Basel Committee
on Banking Supervision

Working Paper No. 15

Studies on credit risk concentration

An overview of the issues and a synopsis of the results from the Research Task Force project

November 2006

BANK FOR INTERNATIONAL SETTLEMENTS

The Working Papers of the Basel Committee on Banking Supervision contain analysis carried out by experts of the Basel Committee or its working groups. They may also reflect work carried out by one or more member institutions or by its Secretariat. The subjects of the Working Papers are of topical interest to supervisors and are technical in character. The views expressed in the Working Papers are those of their authors and do not represent the official views of the Basel Committee, its member institutions or the BIS.

Requests for copies of publications, or for additions/changes to the mailing list, should be sent to:

Bank for International Settlements
Press & Communications
CH-4002 Basel, Switzerland

E-mail: publications@bis.org
Fax: +41 61 280 9100 and +41 61 280 8100

ISSN: 1561-8854

Contents

Research Task Force Concentration Risk Group of the Basel Committee on Banking Supervision

Chairman: Mr Klaus Duellmann, Deutsche Bundesbank, Frankfurt

Mr Per Asberg Sommar	Sveriges Riksbank, Stockholm
Mr Julien Demuynck	French Banking Commission, Paris
Ms Antonella Foglia	Bank of Italy, Rome
Mr Michael B Gordy	Board of Governors of the Federal Reserve System, Washington
Mr Takashi Isogai	Bank of Japan, Tokyo
Mr Christopher Lotz	Federal Financial Supervisory Authority (BaFin), Bonn
Ms Eva Lütkebohmert	Deutsche Bundesbank, Frankfurt
Mr Clément Martin	French Banking Commission, Paris
Ms Nancy Masschelein	National Bank of Belgium, Brussels
Ms Catherine Pearce	Office of the Superintendent of Financial Institutions, Ottawa
Mr Jesús Saurina	Bank of Spain, Madrid
Mr Martin Scheicher	European Central Bank, Frankfurt
Mr Christian Schmieder	Deutsche Bundesbank, Frankfurt
Mr Yasushi Shiina	Financial Services Agency, Tokyo
Mr Kostas Tsatsaronis	Bank for International Settlements, Basel
Ms Helen Walker	Financial Services Authority, London
Mr Martin Birn	Secretariat of the Basel Committee on Banking Supervision, Bank for International Settlements, Basel

Executive summary

Concentration of exposures in credit portfolios is an important aspect of credit risk. It may arise from two types of imperfect diversification. The first type, name concentration, relates to imperfect diversification of idiosyncratic risk in the portfolio either because of its small size or because of large exposures to specific individual obligors. The second type, sector concentration, relates to imperfect diversification across systematic components of risk, namely sectoral factors. The existence of concentration risk violates one or both of two key assumptions of the Asymptotic Single-Risk Factor (ASRF) model that underpins the capital calculations of the internal ratings-based (IRB) approaches of the Basel II Framework. Name concentration implies less than perfect granularity of the portfolio, while sectoral concentration implies that risk may be driven by more than one systematic component (factor).

The Concentration Risk Group of the Research Task Force of the Basel Committee on Banking Supervision undertook a principally analytical project with the following objectives: (i) to provide an overview of the issues and current practice in a sample of the more advanced banks as well as highlight the main policy issues that arise in this context; (ii) to assess the extent to which "real world" deviations from the "stylised world" behind the ASRF assumptions can result in important deviations of economic capital from Pillar 1 capital charges in the IRB approach of the Basel II Framework; and (iii) to examine and further develop fit-for-purpose tools that can be used in the quantification of concentration risk.

The work of the group was divided into three workstreams. The first workstream collected information about the current "state of the art" both in terms of industry best practice and in terms of the developments in the academic literature. A workshop organised in November 2005 was an occasion to exchange views among experts from the supervisory, academic and industry areas. These contacts revealed that there is a great deal of diversity in the way banks measure and treat concentration risk. Some employ sophisticated portfolio credit risk models that incorporate interactions between different types of exposures while some rely on simpler, ad hoc indicators of such risk. Multi-factor vendor models are also used as inputs or benchmarks to internal models. Management of concentration risk typically depends on a variety of tools including limits on single entity exposures either in terms of overall credit limits or economic capital, and pricing tools that are used by a minority of banks. Typical stress tests employed by banks include a concentration risk component although this is not always studied separately. The availability of the necessary bank-level data for the analysis of concentration risk remains an important practical issue especially when it comes to producing stable and reliable estimates of asset correlation across exposures.

