK MODELLING ISSUES

In recent years, the literature on credit risk modelling has grown tremendously; even a concise summary would be well beyond the scope of this report. Instead, this section focuses on the methodologies used by members of the task force and issues of particular relevance to central banks. For a comprehensive introduction into credit risk modelling, the interested reader is referred to one of the standard textbooks, including Bluhm et al. (2003), Cossin and Pirotte (2007), Duffie and Singleton (2003), Lando (2004) or Saunders and Allen (2002), or papers such as O’Kane and Schlögl (2001). Each of these introduces the topic from a slightly different perspective and with its own level of (mathematical) complexity. A good introduction for practitioners is Ramaswamy (2004).


Broadly speaking, credit risk can be quantified in default or in migration mode. In default mode, the only risk that matters is the risk of default. Mark to market losses due to rating migrations are not taken into account. For high quality portfolios, the credit risk in default mode is very low, simply because very few if any high quality issuers default within the risk horizon, which is typically set at one year. By contrast, migration mode deals with all mark to market gains and losses due to changes in ratings. Default is nothing more than a particular, albeit extreme, example of a rating migration, and therefore default mode can be interpreted as a special case of migration mode. Since, empirically, the probability of a rating downgrade exceeds the probability of an upgrade, and the loss associated with a downgrade typically exceeds the gain from an upgrade, the calculated credit risk in migration mode is usually higher than that in default mode.³ The results of Bucay and Rosen (1999) for an international bond portfolio seem to indicate that in migration mode CreditVaR is around 20-40% higher than in default mode,

although these results depend crucially on the nature of the migration matrix (as well as, to a lesser extent, the recovery rate, credit spreads and the duration of the portfolio). In particular, migration matrices such as those derived by KMV, now Moody’s KMV, (based on expected default frequencies) typically find much higher migration probabilities than those computed by the rating agencies. Consequently, migration risk is more relevant in models that use KMV-type migration matrices (while spread risk, discussed below, is smaller). Most of the models implemented or tested by task force members operate in migration mode and use migration probabilities published by the rating agencies.

A central element of credit risk in migration mode is the change in spreads (and, hence, prices) as a result of rating migrations. Spreads can, however, also fluctuate when ratings remain unchanged. Sometimes spread changes reflect the usual market volatility and are not the result of changes in creditworthiness. This risk is known as spread risk. At other times, however, spreads may widen, for instance, in anticipation of a rating downgrade. This situation would clearly reflect credit risk. In practice, it is not always possible to distinguish between spread risk and credit risk. When spreads change for one issuer only, and the rest of the market remains unchanged, this is a clear indication of credit risk. On the other hand, when all spreads change, this may be a reflection of normal market volatility. However, a general spread widening could also, when the economy is deteriorating, reflect an increase in perceived probabilities of default or downgrade. Because of this definition problem, it is not uncommon to refer to all spread changes that do not follow rating changes as spread risk, and to consider as

³ There are, however, technicalities which may partly offset this result, for instance the fact that in default mode, the potential loss from default may be calculated as the difference between the nominal and the recovery value, whereas in migration mode, the loss due to a downgrade is computed as the difference in market value before and after the downgrade. If the market value before downgrade is lower than the nominal value, then the loss in migration mode could be smaller than in default mode. In practice, these technicalities are small and do not change the conclusion that risk in migration mode should be higher than in default mode.

Chart 1 Comparison of typical market and credit returns



Note: The distributions have identical expected returns. The credit return has more probability mass in the left tail, whereas its upside is limited. Due to the (assumed) symmetry of the market return, the chart suggests that the upside of market returns is higher than of credit returns. This is only true for certain types of “market instruments”, such as equities; it is not true for government bonds.

credit risk only those spread changes that are the consequence of a rating change. This report applies the same distinction and does not focus on spread risk.

It is well known that the return distribution of credit instruments is very asymmetric (or “skewed”) towards losses. This is because losses (as a result of defaults or severe downgrades) are potentially much larger (but have a smaller probability) than gains (yield and upgrades). In addition, defaults tend to be positively correlated, limiting the possibilities of diversification. The return distribution of an individual bond that is held until maturity or default is binomial. At the portfolio level, returns are more symmetric, because losses due to downgrades or defaults are partially offset by gains from upgrades on other bonds. “Tail events” (large losses from defaults) can in most cases be avoided through a semi-active approach whereby bonds are sold after being downgraded below a certain threshold (obviously also at the expense of some upside), mirroring the composition of an index. However, a small probability of a sudden default or downgrade remains, and the return distribution is still to some extent skewed, as well as fat tailed

(Chart 1). In order to quantify these risks properly, a credit risk model is needed.⁴


3.2 MODELS AND PARAMETER ASSUMPTIONS USED BY TASK FORCE MEMBERS

There are several commercial systems available to quantify credit risk, the best known of which are probably CreditManager® (based on the CreditMetrics™ methodology developed by the RiskMetrics Group and formerly J.P. Morgan), Portfolio Manager™ (from KMV), CreditRisk+ (developed by Credit Suisse Financial Products) and CreditPortfolioView (from McKinsey). This report focuses on the CreditMetrics™ methodology⁵, since it is used or being tested

4 Several market participants have argued that the return distribution of a well diversified corporate bond index is not dissimilar from the return distribution of a government bond portfolio. Hence, the index return would be more or less symmetric (see, for instance, Loeys, 1999, or Dynkin et al., 2002) and a special credit risk model might not be needed. This symmetry may be hard to achieve in an actual portfolio, especially if the market itself is not well diversified (as in the euro corporate bond market) since correlations among issuers in the same sector are likely to be higher than with issuers in other sectors. Moreover, corporate bond indices are typically based on market capitalisation, with large exposures to heavily indebted companies, further exacerbating downward risks. Another argument why returns may be skewed is that it may not always be possible to sell a position in a distressed market/company at an acceptable (market) price. So, even if an index return seems fairly symmetric, if it cannot be fully replicated, portfolio return may be more skewed in the event that a downgraded bond continues to underperform after being removed from the index. This “survivorship bias” has been studied, among others by Dynkin et al. (2004), who find that over a period of observation (January 1990 - September 2003) the survivorship bias was small (around 0.5 basis point per month) during the first three months after a downgrade, and even reversed if the bonds were held longer, reflecting a general recovery of downgraded bonds after the initial sell-off. A further argument is that, even if the bulk of the distribution appears normal, the returns in the tail of the distribution, which are most relevant especially to conservative investors such as central banks, can still behave far from normally. Finally, symmetry is only possible if a significant proportion of the portfolio has potential to be upgraded, in order to offset losses from downgrades/defaults. A portfolio with AAA issuers only cannot be upgraded, and so, even though defaults are highly unlikely, its return distribution logically exhibits some skewness.

5 Although the methodologies may superficially seem very different, some well-known comparative studies – including Koyluoglu and Hickman (1998), Gordy (2000), Crouhy et al. (2000) and Kern and Rudolph (2001), all of which compare two or more of the main commercially available models – find similarities among them. Note that several of the (earliest versions of the) models operate in default mode only, and that, as a result, some of the comparisons examined the default component of credit risk only.

Chart 2 Asset value and migration



by most central banks participating in the task force, either directly, using the CreditManager[®] software, or through in-house systems (developed in Matlab[®] or Excel[®]) using a similar methodology. The popularity of CreditManager[®] and its methodology is due to a combination of factors: ease and documentation of the methodology, quality and user-friendliness of the software, the reputation of the RiskMetrics Group and familiarity with some of its other products, and sometimes also cost considerations.

The introduction in this section is largely based on the original Technical Document (Gupton et al., 1997), even though the methodology has been updated and improved since then. The CreditMetrics[™] methodology can be classified as a ratings-based (migration) approach, combined with a structural correlation model. Monte Carlo simulation techniques are applied to generate credit loss distributions. To understand the CreditMetrics[™] methodology at the portfolio level, it is best to start with an individual bond. Consider a bond rated A. In order to generate a loss distribution for this bond over a certain horizon, CreditMetrics[™] uses rating migration probabilities such as those regularly published by the rating agencies. It draws random numbers (asset returns) from a standard normal distribution, which are transformed into simulated ratings at the end of the horizon, in such a way that the migration probabilities in the simulation match the

historically observed rating migration probabilities that are used as inputs to the model. This process is illustrated in Chart 2.

As long as the randomly generated asset return is between the thresholds z_{BBB} and z_{AA} , the simulated rating remains unchanged, but when a threshold is exceeded, the rating changes (up or down, depending on the threshold). Thresholds are set in such a way that the simulated migration probabilities are equal to the empirical (input) probabilities. For instance, if the historical probability of an upgrade to AAA is 1%, then the threshold is set at 2.326 (since $\Pr(X > 2.326) = 0.01$ for a standard normal random variable X). On the basis of the simulated rating, the bond is repriced from the relevant forward curve. This process is repeated many times. Two observations are crucial. First, even though asset returns are drawn from a normal distribution, ratings and therefore bond prices are not. Second, for individual bonds, simulation is not really needed, since (in the limit) the simulated rating distribution equals the empirical (input) distribution. The example here merely serves to introduce the methodology at the portfolio level, where simulation techniques are needed to generate correlated migrations.

A similar procedure is applied to portfolios which consist of more than one obligor, but with the additional complexity that asset returns and therefore rating migrations are correlated. Uncorrelated random returns need to be transformed into correlated returns, which can be done in a number of ways. A well-known technique, available in CreditManager[®], is based on the Cholesky decomposition of the correlation matrix.⁶ The normal distribution of asset returns is merely used for convenience –

6 A correlation matrix Σ is decomposed into an upper triangle and a lower triangle matrix L in such a way that $\Sigma = LL^T$. A vector of uncorrelated random returns x_u is transformed into a vector of correlated returns $x_c = Lx_u$. It is easy to see that x_c has zero mean, because x_u has zero mean, and a correlation matrix equal to $E(x_c x_c^T) = E(Lx_u x_u^T L^T) = LE(x_u x_u^T)L^T = LL^T = \Sigma$, as desired. Since correlation matrices are symmetric and (in theory) positive-definite, the Cholesky decomposition can be computed.