The second workstream focused on gauging the impact of departures from the ASRF model assumptions on economic capital and examined various methodologies that can help to bridge the gap between underlying risk and risk measured by the specific model. The workstream had two sub-themes that focused on name concentration risk (imperfect portfolio granularity) and sector concentration risk (imperfect diversification across risk factors).

The empirical studies conducted by the group, all of which used data only on corporate portfolios, suggest that name concentration risk, albeit important in its own sake, is likely to represent a smaller marginal contribution to economic capital than sector concentration for a typical commercial bank with a medium to large sized loan portfolio. For these portfolios, name concentration could add anywhere between 2 and 8% to the credit value-at-risk while sector concentration can increase economic capital by 20-40%. The patterns of asset correlations both across and within sectors are key determinants of this impact. While single-factor credit risk frameworks tend to produce higher measures of risk in certain circumstances because they generally do not account for diversification across credit portfolio types (eg between wholesale and retail) or do not fully allow for diversification gains within portfolio types, there are also situations in which single-factor credit risk models produce lower measures of risk because they do not capture name and sectoral concentrations.

The notion of name concentration risk is generally better understood than sectoral concentration risk and a number of analytical measurement tools have been proposed in the literature. Some are based on ad hoc measures of concentration (such as the Herfindahl-Hirschman index of portfolio exposures) while others are more firmly embedded in formal models of credit risk. The latter are preferred to the former whenever the needed data requirements are met because they represent a more consistent approach to the measurement and management of all dimensions of credit risk for the portfolio. The group elaborated on an adjustment for imperfect portfolio granularity which had been proposed as part of an earlier version of Basel II. The revised method incorporates analytical advancements that have occurred in the meantime and deals with some practical complications of the earlier proposal.

Sector concentration arises from the violation of the single systematic risk factor assumption which represents an elementary departure from the IRB model framework. It arises because business conditions and hence default risk may not be fully synchronised across all business sectors or geographical regions within a large economy. A bank's portfolio may be more or less concentrated on some of these risk factors leading to a discrepancy between the measured risk from a single-factor model and a model that allows for a richer factor structure. Given the calibration of the ASRF model for the IRB formulae, this discrepancy can be positive as well as negative.

The group examined various methods that can deal with sector concentration. Some represent tools that can be considered as extensions of more elementary models while others start from a more general multi-factor structure. An example of the former group of tools is a multiplicative adjustment to the ASRF model which uses a more general calibration to a multi-factor model to incorporate concentration risk and was found to perform quite well. In terms of tools that rely explicitly on multi-factor frameworks the group studied the performance of a simplified version of a model originally proposed by Pykhtin and obtained very favourable results. Overall, the choice of approach depends very much on the purpose of the exercise and the availability of the necessary inputs (such as estimates of differentiated probability of default, loss-given-default and asset correlations for various sectors). All approaches require considerable care and judgment by the analyst.

The third workstream focused mostly on the ability of stress tests to detect excessive concentration (of either type) and to provide estimates of economic capital in stress scenarios. Plausibility, consistency with the credit portfolio model, being adapted to the portfolio under consideration and being reportable to senior management were identified as desirable properties for stress tests. A methodology based on the idea of stressing core factors while other factors move conditional on them demonstrates that it is possible to derive stress tests on the basis of a consistent model and a close link between the model and the real world.

Finally the group highlighted a number of technical issues that while outside the scope of the project, are nonetheless important in dealing with the overall issue of concentration risk in credit portfolios. These were: (i) the choice of an adequate sector scheme for the purpose of concentration risk assessment; (ii) the definition of a "benchmark" for concentration risk correction; and (iii) data-related issues.

Studies on credit risk concentration

Historical experience shows that concentration of credit risk in asset portfolios has been one of the major causes of bank distress. This is true both for individual institutions as well as banking systems at large. The failures of large borrowers like Enron, Worldcom and Parmalat were the source of sizeable losses in a number of banks. Large exposures to less-developed countries' debt were one of the reasons of protracted weakness of major US banks in the 1980s demonstrating that the stability of entire systems can be undermined by the excessive exposure to a single asset class. More intriguing, banks in Texas and Oklahoma suffered severe losses in both corporate and commercial real estate lending in the 1980s. The reason being that in addition to very significant concentrations of lending in the energy industry, the regional dependence on oil implied a strong correlation between the health of the energy industry and local demand for commercial real estate.