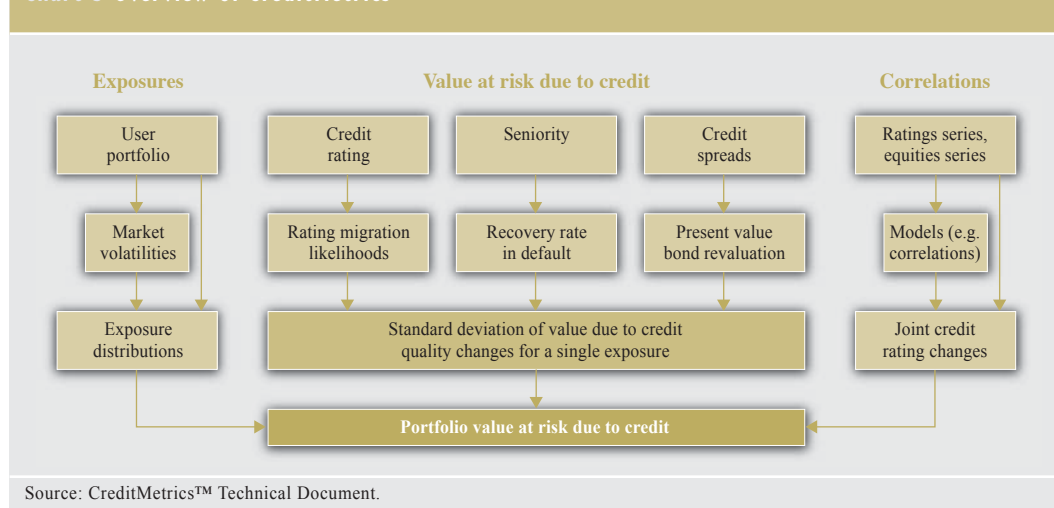
correlation is the only determinant of co-dependence – but in theory it is also possible to use alternative probability distributions for asset returns. These, however, increase the complexity of the model. Sampling from Student-*t* distributions, for example, allows higher tail dependence. Lucas et al. (2001) find that the choice of distribution has significant consequences for the credit loss quantiles, especially far in the tails. Note that using default correlation directly (rather than asset correlation) poses several difficulties, aside from the usual measurement problems and lack of data. Lucas (2004) illustrates why pairwise default correlations are insufficient to quantify credit risk in portfolios consisting of three assets or more. This is a consequence of the discrete nature of defaults.

Under the CreditMetrics™ methodology, credit risk is independent from market (spread) risk. This is because spreads are constant and derived from forward curves. Most members of the task force follow the standard CreditManager® set-up; in-house models sometimes rest on somewhat simplifying assumptions. The approach can be simplified to default mode only, and the number of “ratings” is reduced to two (default/no default only). This may be useful when quantifying the credit risk for non-tradable assets such as deposits, for which

migrations and marking to market are less relevant. Some central banks have “upgraded” their models from default to migration mode fairly recently. One has been testing credit risk models primarily in default mode but has applied migration mode for the simulation exercise in Section 4.

In order to generate reliable estimates of risk (tail) measures, a large number of simulations are needed. The number can be greatly reduced using variance reduction techniques such as importance sampling, which is especially suited to rare event simulations. Importance sampling is based on the idea that one is really only concerned with the tail of the distribution, and will therefore sample more observations from the tail than from the rest of the distribution. With importance sampling, the original distribution from which observations are drawn is changed into a distribution which increases the likelihood that “important” observations are drawn. These observations are then weighted by the likelihood ratio to ensure that estimates are unbiased. The challenge is finding a good transformation of the original distribution, which is an art as well as a science. For a normal distribution this transformation is technically straightforward and involves shifting the mean (and sometimes also scaling the variance).

Chart 3 Overview of CreditMetrics™



CreditManager[®] also uses importance sampling. A good reference is Glasserman (2005).

Several task force members reported that importance sampling can reduce the number of simulations and, hence, computation time, by a factor of 10 or more. However, it was also noted that the likelihood ratio, which adjusts the likelihood of the drawn outcomes to reflect their likelihood under the original distribution, can be unstable, thus reducing the accuracy of simulation results. Most of the results presented in Section 4 are derived from 100,000 to 200,000 simulated scenarios with importance sampling, which can be completed on most computers in a reasonable amount of time (typically a few minutes using CreditManager[®]). In practice, the number of draws needed to reach a certain precision depends crucially on the composition of the portfolio as well as the chosen confidence level.

The CreditMetrics[™] framework is summarised in the well-known Chart 3. More details can be found in its Technical Document (Gupton et al., 1997). The following sub-sections discuss key parameters in CreditMetrics[™] and related methodologies.

3.2.1 PROBABILITIES OF DEFAULT/MIGRATION

It is important to realise that CreditMetrics[™] is not a methodology to estimate probabilities of default (PDs). Instead, these probabilities, together with migration probabilities, are important input parameters, usually obtained from one of the major rating agencies, which publish updated migration matrices frequently.⁷ The migration matrices from different rating agencies are all fairly similar for any given industry. Each of the three major rating agencies is used by at least one of the task force members, sometimes mixing migration matrices from different sources. One central bank uses default probabilities discussed in Ramaswamy (2004), which are based on Moody's data. All measure probabilities over a one-year horizon.

Rating migration probabilities have their limitations, in particular for central banks

whose portfolios are dominated by highly rated sovereign issuers. It is well-known that default and migration probabilities for sovereign issuers are different from probabilities for corporate issuers. Comparing, for instance, the latest updates of migration probabilities by Standard & Poor's (2007a and 2007b) reveals that while, historically since 1981, a few AA and A corporate issuers have defaulted over a one-year horizon (with frequencies equal to 1 and 6 basis points respectively, see Table 13 of S&P 2007a), not a single investment grade (i.e. down to BBB) sovereign issuer has ever defaulted over a one-year horizon (based on observations since 1975, see Table 1 of S&P 2007b). Even after ten years, A or better rated sovereign issuers did not default (Table 5 of S&P 2007b). While these are comforting results, one should also be aware that they are based on a limited number of observations. Hence, their statistical significance may be questioned. Moreover, the rating agencies themselves acknowledge that rating sovereign issuers is considerably more complex and subjective than rating corporate issuers.

7 Default and migration probabilities can also be (and often are) estimated from structural and reduced form models, among others. Structural models are based on the work of Merton (1974), and apply the logic that equity represents a call option on a firm's assets. Debt can be modelled as a short put option, and so option pricing techniques can be applied to value debt and estimate the probability of default. The value of assets is represented by a stochastic process (typically geometric Brownian motion, whereby logarithmic changes in the asset value are normally distributed), based on the assumption that a firm defaults if the value of its assets falls below (the nominal value of) its liabilities. The best known application of this model was developed by KMV, which links "expected default frequency" and "distance to default" to risk-neutral default probabilities. An advantage of structural models over other models is that they can help explain why a company is likely to default. They are, however, less suitable for sovereign issuers or private companies, since the volatility of equity prices is often used to estimate asset volatility. Reduced form models, by contrast, do not try to explain why a firm defaults but exploit information from bond markets to calculate default probabilities or, more precisely, the expected time until default. Default is treated as an unexpected event, the likelihood of which is governed by a default-intensity process. The default intensity measures the conditional likelihood that an issuer will default over the next small interval of time, given that it has not yet defaulted. The parameter (intensity or hazard rate) of this process can be estimated from credit spreads. The simplest example of this approach uses a Poisson process, whereby the time until default is exponentially distributed. Reduced form models are preferred for pricing and hedging credit derivatives.

As a result, investors, including central banks, often use migration probabilities derived from corporate issuers, which leads to conservative but probably more robust risk estimates. But even corporate default probabilities over a one-year horizon are historically equal or close to zero for the highest ratings. Since it seems reasonable to assume that the “true” probabilities are somewhat higher, even for AAA-rated issuers, it is not uncommon for these default probabilities to be adjusted upwards by a few basis points (and for one or more other migration probabilities to be reduced by the same amount). In the absence of better alternatives, this is often done in a rather ad hoc manner. A promising approach, recently proposed by Pluto and Tasche (2006), that derives confidence intervals for PDs, taking into account the number of observations, has not yet found its way to the models used by market participants and task force members.

Task force members apply various adjustments that assign a positive PD to the highest ratings while still respecting the ranking of ratings (i.e. the PD for a AA obligor should be higher than the PD for AAA, etc.). For AAA-rated issuers, the PD is set in the range 0-1 basis point; for AA, it is in the range 0-4 basis points. Note that the upper bounds correspond to the “normalised” PDs in Ramaswamy (2004, Exhibit 5.4, where the 4 basis points is applied to Aa3/AA–). The levels are, however, not based on any empirical evidence; they are merely introduced as a pragmatic solution to allow default correlations to be estimated directly. Sometimes a higher PD is assumed for corporate issuers than for government issuers with the same rating. For example, one task force member assumes that the PD for sovereign and supranational issuers is half that of the PD for corporate issuers with the same rating.

Clearly, the accuracy of the migration probabilities published by the rating agencies is crucial. Their methodology for estimating these probabilities can be described as statistical: the main technique, the “cohort” approach, simply counts the number of migrations for a given

rating within a calendar year and divides this number by the total number of obligors with the initial rating.

Default probabilities for short horizons

An interesting and highly relevant problem, particularly for central bank portfolios with low durations, is how to compute default probabilities for short horizons. When assets mature before the end of the risk horizon (typically one year), then it obviously matters how the expected cash flow at maturity is reinvested. If it were invested in a similar asset from the same obligor at all times, even after a downgrade, then the risk would be identical to a one-year investment. Sometimes this may be a realistic assumption, for instance when a strong relationship with the obligor outweighs increased counterparty risks. It is, however, more common that after a downgrade beyond a certain threshold the cash from the matured asset is reinvested elsewhere. Hence, CreditMetrics™ assumes that matured assets are held in risk-less cash until the end of the horizon. In these cases, the risk of the short maturity asset is lower than the risk of a longer-term position in the same obligor, and it is necessary to scale default probabilities to short horizons. Note that migration risk is irrelevant for instruments with a maturity less than the horizon, since time is assumed to be discrete and positions can only change at the end of the horizon.

Scaling default probabilities to short horizons can be done in several ways. The easiest approach is to assume that the conditional PD (or “hazard rate” in reduced form models) is constant over time. The only information needed from the migration matrix is the right-hand column which contains the probabilities of default over the risk horizon. Assuming the risk horizon is one year, then for each rating the probabilities of default $pd(t)$ for a shorter horizon $t < 1$ follow directly from the one-year probabilities $pd(1)$ using the formula $pd(t) = 1 - [1 - pd(1)]^t$.⁸ This is approximately equal to

⁸ This is the discrete-time equivalent of reduced form models with a constant hazard rate (conditional probability of default) λ , where the probability of default over a period t is given by $1 - e^{-\lambda t}$. From $pd(1) = 1 - e^{-\lambda}$, it follows that $\lambda = -\ln[1 - pd(1)]$, and therefore that $pd(t) = 1 - [1 - pd(1)]^t$.

$pd(1) \times t$. Note that this would be a logical procedure in *default mode*.

Alternatively, one may use all the information embedded in the migration matrix, taking into account that default probabilities are not constant over time but increase as a result of downgrades. Ideally, if \mathbf{M} is the one-year migration matrix, and one is interested in one-month probabilities of default, a one-month migration matrix \mathbf{G} is needed, such that $\mathbf{G}^{12} = \mathbf{M}$. Essentially, this involves computing the root of the migration matrix. Finding this root requires the computation of eigenvalues and eigenvectors. Any $n \times n$ matrix has n (not necessarily distinct) eigenvalues and corresponding eigenvectors. If the matrix is symmetric, then all eigenvalues are real. If \mathbf{C} is the matrix of eigenvectors and $\mathbf{\Lambda}$ is the matrix with eigenvalues on the diagonal and all other elements equal to zero, then any symmetric matrix \mathbf{M} can be written as $\mathbf{M} = \mathbf{C}\mathbf{\Lambda}\mathbf{C}^{-1}$ (where \mathbf{C}^{-1} denotes the inverse of matrix \mathbf{C}). In special cases, a non-symmetric square matrix (such as a migration matrix) can be decomposed in the same way. The square root of the matrix follows from $\mathbf{M}^{1/2} = \mathbf{C}\mathbf{\Lambda}^{1/2}\mathbf{C}^{-1}$. Migration matrices for shorter periods are found analogously. The computation of the root is based on the Markovian property of migration matrices, which means that rating migrations are path-independent and the probabilities are constant over time. This is a very common assumption, used by many, despite empirical evidence to the contrary (see, for instance, Nickell et al., 2000).

The root of a matrix can only be computed if all of its eigenvalues are non-negative. The eigenvalues of a migration matrix are in practice usually positive – although there is no guarantee that they always will be – because migration matrices are diagonally dominant (i.e. the largest probabilities in each row are on the diagonal). A more serious problem, however, is that some of the eigenvectors can have negative elements and generate a root matrix which also has negative elements. Clearly, in such cases, the root is no longer a valid migration matrix.

In fact, it can be shown that if there are ratings r_1 and r_2 such that r_2 is accessible from r_1 , while the probability of migrating from r_1 to r_2 in a single period is zero, then the root is not a valid migration matrix (Kreinik and Sidelnikova, 2001). Unfortunately, this is precisely the structure of most migration matrices that are based on empirical data, as the one-period PD for AAA is typically zero, while the probability over longer periods is clearly higher.

Note that a transformation is only needed if the horizon of default probabilities exceeds the maturity of the shortest asset in the portfolio. Clearly, it would be more efficient to estimate short horizon PDs directly from a ratings database. This can be done in discrete as well as in continuous time. In the limit, as the time interval approaches zero, migration probabilities can be represented by a generator matrix \mathbf{G} , from which the actual migration probabilities over horizon t are derived by computing the matrix exponential $\exp(t \times \mathbf{G})$ (Lando and Skødeberg, 2002). The estimation of generator matrices takes into account the exact timing of each rating migration and therefore uses more information than traditional approaches. A positive spin-off of using generator matrices is that they normally also generate positive probabilities of default for the highest rated issuers, so that fewer manual (and arbitrary) adjustments are needed. They do not solve the limited data availability as regards sovereign issuers, however. Generator matrices are not (yet) very common in practice.

If the root of the migration matrix does not exist, or if it is not a valid migration matrix, then an approximation is needed for the PD over short horizons. The central banks participating in the task force use various approximations for this. A standard approach in CreditManager® is to “scale down” the annual PDs linearly, for example the one-month PD is set equal to the annual probability divided by 12. As noted before, this approach is approximately equal to the “true” formula in default mode. The approach is used by several central banks if the root matrix cannot be found

or is not a valid migration matrix. One central bank transforms maturities into multiples of three months and uses the adjusted maturity to compute the PD. Another assumes that any asset which matures before the end of the horizon is rolled into a similar asset, which implies that the PD of a short duration asset equals the one-year PD. One task force member has recently started computing the “closest three-month matrix generator” to the one-year matrix. This generator is calculated numerically by minimising the sum of the squared differences between the original one-year migration probabilities and the one-year probabilities generated by the three-month matrix. This three-month matrix provides plausible estimates of the short-term migration probabilities and will normally also generate small but positive one-year default probabilities for highly rated issuers.

It seems fair to say that, by most standards, all of these approximations lead to conservative estimates of the “true” short-term PD. In the structural models of default, for instance, the stochastic properties of the asset value imply a probability that is very close to zero over short horizons, since a “jump” in the asset value is not possible and time passes before the default threshold is reached with any significant probability. In the reduced form models, default is an unforeseeable event and so will have a positive probability even over shorter horizons, but, unless very unusual parameter choices are made, the probability will not be higher than the probabilities assumed above.

In reality, the conservative estimates of default probabilities find some justification in the fact that most central banks, like any other investment grade investor, sell a bond once it has been downgraded beyond a certain threshold. This reduces the actual PD, but its impact cannot be addressed directly by single-step models, which do not allow selling before the end of the horizon. One member of the task force has adopted a multi-step approach, whereby 12 monthly sub-periods are simulated

and the model allows downgraded bonds to be sold at the end of each sub-period.

3.2.2 CORRELATION

Probabilities of default and migration are key risk parameters as regards individual obligors, together with recovery rates, which will be discussed in the next sub-section. At the portfolio level, correlation is crucial. Under certain assumptions, the PD and correlation determine the entire loss distribution of a credit-risky portfolio.⁹ Correlation is, however, also the parameter that is most difficult to estimate. Correlation measures the extent to which assets default or migrate together. In the credit risk literature, the parameter often (but loosely) referred to is default correlation, formally defined as the correlation between default indicators (1 for default, 0 for non-default) over some period of time, typically one year. Default correlation can be either positive – for instance because firms in the same industry are exposed to the same suppliers or raw materials, or because firms in one country are exposed to the same exchange rate – or negative, when for example the elimination of a competitor increases a company’s market share. Default correlation is difficult to estimate directly, simply because defaults, let alone correlated defaults, are rare events. It is also, as mentioned before, difficult to apply in practice. For these reasons, CreditMetrics™ (and many other models) estimates correlations of asset returns rather than of defaults.


CreditMetrics™ uses equity returns as a proxy for asset returns, which cannot be observed directly or only infrequently. This is a common approach, used by many others. Rather than using one uniform asset correlation, CreditMetrics™ allows a factor model to be used for correlations. The model is estimated

⁹ A well-known result by Vasicek (1991) is that the cumulative loss distribution of an infinitely granular portfolio in default mode (no recovery) is in the limit equal to:

$$F(x) = N\left(\frac{\sqrt{1-\rho}N^{-1}(x) - N^{-1}(\rho d)}{\sqrt{\rho}}\right),$$

where ρ is the (positive) asset correlation and $N(x)$ represents the cumulative standard normal distribution evaluated at x (N^{-1} being its inverse).


Chart 4 Range of possible default correlations for a given asset correlation (30%)



Source: CreditMetrics™ Technical Document.

Chart 5 Impact of asset correlation on portfolio risk

(hypothetical portfolio with 100 issuers rated AAA-A, confidence level 99.95; percentages of portfolio market value)



Source: ECB calculations.

on equity indices, with individual obligors “mapped” onto countries and sectors. Because of uncertainties and strong assumptions in the computation, CreditManager® also allows users to select their own, possibly uniform, asset correlations.

It is important to note that asset and default correlation are very different concepts. Default correlation is related non-linearly to asset correlation and tends to be considerably lower (in absolute value).¹⁰ While Basel II, for instance, proposes an asset correlation of up to 24%¹¹, default correlation is normally only a few percent. Indeed, Lucas (2004) demonstrates that, for default correlation, the full range of -1 to +1 is only attainable under very special circumstances. Chart 4 below illustrates the range of possible default correlations for a given asset correlation (30%). Note that, for a given level of asset correlation, default correlation is a (generally increasing) function of the individual probabilities of default.

Other things being equal, risks become more concentrated as asset correlations increase, and the probability of multiple defaults or downgrades rises. With perfect correlation among all obligors, a portfolio behaves as a single bond. It should thus come as no surprise that the relationship between asset correlation

and credit risk is positive (and non-linear). Chart 5 plots this relationship, using expected shortfall (see Section 3.3) as the risk measure, for a hypothetical portfolio.

In practice, it is not possible to estimate and use individual correlations for each pair of obligors. First of all, scarcity of data limits the scope for estimating correlations, and second, the large number of correlations ($n(n-1)/2$ for a portfolio of n obligors) makes this approach untenable. Instead, it is common to use industry and country correlations, or simply to assume one uniform correlation. Task force members use various asset correlations when computing CreditVaR. Several use a fixed and uniform correlation equal or very close to the Basel II level of 24%. Others prefer the CreditMetrics™ factor model, which maps obligors to one or more country and industry indices and estimates asset correlations from equity indices, because

10 The formal relationship between asset and default correlation depends on the joint distribution of the asset returns. For normally distributed asset returns, the relationship is given by equations 8.5 and 8.6 in the CreditMetrics™ Technical Document.

11 Under the internal ratings-based approach of Basel II, the formula for calculating risk-weighted assets is based on an asset correlation ρ equal to $\rho = w \cdot 0.12 + (1-w) \cdot 0.24$, where $w = \frac{1 - e^{-50 \cdot pd}}{1 - e^{-50}}$. Notice that ρ decreases as pd increases, which seems to contradict Chart 4. Note, however, that Chart 4 plots *default* correlation (for a given asset correlation), whereas the Basel II formula computes *asset* correlation.

it captures diversification effects between industries and countries. Note that this approach has its limitations for central bank portfolios, which mainly consist of bonds issued by (unlisted) governments. An approximation used maps the government issuer to a broad country equity index and estimates the R^2 with a general (i.e. world) index. While the magnitude of correlations in the CreditMetrics™ factor model obviously depends on the portfolio composition, they tend to be larger than 24% in a typical central bank portfolio dominated by government and other AAA bonds. In fact, some correlations are considered high enough to justify setting them at a discretionary (and conservative) level of 100%.

3.2.3 RECOVERY RATES

The recovery rate measures the proportion of the principal value (and possibly accrued interest) that is recovered in the event of a default. The recovery rate depends on, among other things, the seniority of a loan. For simplicity, the recovery rate is often assumed to be constant across (types of) issuers and issues. More sophisticated approaches use stochastic recovery rates (typically using a beta distribution), possibly even correlated with default/migration probabilities. Clearly, the impact of the recovery rate on estimated losses is significant, particularly in default mode. Well-known empirical studies into recovery rates are from Asarnow and Edwards (1995), Carty and Lieberman (1996), Altman and Kishore (1996), and Altman, Resti and Sironi (2005). Rating agencies also publish studies on recovery rates regularly.

The members of the task force use several alternative assumptions for the recovery rate. Some are taken from papers mentioned in the CreditMetrics™ Technical Document (which are among those cited above). When a fixed recovery rate is used, it is typically set in the range 40-50% for senior bonds. Also, the mean is in this range when a stochastic recovery rate is modelled. One example is a stochastic recovery rate with a beta distribution with a mean of 48% and standard deviation of 26%

(for senior unsecured bonds). It was mentioned that when recovery rates from CreditManager® are used, typically the most conservative levels are selected.

Recently, more evidence has emerged that recovery rates for bank loans are on average substantially higher than for bonds. In response, Moody's (2004) announced a revision of its rating methodology, which is based on expected losses. S&P ratings, by contrast, are based on PDs and do not take into account recovery rates.