These examples illustrate the importance of measuring concentration risk in credit portfolios of banks that arises not only from exposures to a single credit, or asset class, but also from linkages between asset classes. The Asymptotic Single-Risk Factor (ASRF) model[1] that underpins the IRB approach in the new Basel capital framework[2] does not allow for the explicit measurement of concentration risk. A group of researchers from the Research Task Force (RTF) of the Basel Committee on Banking Supervision undertook a project with the goal of analysing the ability of various methods to account for concentration risk in bank loan portfolios and to survey current best-practice in the industry.

This paper provides an overview of the work conducted by the Concentration Risk Group of the RTF ("the group") and its findings. The complete results of the project are to be found in individual research papers and reports listed at the end of this working paper. The various methodologies for the treatment of concentration risk which were analysed or refined by the group aim to reflect the current state of research in the industry and in academia. Importantly, the group does not give recommendations for the use of any specific approach. Instead, the purpose of this paper is to put the various methodologies into perspective. It is stressed that the individual studies as well as this paper largely reflect the views of individual authors, and should not be viewed as representing specific Basel Committee guidance for supervisory authorities or financial institutions.

The structure of the paper is as follows. The next section discusses the main issues related to concentration risk and the limitations of the single-factor model in this respect, which motivate this project. The second section presents the objectives of the group and the overall structure of the project. The third section presents the results of an informal survey of industry best practice conducted by the group. The fourth section presents the main results of the research conducted by group members and is divided in two sub-sections: one deals with the question of single name concentration, and the other with the question of sector concentration. It also includes a brief discussion of the concepts related to contagion risk on which the group has not conducted new research but produced some empirical evidence. The fifth section discusses the modalities of stress testing loan portfolios for concentration risk. The final section lists a number of practical issues related to the measurement of concentration risk which were identified by the group.

---

[1] See Gordy (2003).

[2] BCBS (2006).

1. The assumptions in the IRB model

In the risk-factor frameworks that underpin both industry models of credit value-at-risk (VaR) and the internal ratings-based (IRB) risk weights of Basel II, credit risk in a portfolio arises from two sources, systematic and idiosyncratic:

- *Systematic risk* represents the effect of unexpected changes in macroeconomic and financial market conditions on the performance of borrowers. Borrowers may differ in their degree of sensitivity to systematic risk, but few firms are completely indifferent to the wider economic conditions in which they operate. Therefore, the systematic component of portfolio risk is unavoidable and only partly diversifiable.
- *Idiosyncratic risk* represents the effects of risks that are particular to individual borrowers. As a portfolio becomes more fine-grained, in the sense that the largest individual exposures account for a smaller share of total portfolio exposure, idiosyncratic risk is diversified away at the portfolio level. This risk is totally eliminated in an infinitely granular portfolio (one with a very large number of exposures).

The IRB risk-weight functions of Basel II were developed with the idea that they would be portfolio invariant, ie the capital required for any given loan should only depend on the risk of that loan and must not depend on the portfolio it is added to. This characteristic has been deemed vital in order to make the new IRB framework applicable to a wider range of countries and institutions. In the context of regulatory capital allocation, portfolio invariant allocation schemes are also called ratings-based. This notion stems from the fact that, by portfolio invariance, obligor-specific attributes like probability of default (PD), loss-given-default (LGD) and exposure-at-default (EAD) suffice to determine the capital charges of credit instruments.[3]

In order to achieve portfolio invariance, at least asymptotically, the ASRF model framework that underpins the IRB approach is based on two key assumptions:[4] (a) bank portfolios are perfectly fine-grained, and (b) there is only one source of systematic risk. The first assumption implies that there are no exposure “lumps” in the portfolio. In other words, no single exposure accounts for more than a vanishingly small share of the total portfolio. Idiosyncratic risk is diversified away. The second assumption implies that the commonality of risk between any two individual credits is uniquely determined by the intensity of their respective sensitivities to the single systematic factor. In other words, there are no diversification possibilities beyond the reduction in idiosyncratic risk which comes from increasing the granularity of the portfolio. Strictly speaking, this second assumption pertains to the sources of credit risk for the economy as a whole rather than for the individual bank portfolio, and requires that there be no sectoral or geographic sources of risk that are distinct from the macroeconomy. A somewhat looser interpretation is that bank portfolios are well-diversified across sectors and geographical regions, so that the only remaining systematic risk is to the performance of the economy. It is in this looser sense that the assumption can be seen as a requirement on bank portfolios.

When these two assumptions hold, it is possible to show that the risk assessment of the credit portfolio can be conducted from the bottom up. Since idiosyncratic risk is assumed to be fully diversified one only needs to assess the systematic component of risk. For this latter component an assessment can be made at the level of the individual exposure and the

[3] See BCBS (2004).

[4] See Gordy (2003).

results simply added up to provide the assessment for the entire portfolio. This is the basis for the IRB approach, which relies on such individual credit assessments and does not allow for a rich correlation structure between individual risks. If the two assumptions hold then those correlations simply do not contain any additional information.

When the two assumptions are violated, however, there is no guarantee that the bottom-up approach will be accurate. The marginal contribution to overall risk by any single exposure will likely depend on the risk profile of the rest of the portfolio. In particular, adding up the IRB-based capital requirements relating to individual exposures might over- or under-state the risk of the portfolio depending on whether the portfolio is diversified or concentrated relative to the one used as a calibration benchmark.

There are important reasons why the Committee opted for the particular additive bottom-up framework. These include the relative simplicity of the bottom-up approach, the fact that the stage of development of more realistic portfolio credit models at the time was judged inadequate for regulatory purposes, and the fact that the validation of inputs is easier than the validation of full models. The desire for portfolio invariance, however, makes recognition of institution-specific diversification effects within the framework difficult: diversification effects depend on how well a new loan fits into an existing portfolio. To maintain internal consistency, the ASRF modelling restrictions were embedded in the methodologies used to calibrate the IRB risk weights. In particular, it assumed a fully granular portfolio in terms of single name exposures, and the asset correlation parameters were chosen to match the economic risk in a credit portfolio that is very well-diversified across sectors (see further discussion on this point below).

As mentioned earlier, the specific assumptions behind the ASRF model are unlikely to be exactly met by actual portfolios, especially those of institutions that are smaller in size or relatively specialised. Concentration risk can arise from significant single exposures, from concentration in specific business sectors, and from potential loss dependencies because of direct business links between borrowers or indirectly through credit risk mitigation.

2. The concentration risk project of the RTF

The potential importance of concentration risk in actual bank portfolios highlights the need for supervisors to assess the potential gap between Pillar 1 capital requirements and the “true” underlying risk. The notion and implications of single name concentration risk are reasonably well-understood, despite a few open issues regarding implementation. The measurement of sector concentration, however, which is relatively more important, is technically quite challenging, especially given the lack of guidance from the literature.

The group examined issues related to both types of concentration risk. More specifically, it conducted analytical work on assessing the importance of single name and sector concentration risk and researched possible approaches to deal with these types of risk. This section presents a brief overview of the objectives of the project and the different workstreams. The group’s more specific findings are outlined in sections 3, 4 and 5 and are fully detailed in five technical papers of working group members, listed in the References. The project undertaken by the group had three main objectives. The first was to provide an overview of the issues and current practice in a sample of the more advanced banks as well as highlight the main policy issues that arise in this context. The second objective was to assess the extent to which “real world” deviations from the “stylised world” behind the ASRF assumptions can result in important deviations of economic capital from Pillar 1 capital charges in the IRB approach of the Basel II Framework.[5] The last objective is to examine and further develop fit-for-purpose tools that can be used in the quantification of concentration risk.

The project is divided into three broad workstreams, each with a separate but complementary function and addressing to a different degree one or more of the listed objectives:

1. An informal survey of the state-of-the-art methods that account for concentration risk used by a sample of “best practice” institutions. The objective was to identify advances in technology and improvements in data availability, as well as to outline some policy lessons.

2. An analysis of the impact of departures from the assumptions of the ASRF model on economic capital, and various methodologies that can help to bridge the gap between underlying risk and risk measured by the specific model. This workstream had two sub-themes:

   (a) To gauge the importance of single-name concentration (not fully diversified idiosyncratic risk) and to develop an adjustment to the ASRF model for this type of risk.

   (b) To assess the impact of sector (and country) concentration (ie the existence of multiple systematic risk factors) on overall portfolio risk. The papers in this sub-theme focus on gauging the deviations of “true” capital from the single-factor assumption of the ASRF model. They also researched risk measurement methodologies that could minimise these deviations.

3. The third workstream focused mostly on the ability of stress tests to detect excessive concentration (of either type) and to provide estimates of economic capital in stress scenarios.