A related concept is that of the exposure at default, which may be different from the current exposure as a result of market movements or accrued interest. In CreditMetrics™, exposure at default is deterministic; in practice, one may use the current exposure (possibly plus an add-on) or the expected exposure at the investment horizon. One system can also generate stochastic yield curves, but concepts such as potential future exposure are not (yet) used. Long and short positions versus individual counterparties are sometimes netted.

Some participants consolidate exposures to related counterparties at the group level, and assume that the PD of the group equals the PD of the member with the lowest rating. Sometimes counterparties with close links are connected indirectly. As many branches don't have individual ratings, limits are assigned at the group level.

3.2.4 YIELDS/SPREADS

The final parameter, which is only needed in migration mode, is the (forward) spread in yields or (zero-coupon) interest rates. This determines the mark to market loss (gain) in the event of a downgrade (upgrade). In essence, a bond that is downgraded is repriced against the curve for the new rating, and the credit loss is approximately equal to the spread widening multiplied by the modified duration. The quality of the spread is thus crucial for the quantification of credit risk under migration mode. Finding reliable data may be a challenge, in particular

for lower ratings and certain currencies. A certain “smoothness” in the spreads is desired, to avoid a bias in simulation results due to a few outliers in bond prices/yields. For this, a large number of curve fitting techniques are available; see Bank for International Settlements (BIS) (2005) for an overview. An issue for central banks with short duration portfolios is that the quality of the spreads at the short end of the curve is more important than the quality at the long end.

The CreditMetrics™ Technical Document is not very specific as regards its curve methodology, although it mentions various data contributors and serious efforts to ensure accuracy and consistency of curves and spreads. Some recent publications by the RiskMetrics Group shed more light on how this may be done. Stamicar (2007) discusses a “spread overhaul”, largely based on the Hull-White framework (see Hull and White, 2000), which

harmonises methodologies across RiskMetrics products. Rather than using spreads directly to reprice assets upon rating migrations, this framework derives term structures of risk-neutral default probabilities for each rating from either bond or equity prices or from CDS spreads.¹² Together with the risk-free rate, these probabilities are an alternative way to price each bond in the portfolio. Hazard rates

¹² The intuition behind risk-neutral default probabilities comes from option pricing and risk-neutral valuation. The price P of a credit-risky bond with maturity t can be obtained in two ways. The first, traditional approach derives the present value of all cash flows, discounted at the relevant rate. Assuming, for simplicity, only one cash flow at time t , the price is given by $P = [pd \times RR + (1 - pd) \times N] / (1 + rf + s)^t$, where pd is the actual probability of default, RR is the recovery rate, N is the principal (+ interest), rf is the risk-free rate and s is the spread (“risk premium”). Alternatively, the price may also be computed as $P = [q \times RR + (1 - q) \times N] / (1 + rf)^t$. Here, q is the risk-neutral default probability. Note that the spread or risk premium is omitted from the denominator. Since the price is given, it follows that $q > pd$: risk-neutral default probabilities are (much) larger than actual default probabilities.

Table 1 Summary of key parameters in CreditVaR models

	CB1	CB2	CB3	CB4	CB5
PD/migration	Source: Moody's, PD adjusted upwards for AAA, AA and A	Source: mix of Fitch, S&P and Moody's	Source: S&P, PD adjusted upwards for AAA and possibly lower ratings, different for government and non-government	Source: Ramaswamy (= Moody's), PD adjusted upwards for AAA	Source: S&P, PD adjusted upwards for AAA and AA
Assets with maturity below one year	Timing of default uniformly distributed across the year (e.g. annual PDs divided by 4 to obtain quarterly PD)	“Closest three-month matrix generator”	Annual PDs divided by 4 to obtain quarterly PD, maturity rounded upwards to multiples of three months	PD assumed equal to annual PD, based on assumption that matured asset is rolled into similar asset	Monthly PD based on assumption of constant conditional PD
Correlation	Asset correlation fixed (25%) or estimated from industry and country indices (CreditMetrics™ factor model); in-house system uses fixed correlation	Asset correlations estimated from factor model based on correlation of industry and country indices; fixed for certain issuers	Asset correlations estimated from country and industry equity indices	Asset correlation fixed at 24%	Asset correlation fixed at 24%
Recovery rate	Fixed or variable; parameters from CreditManager® based on study by Altman and Kishore	Parameters from CreditManager®; most conservative option chosen per instrument type and seniority	Beta distribution; parameters from CreditManager® based on study by Carty and Lieberman	Fixed at 45% (based on several studies from rating agencies)	Fixed at 40%

(conditional risk-neutral PDs) and forward risk-free rates are assumed to be constant between adjacent curve notes, which has an effect similar to imposing smoothness on spreads directly through some curve fitting technique. Here, it suffices to note that the “overhaul” has little impact on plain vanilla instruments typically used in central bank portfolios, but that it may improve accuracy for more complex structured products. Using credit spreads directly is therefore still a valid approach.

The spreads used by task force members are obtained from either CreditManager® directly or from Bloomberg and Reuters. One member tests CreditVaR in default mode, but plans to extend the model to include migration risk.


The main parameter choices of task force members using or implementing CreditVaR models are summarised in Table 1, where CB1 to CB5 refer to the five central banks with models that have been implemented or are being implemented.

3.3 OUTPUT

Typical output from credit risk models includes expected and unexpected loss, (Credit) value at risk (in the remainder of this report simply referred to as VaR, unless confusion with other types of risk could arise) and expected shortfall (ES) (Chart 6). Expected and unexpected losses are the first and second moments (mean and standard deviation) of the loss distribution and can be calculated analytically. Expected portfolio loss is simply equal to the weighted average of expected losses on individual positions. The analytical computation of unexpected loss is more cumbersome and involves correlations. Sometimes it is more efficient to derive unexpected loss by simulation. Strictly speaking, expected loss is not a risk measure, since risk is by definition restricted to unexpected events.

Like VaR for market risks, CreditVaR is defined as a certain quantile of the credit loss

Chart 6 Return distribution and credit risk measures



distribution. It measures the loss that is not exceeded at a given confidence level over a given time period. In other words, it is the minimum loss that may be suffered with a certain probability. In credit risk modelling, it is common to refer to VaR as the loss *in excess of* the expected loss. Expected shortfall, sometimes also referred to as conditional VaR or expected tail loss, measures the loss in the tail of the distribution, conditional on the fact that the loss exceeds the VaR. It can be calculated as the average VaR at higher confidence levels, and is therefore equal to the average loss with a certain probability. In Chart 6, ES is represented by the surface under the distribution to the left of the VaR.

It is often argued that ES is more appropriate than VaR for the analysis of rare events (such as default). CreditVaR is typically also computed at higher confidence levels than VaR for market risk. This is because issuers with an external rating need to have very low probabilities of default if they aim for a high rating. If, for instance, a bank aims at a single-A rating, corresponding roughly to a PD of 10 basis points, then it should calculate its VaR at a 99.9% confidence level to determine its capital needs. The same confidence level is used in the Basel II formulas for the internal ratings-based approach for credit risk, whereas “only” a 99% confidence level is applied to determine the capital requirements for market risk (Basel Committee on Banking Supervision, 2006).

In addition, the use of VaR is increasingly criticised because VaR is not a coherent risk measure (Artzner et al., 1999)¹³, since it is not necessarily sub-additive. This means that it is possible to construct two portfolios, A and B , such that $\text{VaR}(A + B) > \text{VaR}(A) + \text{VaR}(B)$. In other words, the VaR of a combined portfolio may exceed the sum of the individual VaRs, thus discouraging diversification. Naturally, a risk measure that rewards diversification would be preferable. The sub-additivity problem is particularly acute for portfolios with fat-tailed or discrete return distributions, such as credit-risky portfolios. By contrast, it can be shown that ES is always sub-additive, and also satisfies all other properties of coherent risk measures. Nevertheless, in communication to senior management, VaR still plays a pivotal role, as it is clearer and more comprehensible than ES. Moreover, any risk measure that tries to capture the whole loss distribution in a single number, whether it is VaR or ES, has its limitations. It therefore makes sense to analyse several risk measures at the same time, or in fact use the full return distribution. Finally, it is noted that in the limit, as the confidence level is increased to very high levels, VaR and ES converge. To understand why, note that at a confidence level of 100%, all issuers with a positive PD default, and the VaR as well as ES are equal to the loss given default.

CreditManager[®] and the other systems mentioned compute all of these risk measures. CreditManager[®] offers two definitions of expected loss, both of which are computed analytically. The first (“expected loss” in CreditManager[®]) is equal to the difference between the portfolio value at the start of the simulation (“current value” in CreditManager[®] terminology) and the average portfolio value (over all scenarios) at the end of the simulation horizon (“mean horizon value”). The other definition (“expected loss from horizon value”) equals “horizon value” (if the rating stays the same) minus “mean horizon value”. The difference between the two definitions is that the first is “biased” by interest returns, and can actually be a net gain if default and downgrade

probabilities are small. Hence, all task force members using CreditManager[®] prefer the second definition. Those who do not use CreditManager[®] derive expected loss by simulation and use somewhat different definitions, one of which resembles the first CreditManager[®] definition.

Other risk measures used by task force members include unexpected loss (UL, “standard deviation of horizon value” in CreditManager[®]), VaR (including incremental VaR) and ES at various confidence levels. All of these are derived by simulation. In practice a subset of these is used.

13 A risk measure is said to be coherent if it satisfies the four properties of sub-additivity, (positive) homogeneity, monotonicity and translation invariance.

4 SIMULATION EXERCISE


4.1 INTRODUCTION

The core of this report consists in the analysis of several simulation exercises with the aim of comparing results and quantifying, albeit roughly, the sensitivity of output to changes in parameters. To this end, two very different portfolio have been analysed. The first portfolio (in the following “Portfolio I”) is a subset of the aggregate ECB US dollar portfolio, as it existed some time ago. The portfolio contains government bonds, bonds issued by the BIS, government-sponsored enterprises (GSEs) and supranational institutions, all rated AAA/Aaa, and short-term deposits with 32 different counterparties rated A or higher and with an assumed maturity of one month. Hence, the credit risk of the portfolio is expected to be low. The modified duration of the portfolio is low.

The other portfolio (“Portfolio II”) is fictive. It contains 62 (mainly private) issuers, spread across regions, sectors, ratings as well as maturities. It is still relatively “chunky”, in the sense that the six largest issues make up almost 50% of the portfolio, but otherwise more diversified than Portfolio I. It has a higher modified duration than Portfolio I. The lowest rating is B+/B1. Chart 7 compares the composition of the two portfolios, by rating as well as by sector (the sector “banking” includes positions in GSEs). From the upper chart (distribution by rating), one would expect Portfolio II to be more risky.