Given resource constraints and areas of comparative expertise, the group decided not to address certain issues in this project. It focused on questions of concentration risk in credit portfolios (ie the asset side of the balance sheet) and did not address issues related to the management of this risk arising from liabilities or transactions. Moreover, it focused its efforts more on questions related to sector concentration risk, judging this to be an area where, despite its materiality for banking institutions, progress in research has been relatively limited. At the same time, no analysis was conducted on sector concentration risk that arises indirectly via credit risk mitigation. Neither did the group carry out empirical analyses on regional concentration.

The work was mainly research-oriented and comprised the enhancement of methodologies and empirical tests. Without compromising scientific rigour, the group focused primarily on fit-for-purpose solutions that take into account typical data limitations. In addition, a research workshop with external presenters was hosted by the Deutsche Bundesbank in November

---

[5] In this paper “economic capital” always refers to the difference between the value-at-risk of a credit portfolio on a 99.9% confidence level and the expected loss, given a certain model. It corresponds to the term “unexpected loss” which is used in the Basel II Framework as the conceptual basis of the IRB risk-weight functions for credit risk.

2005 to initiate discussion with practitioners and to spur further academic research in this area.[6]

3. Survey of best practice

To gain a better understanding of how concentration risk is treated at major banks, the group undertook an informal survey of a small number of best practice institutions.[7] Further feedback about industry practice was gathered at the workshop with bankers, supervisors and academics. This section provides a brief summary of the information gathered through these channels.

A general impression needs to be highlighted first. Banks and supervisors often do not have the same understanding about concentration risk, and in particular about its relation to the Basel II Framework. Supervisors interpret concentration risk as a positive or negative deviation from Pillar 1 minimum capital requirements derived by a framework that does not account explicitly for concentration risk. Banks perceive that sector concentration (often referred to, with a positive connotation, as "diversification") warrants capital relief relative to Pillar 1, which they take as the non-diversified benchmark. This difference in perspective is discussed in more detail below.

Overall, business-sector concentration has traditionally received less attention by banks as a source of concentration risk than exposure concentration in geographic regions.

In general, banks have different measures in place to capture and manage concentration risk: (i) exposure limit systems, which also depend on the strategic goals of the bank; (ii) internal economic capital models that measure the risk contribution of exposures for risk management purposes; and (iii) "pricing tools" that allow banks to account for concentration risk in the pricing of a new exposure. Whereas limit systems and internal models are commonly applied across best practice banks, incorporating concentration risk in the pricing of new loans is practiced by less than half of the banks.

There is also a disparity across the best practice banks in the methodological treatment of concentration risk. The more sophisticated banks employ internal economic capital models that can in principle adequately measure concentration risk but they are often constrained by data problems, for example, by grouping exposures to risk entities. The less sophisticated institutions surveyed employ simpler concentration measures, such as the Herfindahl-Hirschman index, which do not allow the translation of concentration risk into an economic capital figure (see discussion on this topic below).

Banks which capture concentration risk by internal multi-factor models do not necessarily recognise concentration risk explicitly as a separate risk category. Credit risk from large exposures to individual industry sectors is often perceived as a risk that arises from asset correlations between exposures rather than from exposure concentrations. Therefore, it is often not captured by the limit system and instead accounted for indirectly through the (marginal) risk contribution of an exposure, given by the internal model.

---

6 Research papers from this workshop were published in a special issue of the Journal of Credit Risk in Fall 2006.

7 The surveyed banks were from: Belgium, Canada, Germany, Italy, Japan, Spain, Sweden, and the U.K.

Limit systems often do not capture concentration risk that arises from distinct but correlated exposures. Moreover, they are usually applied in the context of exposures to single obligors or to specific geographical regions rather than to exposures to business sectors. Finally, limits are often decided on the basis of a variety of business considerations and strategic objectives of which controlling concentration risk is only one aspect.

Banks use a mix of vendor models and in-house built models to capture concentration risk in their economic capital calculations. Vendor models are also used as a benchmark for internal models. Typically these are multi-factor asset value models and sensitivity to industry and/or geographical factors is measured through asset correlations. These correlations are in turn typically estimated on the basis of either equity correlations, or correlation estimates derived from rating migrations or default events. The number of employed factors can vary from as few as seven to as many as 110. Stability of the estimated correlations is an issue that banks often have to cope with.

Credit risk mitigation techniques are taken into account if economic capital models are used. They are also accounted for, although generally to a lesser extent, when concentration risk is controlled by a limit system.

Concentration risk is generally managed on a centralised basis through the monitoring of exposures. However, at some banks business units are given discretion to impose their own controls over concentration risk. Practice regarding incentives in the management of concentration risk varied across institutions, albeit many mentioned that performance measurement is already, or will soon be linked to the return on economic capital.