Five task force members participated in the simulation exercise. It is recalled that not all had already fully implemented a portfolio credit risk system. Most participants in the simulation exercise analysed the portfolios using at least two sets of parameters, a common set to be used by all participants and one or more sets of individual model parameters. Simulation results were reported using a common template, which included, among other things, the following risk measures: expected loss, unexpected loss, VaR and ES, at various confidence levels and

Chart 7 Comparison of portfolios by rating and by industry



all for a one-year investment horizon. In addition, the probability of at least one default was computed by some participants, since a default might have reputational consequences for a central bank invested in the defaulted company (see also Section 5). However, since the latter is not standard output from any of the systems used, participants had to recourse to ad hoc solutions for computing this statistic, and the numbers should be treated with care.

4.2 SIMULATION RESULTS FOR PORTFOLIO I USING THE COMMON SET OF PARAMETERS

The first simulations were conducted on the basis of a common set of parameters. The results provide a starting point for the scenario analysis in Section 4.4 and can also be used to spot differences in modelling assumptions for parameters not prescribed by the parameter set, in particular short horizon PDs. The common

set includes a fixed recovery rate (40%) and a uniform asset correlation (24%). The credit migration matrix (Table 2) was obtained from Bucay and Rosen (1999) and is based on S&P ratings, but with default probabilities for AAA and AA revised upwards (from 0) as in Ramaswamy (2004, AA set equal to AA-). Spreads were derived from Nelson-Siegel curves (Nelson and Siegel, 1987), where the zero-coupon rate $r(t)$ for maturity t (in months) is given by $r(t) = \beta_1 + (\beta_2 + \beta_3) \frac{1 - e^{-\lambda t}}{\lambda t} - \beta_3 e^{-\lambda t}$. The curve parameters are shown in Table 3.

Note that under this common scenario set, individual assumptions were still needed for a number of parameters. The list includes the computation of the mark to market gain/loss in the event of a rating migration (linear approximation using the modified duration versus full revaluation), the number of simulation runs and whether or not to use variance reduction techniques. A key parameter left to the participants was how to apply annual default probabilities to short duration positions (mainly deposits).

Table 4 displays the simulation results, expressed as a percentage of market value, for Portfolio I, based on the common set of parameters. For each confidence level, the highest VaR and ES are displayed in italics.

The starting point for the analysis of Table 4 is the validation of the models, using an analytical approximation for expected loss. Recall from Section 3.3 that not every participant uses the same definition of expected loss. In absolute terms, all participants reported similar expected losses (i.e. very close to 0). Ignoring, for simplicity, time decay, it is easy to validate these results analytically. Approximately 80% of the portfolio is rated AAA, 17% has a rating of AA and the remaining 3% is rated A. If one multiplies these weights by the PDs (1, 4 and 10 basis points, respectively) and the loss given default (i.e. one minus recovery rate), then the expected loss in default mode and assuming a one-year maturity of deposits would be $(0.80 \times 0.0001 + 0.17 \times 0.0004 + 0.03 \times 0.0010) \times 0.6 = 1.1$ basis points. In migration mode, the expected loss would be somewhat higher, but

Table 2 Common migration matrix (one-year migration probabilities)

(percentages)								
To \ From	AAA	AA	A	BBB	BB	B	CCC/C	D
AAA	90.79	8.30	0.70	0.10	0.10	-	-	0.01
AA	0.70	90.76	7.70	0.60	0.10	0.10	-	0.04
A	0.10	2.40	91.30	5.20	0.70	0.20	-	0.10
BBB	-	0.30	5.90	87.40	5.00	1.10	0.10	0.20
BB	-	0.10	0.60	7.70	81.20	8.40	1.00	1.00
B	-	0.10	0.20	0.50	6.90	83.50	3.90	4.90
CCC/C	0.20	-	0.40	1.20	2.70	11.70	64.50	19.30
D	-	-	-	-	-	-	-	100.00

Source: Bucay and Rosen (1999), PD for AAA and AA adjusted as in Ramaswamy (2004).

Table 3 Parameters for Nelson-Siegel curves

	AAA	AA	A	BBB	BB	B	CCC/C
λ	0.0600	0.0600	0.0600	0.0600	0.0600	0.0600	0.0600
β_1 (level)	0.0660	0.0663	0.0685	0.0718	0.0880	0.1015	0.1200
β_2 (slope)	-0.0176	-0.0142	-0.0149	-0.0158	-0.0242	-0.0254	-0.0274
β_3 (curvature)	-0.0038	-0.0052	-0.0061	-0.0069	-0.0139	-0.0130	-0.0080

Table 4 Simulation results for Portfolio I, using common set of parameters

(percentages)		CB1	CB2	CB3	CB4	CB5
Expected loss		0.02	0.01	0.01	0.03	0.01
Unexpected loss		0.26	0.25	0.25	0.30	0.27
VaR	99.00	0.19	0.04	0.06	0.37	0.26
	99.90	0.57	0.43	0.51	1.21	1.35
	99.99	17.52	17.03	18.57	21.98	12.97
ES	99.00	0.69	0.55	0.61	1.18	1.08
	99.90	4.39	4.27	4.72	5.68	4.98
	99.99	22.42	21.87	21.74	22.15	21.59
Probability at least 1 default				0.18	1.64	1.47

more relevant here is the conversion of one-year default probabilities into one-month probabilities. Different conversion techniques explain most of the differences in expected losses across participants. CB4 estimated the highest expected loss, consistent with the most conservative assumption for short-term deposits (see Section 3.2.1 and Table 1).

An even stronger impact of this parameter is on the probability of at least one default, where the range of outcomes is much wider between, on the one hand, CB3 and, on the other hand, CB4 and CB5. Computing the probability of at least one default analytically is complicated if correlation is taken into account, but a crude first approximation can be found when the simplifying assumption is made that defaults are independent. The portfolio consists of six obligors rated AAA, 22 with a AA rating and eight which have a rating equal to A. The probability of at least one default equals one minus the probability of no defaults. If, as a starting point, the assumption is made that the maturity of all assets exceeds the holding period of one year, then it is easy to see that the probability of at least one default should be equal to $1 - (1 - 0.01\%)^6 \times (1 - 0.04\%)^{22} \times (1 - 0.10\%)^8 = 1.73\%$, i.e. reasonably close to the results of CB4 and CB5. However, all 30 AA and A obligors represent one-month deposits, and so do two of the six AAA obligors. If the assumed PD over a one-month period is only 1/12th of the annual probability, then the probability of at least one default is reduced to

$1 - (1 - 0.01\%)^4 \times (1 - 0.01\% / 12)^2 \times (1 - 0.04\% / 12)^{22} \times (1 - 0.10\% / 12)^8 = 0.18\%$ only, equal to the result reported by CB3.

The calculations in the previous paragraph are based on assumed default independence. The impact of correlation is rather complex and crucially depends on whether the correlation model deals with asset correlation (as is typically the case) or default correlation. Since the computations above are concerned with default only, it is useful to discuss the impact of default correlation. Consider a very simple although rather extreme example of a portfolio composed of two issuers, *A* and *B*, each with a PD equal to 50%.¹⁴ If the two issuers default independently, then the probability of at least one default equals $1 - (1 - 50\%)^2 = 75\%$. If, however, defaults are perfectly correlated, then the portfolio behaves as a single bond and the probability of at least one default is simply equal to 50%. On the other hand, if there is perfect negative correlation of defaults, then if one issuer defaults, the other does not, and vice versa. Either *A* or *B* defaults and the probability of at least one default equals 100%. Table 5 summarises these results, which show that the probability of at least one default decreases as the default correlation increases. Note that these findings correspond to a well-known result in structured finance, whereby the holder

¹⁴ This rather extreme PD is chosen for illustration purposes only, because perfect negative correlation is only possible with a PD equal to 50%. The conclusions are still valid with other PDs, but the example would be more complex. See also Lucas (2004).

of the equity tranche of an asset pool, who suffers from the first default(s), is said to be “long correlation”. Given the complexity of the computations with multiple issuers, it suffices to conclude that one should expect simulated probabilities of at least one default to be somewhat lower than the analytical equivalents based on zero correlation, but that, more importantly, the assumptions for short duration assets can have a dramatic impact on this probability.

It is instructive to analyse what proportion of expected losses (or any other simulation result) is due to default and how much is due to migration (downgrades). This is not standard output from any of the models, but two participants ran their simulations in default (as well as migration) mode. The results reported by one of them were obtained with the same system (CreditManager®) and the same parameters, except that migration probabilities other than migration to default were set to 0 and the probabilities of ratings remaining unchanged were increased accordingly. Hence, these results can be used to isolate the contribution of default to the total loss. On the other hand, the default and migration mode results reported by the other participant were computed with two different models and can therefore not be used to decompose simulation results. Instead, these results are presented in Section 4.4 on sensitivity analyses.

The result of the decomposition is displayed in Table 6 below. Note that expected losses due to defaults are three times larger than expected losses due to migration, even for a high-quality portfolio such as Portfolio I. This result confirms that in this case the analytical validation of expected loss based on defaults only is sufficiently accurate as a first approximation. Note also that at lower confidence levels, migration is an important source of risk, but that default becomes more relevant as the confidence level is increased. At 99.99%, virtually all the risk comes from default.

Table 5 Probability of at least one default for a hypothetical portfolio (two issuers, each with PD = 50%)

Default correlation	Probability of at least one default (percentages)
-1	100
0	75
1	50

From Table 4, a number of further interesting observations can be made. One of the first things that can be seen is that the VaR and, to a lesser extent, ES are well contained until the 99.90% level, but that these risk measures increase dramatically when the confidence level is raised to 99.99% (which corresponds to the assumed probability of survival (non-default) of AAA-rated instruments, i.e. the majority of the portfolio). Evaluated at the 99.90% confidence level, the CreditVaR is almost irrelevant when compared with the VaR for market risks (in particular currency and gold price risks). However, once the confidence level is raised to 99.99%, credit risk becomes a significant source of risk too. With 0.01% probability, potential losses as measured by the VaR are estimated in the region of 20%. As confirmed by the results in Table 6, defaults have a significant impact on portfolio returns at this confidence level.

Table 6 Decomposition of simulation results into default and migration

(percentages)		
	Default	Migration
Expected loss	75.6	24.4
Unexpected loss	99.7	0.3
VaR 99.00	- ¹⁾	100.0
99.90	45.6	54.4
99.99	99.8	0.2
ES 99.00	77.0	23.0
99.90	98.6	1.4
99.99	99.5	0.5


1) At 99%, there are no defaults. Recall that VaR has been defined as the tail loss exceeding expected losses. As a consequence, the model in default mode reports a negative VaR (i.e. a gain offsetting expected loss) at 99%. For illustration, this result is shown in the table as a 0% contribution from default (and, consequently, 100% from migration).

VaR and ES estimates reported by individual task force members are of the same order of magnitude. To some extent, this may not be surprising, as the participants use similar or even identical systems. Note that the methodology for scaling default probabilities, which has a relatively large impact on expected losses, barely affects tail measures such as VaR and ES, because the portfolio weight of short maturity deposits is relatively small, and the tails of the return distribution are largely determined by defaults of large issuers. Note also that the similarity of simulation results rises with the confidence level. For instance, the ratio of the highest to the lowest ES at the 99.99% confidence level is only 1.04, whereas the same statistic is 2.16 at the 99.00% confidence level. A similar observation can be made for the VaR. While this may seem counterintuitive at first sight, in reality it is not, because the maximum loss is bounded by the default of all issuers in the portfolio. The portfolio is concentrated in a limited number of issuers, and the three largest issuers comprise nearly 80% of the portfolio. As the confidence level is increased, defaults (and downgrades) accumulate, and this result is found by every system. In the limit (confidence level approaches 100%), all issuers have defaulted. At lower confidence levels, the (random) inclusion or exclusion of defaults (or downgrades, which contribute most to the overall VaR and ES at these levels) has a large impact on the simulated credit risk measures.

A plot of the cumulative return distribution sheds more light on these results. Chart 8 shows the return distributions derived from the simulations with the largest disparity of the 99.99% VaR, reported by CB4 and CB5. The return distributions are very similar, and yet, because defaults and rating migrations are discrete events, the two lines happen to cross the 0.01% probability (corresponding to the 99.99% confidence level) at very different return levels. At 99.99%, CB4 reports the highest VaR, but at 99.995% for instance, the order of magnitude is reversed. This example illustrates the importance of using the full return distribution.

Chart 8 Simulated tails of return distribution for Portfolio I (losses > 5% of portfolio only)


(x-axis: return; y-axis: cumulative probability; percentages)



In order to determine the statistical significance of (differences in) simulation results, one participant reported confidence bounds for the VaR estimates, based on standard CreditManager® output. The confidence bounds are based on the observation that the number of scenarios with losses exceeding the VaR is a random variable which follows a binomial distribution with mean $n(1 - \alpha)$, where n equals the number of draws in the simulation and α corresponds to the confidence level of the VaR. For example, if the 99.99% VaR is estimated from 100,000 simulations, then the expected number of scenarios with losses exceeding this VaR is $100,000 \times (1 - 0.9999) = 10$. A binomial distribution with mean $n(1 - \alpha)$ has a standard deviation of $\sqrt{n\alpha(1 - \alpha)}$. CreditManager® computes this standard deviation and finds the corresponding simulation results above and below the VaR (using interpolation when the standard deviation is not an integer number). The difference between the upper and lower bound, expressed as a percentage of the VaR and divided by 2, is reported. For a very large sample, it is reasonable to approximate the distribution of the number of losses exceeding the VaR by a normal distribution, and conclude there is a 68% probability that the “true” VaR will fall within one standard deviation around the estimated VaR. Note that the standard deviation of the binomial distribution increases

Chart 9 Dispersion of simulation results for Portfolio I

(x-axis: confidence level; y-axis: losses as a percentage of portfolio market value, logarithmic scale)



deviation of the binomial distribution increases with the number of simulations n , but that this value represents only the index of observations. As the number of simulations increases, individual simulation results are less dispersed. As a result, the standard deviation of the VaR is expected to decrease.

The reported confidence bounds indicate that the estimates of the 99.00% and 99.90% VaR are very accurate. After 100,000 simulations, the reported standard deviation, rounded to the nearest percentage is 0%. However, the uncertainty surrounding the VaR increases dramatically as the confidence level rises to 99.99%: one standard deviation equals 24% of the VaR. Not surprisingly, given the lack of data, simulation results at very high confidence levels should be treated with care, even though many participants find similar results. The dispersion of simulation results is also displayed graphically in Chart 9.

4.3 SIMULATION RESULTS FOR PORTFOLIO II USING THE COMMON SET OF PARAMETERS

Portfolio II was designed in such a way as to reflect a portfolio for which credit risk is more relevant than for Portfolio I. It is therefore to be expected that risks are higher than in the previous section (see also Chart 7). In this section, the portfolio is analysed using again

the common set of parameters. Note that the proportion of assets with a maturity of less than one year is smaller than for Portfolio I. At the same time, the average maturity of these assets is more than one month (which was the assumed maturity of deposits in Portfolio I). The conversion of annual default probabilities into probabilities for shorter horizons is therefore less relevant here.

The simulation results are shown in Table 7 and Chart 10. Many of the observations from the previous section can also be made for Portfolio II. There is consensus over the probability of at least one default among those who reported it. The simulations all show that the probability is around 12%. The consensus is the consequence of the small number of short duration assets. It is not difficult to verify the consensus level analytically. The portfolio contains 62 issuers with an average PD equal to 0.22% (found by linear interpolation for notches). If, for presentational convenience, this average (rather than individual default probabilities for each rating) is used and defaults are assumed to be independent, then the estimated probability of at least one default is approximately equal to $1 - (1 - 0.22\%)^{62} = 12.8\%$. Again, the “true” probability is slightly different, because defaults are correlated and because averages were used.

Table 7 Simulation results for Portfolio II, using common set of parameters

(percentages)		CB1	CB2	CB3	CB4	CB5
Expected loss		0.17	0.21	0.17	0.25	0.33
Unexpected loss		0.60	0.73	0.60	0.61	0.79
VaR	99.00	2.20	2.72	2.18	2.72	3.85
	99.90	8.52	7.95	8.61	8.08	7.64
	99.99	11.24	11.55	11.24	11.36	10.74
ES	99.00	4.02	5.10	3.96	4.48	5.03
	99.90	9.66	9.59	9.68	9.45	8.86
	99.99	12.97	13.99	12.99	13.08	12.08
Probability at least 1 default				11.89	12.13	12.70

Next, consider expected losses. The average reported expected loss is around 20 basis points. CB1, CB2, CB3 and CB4 are all reasonably close to this average; CB5 is an outlier. Still, all systems find similar tail measures, suggesting that the differences may not be due to computational errors, but to differences in definition (see Section 3.3). It appears that the analytical approximation based on defaults only that was used in the previous section is not as good for Portfolio II, because migration risk is more relevant (see below). Simply multiplying the proportion of the portfolio in each rating by the corresponding PD and adding up the results gives an expected loss due to default of 10 basis points only. A more accurate approximation, still using a simplifying assumption (ratings migrate instantaneously) and using a linear approximation of price changes (i.e. using modified duration), shows that the expected loss would be around 22 basis points, which validates most simulation results.

Table 8 provides a decomposition of simulation results into default and migration. The table shows that default represents less than 50% of the overall expected loss. The proportion is much lower than for Portfolio I (almost 75%), mainly because the duration of Portfolio II is substantially higher. In addition, credit spreads between A (average rating of Portfolio II) and BBB are somewhat larger than between AAA (bulk of portfolio I) and AA. Note also that even at 99.99%, the contribution of migration to VaR and ES is non-negligible.

As before, there is more agreement on VaR and ES at the 99.99% level than at lower confidence levels. Moreover, the simulation results appear to be more precise than for Portfolio I. The reported confidence bound (one standard deviation) around the 99.99% VaR was $\pm 1\%$ only, which indicates that the loss distribution of Portfolio II is much smoother.


The most surprising observation is, however, that while VaR and ES are higher than for Portfolio I at the 99.00% and 99.90% confidence levels (as expected), the numbers are actually lower at the 99.99% confidence level. It turns out that the explanation is the same as that given for the steep rise in VaR and ES at the 99.99% confidence level for Portfolio I: concentration. At very high confidence levels, credit risk is not driven by average ratings or credit quality, but by concentration. Even with low probabilities of

Table 8 Decomposition of simulation results into default and migration

(percentages)		Default	Migration
Expected loss		47.6	52.4
Unexpected loss		77.9	22.1
VaR	99.00	83.9	16.1
	99.90	97.0	3.0
	99.99	91.1	8.9
ES	99.00	87.5	12.5
	99.90	92.2	7.8
	99.99	92.3	7.7

Chart 10 Dispersion of simulation results for Portfolio II

(x-axis: confidence level; y-axis: losses as a percentage of portfolio market value, logarithmic scale)




default, at certain confidence levels defaults will happen, and when they do, the impact is more severe if the obligor has a large weight in the portfolio. Since Portfolio I is more concentrated in terms of the number as well as share of individual obligors, its VaR and ES can indeed be higher than the risk of a portfolio with lower average ratings, such as Portfolio II. In other words, a high credit quality portfolio is not necessarily the least risky. Diversification matters, in particular at high confidence levels. This result is also discussed in Mausser and Rosen (2007).

Chart 11 compares the concentration of Portfolios I and II. Lorenz curves plot the cumulative proportion of assets as a function of the cumulative proportion of obligors. An equally weighted portfolio is represented by a straight diagonal line (note that such a portfolio may still be poorly diversified, as the obligors could all be concentrated in one sector, for instance); at the other extreme, a portfolio concentrated in a single obligor is represented by a horizontal line followed by an almost vertical line. The greater the disparity between the curve and the diagonal, the more the portfolio is concentrated. Chart 11 confirms that Portfolio I is indeed more concentrated than Portfolio II (although the latter also displays a fairly high degree of concentration).

Note that the relative size of individual obligors does not affect the probability of at least one default, which is much higher for Portfolio II than for Portfolio I and rises to a level that may concern investors who fear reputational consequences from a default in their portfolio. Statistically, this result is trivial: the larger the number of (independent) issuers in the portfolio, the larger the probability that at least one of them will default. Mathematically, the probability of at least one default in a portfolio of n independent obligors, each with identical default probability pd , equals $1 - (1 - pd)^n$. For

Chart 11 Lorenz curves for Portfolios I and II

(x-axis: cumulative proportion of issuers; y-axis: cumulative proportion of assets; percentages)



small n and pd , this probability can be approximated by $n \times pd$, and so rises almost linearly with the number of independent obligors. Clearly, increasing the number of independent obligors improves the diversification of the portfolio, reducing VaR and ES. It follows that financial risks (as measured by the VaR and ES) and reputational consequences (if these are related to the probability of at least one default) move in opposite directions as the number of obligors rises.

4.4 SENSITIVITY ANALYSIS USING INDIVIDUAL SETS OF PARAMETERS

Simulations for Portfolios I and II were repeated with one or more alternative sets of parameters, chosen by the participants in the simulation

exercise. The following alternative settings were used:

- Government bonds excluded from the computations, as these are perceived to be credit risk-free.
- Alternative recovery rates and/or asset correlations. These were either fixed at different levels, or simulated, using a stochastic model for the recovery rate. In addition, an alternative correlation matrix for Portfolio II was proposed, which was used by some participants as the basis for their alternative scenarios (also for Portfolio I). This correlation matrix was obtained from one of the rating agencies and has, for Portfolio II, on average substantially lower correlations than the

Table 9 Relative changes in VaR and ES from alternative parameter sets for Portfolio I

(percentages)

		CB1			CB2	
		Alternative 1	Alternative 2	Alternative 3	Alternative 1	Alternative 2
Δ VaR	99.00	-	-94.74	-89.47	-0.11	-43.00
	99.90	+8.77	-85.96	-87.72	-3.34	-83.95
	99.99	+7.08	-95.61	-94.86	-61.55	-99.47
Δ ES	99.00	+5.80	-85.51	-86.96	-29.96	-89.38
	99.90	+7.06	-83.83	-85.65	-43.00	-97.71
	99.99	+0.80	-83.01	-84.26	-31.21	-99.59
Parameter changes	PD/ migration		Government bonds assumed credit risk-free	Government bonds assumed credit risk-free		Government bonds assumed credit risk-free, migration matrix from other rating agency
	Short maturity assets					
	Correlation	Correlation matrix from rating agency (on average lower correlations)	25	CreditMetrics™ factor model (on average higher correlations)	CreditMetrics™ factor model (on average higher correlations)	CreditMetrics™ factor model, but fixed for certain issuers (on average higher correlations)
	Recovery rate		50	Based on study by Altman and Kishore (97)		Based on studies by Altman and Kishore (97) for bonds and Asarnow and Edwards (95) for deposits
	Spreads					CreditManager® (source: Reuters)

uniform correlation in the common parameter set.