Banks reported using different methods of stress testing for concentration risk. Test scenarios include the downgrade of all large exposures or of a large sector, the increase of exposures to a cluster of borrowers, or the increase of the PD and/or the LGD for a group of exposures. However, it is often difficult to distinguish stress tests that are specific to addressing concentration risk from more general stress tests of credit risk. For the most part, concentration risk stress tests are conducted on an ad hoc rather than a regular basis.

Measuring concentration risk relative to Pillar 1 capital charges will remain a challenge even for the most sophisticated, best-practice banks. The availability of data is always an important issue. In emerging markets, risk estimation is more difficult and possibly less reliable since markets are often less liquid. Apart from data constraints, the growing complexity of banks' business, in particular the increasing use of credit risk transfer instruments, limits the accuracy of simple tools.

The measurement methodology for concentration risk also needs to be commensurate with the complexity of the banks' business and the environment in which they operate. These issues highlight the importance of gaining a firm understanding of the structure and characteristics of the risk measurement model.

4. Economic capital issues

The bulk of the group's work focused on the measurement and the modelling of concentration risk arising either from imperfect granularity (large single name exposures) or imperfect sectoral diversification. The following two sub-sections present an overview of the main results of the group's efforts in this respect.

Prior to the discussion of the specific approaches, it is useful to briefly describe the Herfindahl-Hirschman Index (HHI) which is extensively used in the context of different methodologies presented below. The HHI is a popular measure of concentration that has found many and varied applications. It is used extensively in the empirical industrial organisation literature and as a diagnostic tool by competition authorities in some jurisdictions. It is calculated as the sum of squared market shares (measured in fractions) of each market participant, and often expressed in a scale of 0 to 1. It is a continuous measure with zero corresponding to the fully granular case (each participant has an infinitesimal share) and unity corresponding to monopoly (there is only one participant). In the context of the measurement of (single name or sector) concentration risk the HHI formula is included as a component of a number of approaches. Its specific use will be discussed in the appropriate context below.

4.1 Imperfect granularity (or name concentration)

As discussed above, the ASRF model underpinnings of the IRB capital rules presume that the bank portfolio is fully diversified with respect to individual borrowers. When there are material name concentrations of exposure, there will be a residual of undiversified idiosyncratic risk in the portfolio, and the IRB formula will understate the required economic capital. This form of credit concentration is sometimes known as lack of granularity. This section discusses how to extend the ASRF model to incorporate the effect of granularity.

To fix ideas, consider how economic capital (credit VaR) varies over a sequence of loan portfolios with the following structure: they all contain a number of exposures to similar credits which are all of the same size with the exception of one that is ten times that size. Table 1 depicts the tail of the simulated loss distribution for seven such portfolios of different sizes ranging from 10 to 3000 credits. As the number of credits increases the importance in the portfolio of the single large exposure declines and the economic capital converges to the one corresponding to the infinitely granular case.

Table 1

A stylised example of the effect of granularity on portfolio risk

Number of loans	**10**	**50**	**100**	**500**	**1,000**	**2,000**	**3,000**
VaR(95%)	.0526	.0508	.0459	.0393	.0386	.0378	.0389
VaR(99%)	.5263	.1695	.1009	.0786	.0773	.0762	.0758
VaR(99.9%)	.5263	.1864	.1284	.0982	.0971	.0950	.0947

Note: Credit VaR at the specified level of confidence expressed as a fraction of total portfolio exposure. The calculations assume PD=1% and asset correlation of 20%.

How important is the effect of name concentration on economic capital?

A number of studies produced by the group provide either direct or indirect estimates of the importance of granularity risk for bank portfolios. The effect is clearly more pronounced for smaller portfolios. An indicative calculation of the upper bound of the contribution of idiosyncratic risk to economic capital can be performed by reference to a portfolio having the maximum permissible concentration under the EU large exposure rules.[8] Such calculations give estimates of 13% to 21% higher portfolio value-at-risk for this highly concentrated portfolio versus a perfectly granular one that is comparable in all other dimensions.[9]

For portfolios that are more typical for actual banks, the impact of name concentration is substantially lower. Gordy and Lütkebohmert (2006) use characteristics of loans from the German credit register (including PDs) to compare the effect of name concentration on loan portfolios of the size that can be found in actual banks. For large credit portfolios of more than 4000 exposures, it is estimated that name concentration can contribute about 1.5% to 4% of portfolio value-at-risk. For smaller portfolios (with 1,000 to 4,000 loans) a range between 4 and 8% is more likely.