- Different migration matrices. A special form of these is the use of default rather than migration mode.
- Other sources for yield curves and spreads.

Some participants chose to run several alternative scenarios. Each of these is shown in Table 9 below, which displays the relative changes in the tail measures VaR and ES for Portfolio I as a result of the parameter change. Empty cells indicate that a parameter was left unchanged. As expected, VaR and ES decrease when the simulation is run in default mode. However, they fall even further in many of the other alternative simulations. Whenever more than one alternative

scenario is shown, the largest impact on each risk measure is shown in italics.

The simulations with the largest impact on (i.e. the largest reduction of) VaR and ES all have one thing in common: the assumption that government bonds bear no credit risk and can therefore be excluded from the analysis. Once these have been excluded, other parameter variations are of secondary importance, although increasing the recovery rate – in particular almost doubling it for deposits (CB3) – obviously matters. The impact of changes in correlations is relatively small, unless they are changed dramatically.

It is recalled that the migration matrix used in the common scenarios was based on empirical rating migrations from S&P, but with the PD for

	CB2				CB3		CB4
	Alternative 3	Alternative 4	Alternative 5	Alternative 6	Alternative 1	Alternative 2	Alternative 1
	-43.00	-38.48	-20.35	+142.60	+40.67	+25.50	-49.46
	-83.95	-74.44	-56.98	-52.12	-38.24	-51.07	-25.02
	-99.47	-98.89	-98.37	-98.15	-46.79	-89.76	-33.83
	-89.25	-85.44	-77.70	-71.17	-26.63	-54.99	-19.73
	-97.70	-96.04	-93.51	-94.09	-29.48	-66.74	-9.65
	-99.59	-99.15	-98.75	-97.58	-46.03	-47.05	-2.94
Government bonds assumed credit risk-free, migration matrix from other rating agency	Migration matrix from other rating agency	Migration matrix from other rating agency	Migration matrix from other rating agency	Migration matrix is mix of three rating agencies	Migration matrix from other rating agency	Govt bonds assumed credit risk-free, migration matrix from other rating agency	Default mode
				"Closest three-month matrix generator"	Maturity of deposits extended to three months	Maturity of deposits extended to three months	
	CreditMetrics™ factor model, but fixed for certain issuers (on average higher correlations)			CreditMetrics™ factor model, but fixed for certain issuers (on average higher correlations)			
Based on studies by Altman and Kishore (97) for bonds and Asarnow and Edwards (95) for deposits	Based on studies by Altman and Kishore (97) for bonds and Asarnow and Edwards (95) for deposits	Based on studies by Altman and Kishore (97) for bonds and Asarnow and Edwards (95) for deposits	Based on studies by Altman and Kishore (97) for bonds and Asarnow and Edwards (95) for deposits	Based on studies by Altman and Kishore (97) for bonds and Asarnow and Edwards (95) for deposits	48% for bonds, 71% for deposits	48% for bonds, 71% for deposits	
CreditManager® (source: Reuters)	CreditManager® (source: Reuters)	CreditManager® (source: Reuters)	CreditManager® (source: Reuters)	CreditManager® (source: Reuters)	Own selection	Own selection	

AAA and AA issuers increased manually – one could argue arbitrarily – from 0 to 1 and 4 basis points per annum respectively (and the probability of the rating remaining unchanged reduced by the same amount). The sensitivity analysis in this section clearly demonstrates the impact of this rather subjective choice.

The sensitivity analysis is repeated in Table 10 for Portfolio II. The impact of alternative parameter sets is much smaller than on Portfolio I. Two participants report increases as well as decreases in the VaR and ES, depending on the confidence level applied. Since there are multiple changes in the parameters – stochastic recovery rates and a different migration matrix

among other things – it is not a priori clear which parameter change dominates at which confidence level, or why.

The largest changes (although not at the 99.99% confidence level) are reported by CB1, which has excluded the few (nine) government bonds from the simulation. Equally important, in this case, is that the recovery rate was increased from 40% to 0.50%. Overall though, the relatively small changes in Table 10 reflect the minor share of government bonds in Portfolio II and the limited impact of alternative PD assumptions for these issuers. Indirectly, this confirms the earlier conclusions.

Table 10 Relative changes in VaR and ES from alternative parameter sets for Portfolio II

(percentages)									
		CB1			CB2		CB3	CB4	CB5
		Alternative 1	Alternative 2	Alternative 3	Alternative 1	Alternative 2	Alternative 1	Alternative 1	Alternative 1
Δ VaR	99.00	-13.64	-17.27	-39.55	-16.53	-17.31	+7.63	-24.26	-0.39
	99.90	-20.07	-17.84	-44.84	-6.24	+2.07	-28.77	-2.35	-16.66
	99.99	-17.88	-16.73	-1.69	-20.38	+18.11	+16.63	-12.41	-19.63
Δ ES	99.00	-23.38	-16.92	-30.85	-16.20	-12.09	+1.59	-17.63	-7.12
	99.90	-14.29	-17.18	-21.84	-16.05	+7.61	-0.66	-9.74	-7.78
	99.99	-27.06	-16.35	-7.25	-26.19	+22.42	+11.73	-10.47	-14.70
Parameter changes	Probability of default/migration	Government bonds assumed credit risk-free		Government bonds assumed credit risk-free	Migration matrix is mix of three rating agencies		Migration matrix from other rating agency	Default mode	
	Short maturity assets	“Closest three-month matrix generator”							
	Correlation	Correlation matrix from rating agency (on average lower correlations)	25	CreditMetrics™ factor model	Correlation matrix from rating agency (on average lower correlations)	CreditMetrics™ factor model, but fixed for certain issuers (on average higher correlations)	CreditMetrics™ factor model (average 48%, standard deviation 26%)	Correlation matrix from rating agency (on average lower correlations)	
	Recovery rate	50 Based on study by Altman and Kishore (97)			Based on study by Altman and Kishore (97)				
	Spreads					CreditManager® (source: Reuters)	Own selection		

5 CONCLUSIONS AND LESSONS LEARNED

Credit risk modelling will gain in importance within the central banking community. From surveys of central bank reserves management practices that are published regularly, it is clear that many central banks are expanding into non-traditional assets, often implying more credit risk taking. Still, central banks are likely to remain conservative investors (as they should) and their overall portfolio risks are unlikely to increase much (indeed, measured in terms of standard deviation of returns, the risk may even be reduced as a result of better diversification). Nevertheless, the special characteristics of credit return distributions warrant the acquisition of expertise in credit risk modelling and suggest that systems be put in place to measure credit risk. An increasing number of the NCBs represented in the task force are using portfolio credit risk models. These models are intended to complement existing market risk models, which are by now commonplace in any central bank. Given the importance of credit risk models in commercial banks, expertise within the investment and risk management divisions of central banks may also have positive spin-offs for other areas of the central banks.

The task force has identified several important lessons that can be learned from its work, and in particular from the simulation exercise. Some of these lessons may already be known, as they apply to every user of a credit risk system; others, however, are more specific to central banks. The lessons are summarised one by one below.

Lesson 1: A portfolio credit risk model is recommended for central banks with credit-risky assets.

While credit risk has traditionally been perceived as a minor part of the overall financial risks in most central bank portfolios, the expansion of the investment universe of central banks and increased awareness of concentration risks have gradually changed the risk assessment. To measure credit risks, and to

compare them quantitatively with other types of central bank risk, a portfolio credit risk model is needed. Such models, which have to apply simulation techniques, can be developed in-house or purchased from external vendors. Several members of the task force have positive experience with the CreditMetrics™ methodology and/or software (CreditManager®). It is easy to use and comprehend, and seems, despite some limitations (some of which apply to any methodology, see also the other lessons below) and the fact that it is designed primarily for corporate bond portfolios, relatively well positioned for central bank purposes.

Yet, it should be emphasised that central banks' experience with credit risk models is relatively limited. Models are regularly improved and new techniques are developed, some of which may be better positioned than current models for the specific needs of central banks. Central banks should closely monitor these developments, as well as the proliferation of different types of ratings (default ratings, recovery ratings, bank deposit ratings, support ratings, etc.), key parameters of the models discussed in this report.

Lesson 2: Measured by CreditVaR, a typical central bank portfolio may exhibit more portfolio credit risk than expected, especially at very high confidence levels.

Central bank reserves are predominantly composed of high quality assets, which should ensure security (and liquidity), especially in the event of market disruptions. A large proportion is invested in government bonds, with low PDs. One would therefore expect credit risk in a central bank portfolio to be very low. However, this need not always be the case, in particular if no obligor is considered default-free. As the comparison of Portfolios I and II demonstrates, credit risk at very high confidence levels is not determined primarily by the average rating level or quality of individual issuers, but by concentration. A portfolio with a small number of large issuers may have a very high CreditVaR, even if the issuers are all AAA governments. Of

course this is only true if they are assumed to carry at least some form of credit risk and are held in the portfolio even after several downgrades. At very high confidence levels, the credit risk of such portfolios may be reduced by replacing some of the government bonds by bonds from other issuers, possibly with a lower rating.

The choice of the confidence level is crucial. For a private financial institution the desired credit rating is one of the main determinants of its risk appetite, required economic capital and, hence, the choice of a confidence level for VaR. A central bank's risk appetite depends on its mandate, exchange rate mechanism, balance sheet structure and relationship with the shareholder(s), among other things. Arguably, a central bank – with reputation as its main asset – should aim for high confidence levels, also in comparison with commercial institutions.

Lesson 3: The quality of results crucially depends of the quality of assumptions on parameters. Some of these are of particular relevance to central banks.

The first part of this lesson is trivial. In credit risk modelling, the lack of data is a problem shared by all market participants. Gordy (1999), in his comparative anatomy of credit risk models, concluded: “This sensitivity ought to be of primary concern to practitioners. It is difficult enough to measure expected default probabilities and their volatility. Capital decisions, however, depend on extreme tail percentile values of the loss distribution, which in turn depend on higher moments of the distribution of the systemic risk factors. These higher moments cannot be estimated with any precision given available data. Thus, the models are more likely to provide reliable measures for comparing the relative levels of risk in two portfolios than to establish authoritatively absolute levels of capital required for any given portfolio”.

Nevertheless, the task force identified two (to some extent related) types of parameter choice that are of particular relevance to central banks

and at the same time very hard to estimate. These parameters are the PD for high quality issuers, in particular the difference between government and corporate issuers, and the application of any PD to short duration assets.

The sensitivity analyses of the simulation exercise reveal the impact of the treatment of government (and related) issuers on VaR and ES. Empirically, defaults of AAA or AA-rated issuers over a period of one year are (virtually) non-existent. For government issuers, these are rarer still, even at lower rating levels. Consequently, migration matrices predict a (close to) zero PD for these high-quality issuers, which may be too optimistic. In response, investors (including central banks) have typically made rather ad hoc adjustments to the migration matrices, whereby PDs are adjusted manually to ensure they are positive, in such a way that the ranking of ratings is still respected. The size of the adjustment is, however, not based on any empirical evidence, and may have a dramatic impact on the results, in particular for central bank portfolios dominated by high-quality issuers. Clearly this is not totally satisfactory, and it raises the question whether other “null” migration probabilities must be raised as well.

The typical central bank portfolio is not only highly rated but also has a relatively low duration and a significant proportion of short maturity assets. Normally, the horizon for the credit risk analysis exceeds the maturity of these assets, and if one uses annual probabilities of default, it will be necessary to convert these into short-term equivalents. This can be done in several ways; the choice can have a significant impact on the calculated risk measures. In practice, approximations of the “true” short-term probabilities are needed, and members of the task force generally use an approach that results in conservative PD estimates. It seems, however, more efficient for short horizon PDs to be estimated directly from ratings databases, rather than derived from annual PDs.

Lesson 4: There may sometimes be a trade-off between financial and reputation risk.

A central bank that invests in credit-risky bonds may be especially concerned with default risk, not necessarily for the potential loss of money, but for any reputational consequences should it become known to the outside world that the central bank had invested in a company that had defaulted. It therefore seems logical to compute the probability that such an event might occur, or even to limit this probability. From an investment perspective, however, as demonstrated in this report, the probability of at least one default is not a good risk measure, because it discourages diversification. Statistically, it is obvious that as the number of issuers in the portfolio increases, the probability of a default by at least one of them increases. If the aim were to reduce this probability, it would be optimal to invest only in the issuer with the lowest credit risk (i.e. government bonds from one country only). Clearly, financial risk management would recommend the opposite, and it can be concluded that if reputation risk is related to being involved in a corporate default, then there is a trade-off between financial and reputation risks.

The lesson is that this probability must be treated with care in determining the composition of a credit portfolio. The likelihood as well as the magnitude of a potential loss are relevant for reputation risks. A central bank more concerned about reputational consequences of a single default than about the financial consequences could consider buying credit exposure indirectly, through participation in a fund or the use of a derivative on an index. A default by one of the fund/index members would still have a negative impact on performance (albeit a small one if the fund/index is well diversified) and buying credit exposure indirectly has its disadvantages, such as limited control over the composition of the fund or index. Still, a default by a fund/index member should not have the same reputational consequences as a direct involvement in a corporate bankruptcy. The magnitude of a potential loss would then be more relevant than

the likelihood, and reputation and financial risk management would be more aligned.

Lesson 5: Integrated risk management is important.

Although central banks have idiosyncratic features, integrated risk management should be a goal for them as for any other financial institution. Central banks' balance sheets and, hence, financial risks are largely dictated by their mandate. To some extent, they also reflect past developments. Currency and gold price risks can only be hedged to a limited extent, if at all. At the same time, external stakeholders demand prudent investment styles and risk management, and sometimes even stable financial results. Reducing other risks, such as interest rate risk or credit risk, will often have a small impact only, and may even increase overall risks. A better, but also more complicated, way forward therefore seems a framework of integrated risk management, which takes into account all possible risks and thus allows a proper modelling of diversification (into, for example, additional currencies or new asset classes) and joint optimisation techniques. This work is still in its early stages. However, modelling portfolio credit risk is clearly a first, crucial step in integrating the different types of risk, which starts by measuring them homogeneously.

REFERENCES

- Altman, E. I. and V. M. Kishore (1996), "Almost everything you wanted to know about recoveries on defaulted bonds", *Financial Analysts Journal*, (November/December), pp. 57-64.
- Altman, E. I., A. Resti and A. Sironi (eds.) (2005), *Recovery risk: the next challenge in credit risk management*, Risk Books.
- Artzner, P., F. Delbaen, J.-M. Eber and D. Heath (1999), "Coherent measures of risk", *Mathematical Finance*, 9 (3), pp. 203-228.
- Asarnow, E. and D. Edwards (1995), "Measuring loss on defaulted bank loans: A 24-year study", *Journal of Commercial Lending*, 77 (7), pp. 11-23.
- Bank for International Settlements (2005), "Zero-coupon yield curves: technical documentation", BIS Papers, 25.
- Basel Committee on Banking Supervision (2006), "International convergence of capital measurement and capital standards: a revised framework", Bank for International Settlements.
- Bindseil, U. and F. Papadia (2006), "Credit risk mitigation in central bank operations and its effects on financial markets: the case of the Eurosystem", ECB Occasional Paper, 49.
- Bluhm, C., L. Overbeck and C. Wagner (2003), *An introduction to credit risk modeling*, Chapman & Hall.
- Bucay, N. and D. Rosen (1999), "Credit risk of an international bond portfolio: a case study", *Algo Research Quarterly*, 2 (1), March, pp. 9-29.
- Carty, L. V. and D. Lieberman (1996), "Defaulted bank loan recoveries," Moody's Investors Service, *Global Credit Research*, special report.
- Cossin, D. and H. Pirotte (2007), *Advanced credit risk analysis: financial approaches and mathematical models to assess, price and manage credit risk*, second edition, Wiley.
- Crouhy, M., D. Galai and R. Mark (2000), "A comparative analysis of current credit risk models", *Journal of Banking & Finance*, 24 (1-2), pp. 59-117.
- de Beaufort, R., S. Benitez and F. Palomino (2002), "The case for reserve managers to invest in corporate debt", *Central Banking*, XII, 4, pp. 79-87.
- Duffie, D. and K. J. Singleton (2003), *Credit risk: pricing, measurement and management*, Princeton University Press.
- Dynkin, L., J. Hyman and V. Konstantinovskiy (2002), "Sufficient diversification in credit portfolios", Lehman Brothers Fixed-Income Research.
- Dynkin, L., J. I. Mann, S. Mody and B. D. Phelps (2004), "Portfolio and index strategies during stressful credit markets", Lehman Brothers Fixed-Income Research.

- Glasserman, P. (2005), *Monte Carlo methods in financial engineering*, Springer.
- Gordy, M. B. (2000), “A comparative anatomy of credit risk models”, *Journal of Banking & Finance*, 24 (1-2), pp. 119-149.
- Grava, R. L. (2004), “Corporate bonds in central bank reserves portfolios: a strategic asset allocation perspective”, in: C. Bernadell, P. Cardon, J. Coche, F. X. Diebold and S. Manganelli (eds.), *Risk management for central bank foreign reserves*, European Central Bank.
- Gupton, G. M., C. C. Finger and M. Bhatia (1997), “CreditMetrics™ – Technical Document”, JP Morgan.
- Hull, J. and A. White (2000), “Valuing credit default swaps I: no counterparty default risk”, *Journal of Derivatives*, 8 (1), pp. 29-40.
- Kern, M. and B. Rudolph (2001), “Comparative analysis of alternative credit risk models: an application on German middle market loan portfolios”, Center for Financial Studies Working Paper 2001/03.
- Koyluoglu, H. U. and A. Hickman (1998), “Reconcilable differences”, *Risk*, 11 (10): pp. 56-62.
- Kreinin, A. and M. Sidelnikova (2001), “Regularization algorithms for transition matrices”, *Algo Research Quarterly*, 4 (1/2), pp. 23-40.
- Lando, D. and T. M. Skødeberg (2002), “Analysing rating transitions and rating drift with continuous observations”, *Journal of Banking & Finance*, 26, pp. 481-523.
- Lando, D. (2004), *Credit risk modeling: theory and application*, Princeton University Press.
- Loeys, J. (1999), “How much credit?”, JP Morgan.
- Lucas, A., P. Klaassen, P. Spreij and S. Straetmans (2001), “An analytic approach to credit risk of large corporate bond and loan portfolios”, *Journal of Banking & Finance*, 25, pp. 1635-1664.
- Lucas, D. (2004), “Default correlation: from definition to proposed solutions”, UBS CDO Research.
- Mausser H. and D. Rosen (2007), “Economic credit capital allocation and risk contributions”, in: J. Birge and V. Linetsky (eds.), *Handbook of financial engineering* (forthcoming).
- Merton, R. (1974), “On the pricing of corporate debt: the risk structure of interest rates”, *Journal of Finance*, 29 (2), pp. 449-470.
- Moody’s (2004), “Recent bank loan research: implications for Moody’s bank loan rating practices”, Moody’s Investors Service Global Credit Research.
- Nelson, C. R. and A. F. Siegel (1987), “Parsimonious modeling of yield curves”, *Journal of Business*, 60 (4), pp. 473-489.

- Nickell, P., W. Perraudin and S. Varotto (2000), “Stability of rating transitions”, *Journal of Banking & Finance*, 24, pp. 203-227.
- O’Kane, D. and L. Schlögl (2001), “Modelling credit: theory and practice”, Lehman Brothers Analytical Research Series.
- Pluto, K. and D. Tasche (2006), “Estimating probabilities of default for low default portfolios”, in: B. Engelmann and R. Rauhmeier (eds.), *The Basel II risk parameters*, Springer, pp. 79-103.
- Pringle, R. and N. Carver (2005), “Trends in reserve management – survey results”, in: *RBS Reserve Management Trends 2005*, Central Banking Publications.
- Ramaswamy, S. (2004), *Managing credit risk in corporate bond portfolios: a practitioner’s guide*, Wiley.
- Saunders, A. and L. Allen (2002), *Credit risk measurement: new approaches to value at risk and other paradigms*, second edition, Wiley.
- Stamcar, R. (2007), “Portfolio credit spread risk”, *RiskMetrics Journal*, 7 (1), pp. 3-16.
- Standard & Poor’s (2007a), *Annual 2006 global corporate default study and rating transitions*.
- Standard & Poor’s (2007b), *Sovereign defaults and rating transition data: 2006 update*.
- Vasicek, O. A. (1991), “Limiting loan loss probability distribution”, KMV Corporation.
- Wooldridge, P. D. (2006), “The changing composition of official reserves”, *BIS Quarterly Review*, September, pp. 25-38.

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. Wredenberg, 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. Przybyła, 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, May 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.

ISSN 1607148-4



9 771607 148006