Methodologies of dealing with name concentration

The various methodologies, proposed by practitioners and researchers, for dealing with name concentration risk can be generally classified into those that are more ad hoc, based on heuristic measures of risk concentration, and those that are based on more rigorous models of risk. Model-based approaches are strictly preferable, as long as they are feasible to implement.

The HHI calculated in terms of portfolio exposures has been used occasionally to measure the distance between a particular portfolio's distribution of exposures from the infinitely granular ideal. The further the HHI of a portfolio is from zero the more concentrated the portfolio would be. It must be noted that the HHI does not measure the increase in credit risk for the portfolio that arises from this lack of perfect granularity. It can only provide a basis for *ad hoc* adjustments to economic capital that attempt to capture this risk. In the stylised setting of Table 1, the loans in the portfolio differ only in EAD and otherwise are homogeneous in their characteristics. When this is the case, the HHI becomes a natural and effective measure of the degree of portfolio granularity. Real-world portfolios, of course, can exhibit marked heterogeneity in PD, LGD, EAD and maturity, and one finds that simple ad hoc measures based on the HHI are unable to capture reliably the impact of granularity on value-at-risk. Weighting the squared portfolio shares by the ratings of the individual obligors may appear to going some way towards dealing with this shortcoming, but lacking the direct link to a formal risk model it can also generate misleading results.

Model-based approaches can deal more explicitly with exposure distribution, credit quality, and default dependencies. They definitely present a preferable option provided that they retain as much as possible the tractability and transparency of simpler ad hoc rules. In model-based methods HHI-type parameters appear in the calculation of the adjustment, but the inputs and the possible weighting are consistent with the overall framework of risk measurement.

The granularity adjustment described and tested in the paper by Gordy and Lütkebohmert (2006) is firmly linked to a risk model. It shares some essential features with the granularity adjustment that was included in the second consultative paper (CP2).[10] It is derived as a first-order asymptotic approximation for the effect of diversification in large portfolios within the CreditRisk$^+$ model of portfolio credit risk. The theoretical tools for this analysis were

---

8 Directive 93/6/EEC of 15 March 1993. An estimate of the HHI for such a portfolio would be about 0.0156.

9 See Duellmann and Masschelein (2006) and Gordy and Lütkebohmert (2006).

10 See BCBS (2001).

proposed first by Gordy (2004) and refined significantly by Martin and Wilde (2002). In addition, the data inputs to the granularity adjustment are drawn from quantities already required for the calculation of IRB capital charges and reserve requirements.

This last point requires some explanation. For the purpose of calculating IRB capital requirements, the identity of the obligor is immaterial, as capital charges depend only on characteristics of the loan and obligor (eg type of loan, PD, LGD, maturity) and not on the name of the borrower per se. This is a great convenience when data on different sorts of exposures are held on different computer systems, as the job of calculating capital may be delegated to those individual systems and reported back as sub-portfolio aggregates which can then be added up in a straightforward fashion to arrive at the bank-level capital and reserve requirements. When the objective is to measure granularity, however, borrower identity can no longer be ignored. From the perspective of single name concentration, ten loans of 1 million euros each to ten distinct borrowers jointly carry much less idiosyncratic risk than the same ten loans made to a single borrower. The need to aggregate information across computer systems on multiple exposures to a single borrower is arguably the most significant challenge for banks in implementing a granularity adjustment. It must be noted, however, that this aggregation requirement would be necessary in *any* effective measure of granularity (be it ad hoc or model-based), and so is not a drawback peculiar to any specific methodology. In addition, one might ask how a bank can effectively manage its name concentrations without the ability to aggregate exposures across different activities.

While the adjustment is well-understood in principle, in practice there are challenges in its implementation. Gordy and Lütkebohmert (2006) point out that a number of the shortcomings of the earlier version of the granularity adjustment have been addressed in its revised form:

- The granularity adjustment of CP2 required a first-stage calculation in which the portfolio would be mapped to a homogeneous portfolio of similar characteristics. In the revised granularity adjustment, the heterogeneous portfolio is used directly in the formula. The resulting algorithm is both simpler and more accurate.
- At the time of CP2, capital was expressed in terms of expected losses (EL) plus unexpected losses (UL), whereas the finalised Basel II Framework distinguishes UL capital from EL reserve requirements. The revised granularity adjustment has been adapted for this change in the definition of its inputs.
- In the revised form presented by Gordy and Lütkebohmert, the granularity adjustment provides for the possibility that banks be allowed to calculate the granularity adjustment on the basis of the largest exposures in the portfolio, and thereby be spared the need to aggregate data on each and every obligor. To permit such an option, regulators must be able to calculate the *largest possible* granularity adjustment that is consistent with the incomplete data provided by the bank. Gordy and Lütkebohmert, therefore, construct an upper bound formula for the granularity adjustment as a function of data on the m largest capital contributions out of a portfolio of n (with $m \leq n$). As m grows towards n (ie, as the bank provides data on a larger share of its portfolio), the upper bound formula converges to the "full portfolio" granularity adjustment. The advantage of this approach is that the bank can be permitted to choose m in accordance with its own trade-off between higher capital charges (for m small) and higher data aggregation effort (for m large).
- Work in progress (but not yet complete) is intended to incorporate credit risk mitigation activities in the granularity adjustment. Such activities can decrease name concentration (say, through purchase of credit default swaps on the largest exposures in the portfolio) or actually indirectly give rise to name concentration in exposure to the providers of credit protection.

The accuracy of the granularity adjustment is studied closely by Gordy and Lütkebohmert (2006). Two particular sources of inaccuracy must be considered. First, as an asymptotic approximation, the granularity adjustment formula might not work well on small portfolios. Fortunately, this issue is not a material concern. In general, the granularity adjustment errs on the conservative side (ie it overstates the effect of granularity), but is quite accurate for modest-sized portfolios of as few as 200 obligors (for a low-quality portfolio) or 500 obligors (for an investment-grade portfolio).

Second, the IRB formulae are based on a rather different model of credit risk. As a result, the granularity adjustment entails a form of "basis risk" (or "model mismatch"). Unfortunately, one cannot test the accuracy of the granularity adjustment against the IRB model, because it is not possible to construct a "non-asymptotic" generalisation of the IRB model. This is due to the linearisation of the maturity adjustment, which breaks the correspondence between the IRB formula and the model used in its calibration.[11] In order to minimise the potential inaccuracy due to the differences between the mark-to-market basis of the IRB and the default-mode origins of the proposed granularity adjustment, the granularity adjustment formula is based on IRB inputs (the IRB capital charge in particular) that have been maturity-adjusted. Thus, the output of the granularity adjustment formula is implicitly maturity-adjusted, albeit in a not very rigorous manner. As it is not possible to assess directly the accuracy of the granularity adjustment against the IRB model, Gordy and Lütkebohmert (2006) reinterpret questions of accuracy as questions concerning the robustness of the granularity adjustment to its parameterisation.

The group reviewed two other model-based approaches to the granularity adjustment in the credit risk literature, namely those of Vasicek (2002) and Emmer and Tasche (2003). The intuition behind the Vasicek method is to augment systematic risk in order to compensate for ignoring idiosyncratic risk. An important problem is, however, that the systematic and idiosyncratic components of risk have very different distribution shapes. This method is known to perform poorly in practice. The approach proposed by Emmer and Tasche (2003) is based on the default-mode version of CreditMetrics and so shares the Merton model foundation with the IRB model. In contrast to the approach proposed by Gordy and Lütkebohmert (2006), it does not maturity-adjust the input parameters and does not account for idiosyncratic recovery risk. However, in principle it could be extended to capture both aspects. Its major drawback is that the formula itself is quite complex, especially compared to the one proposed by Gordy and Lütkebohmert.

Finally, Gordy and Lütkebohmert analyse the effect of the chosen exposure size cut-off point on the accuracy of the upper bound formula. It must be emphasised here that the upper bound formula *always* delivers a higher adjustment value, since it overstates the concentration of the exposures below the cut-off point. A bank might select the cut-off point to strike a balance between higher granularity adjustment but lower computational costs. From a regulatory perspective, this choice is not of direct consequence. A higher cut-off will deliver a more conservative measure of name concentration risk. Nonetheless, the bank's choice of cut-off may be indicative of the ease with which the bank's IT infrastructure is able to aggregate exposures by name, which might be interesting information on its own sake.

---

[11] Had the IRB formula not been linearised this way, then one could implement a single-factor version of the underlying model and quite easily test the accuracy of the granularity adjustment. However, the linearisation makes it impossible to "work backwards" to the underlying model. It should be noted as well that the "true" term-structure of capital charges in mark-to-market models tends to be strongly concave, so the linearisation was not at all a minor adjustment.

The inclusion of a modest number of the largest exposures in the portfolio would result in a relatively small deviation of the upper bound calculation from the simplified granularity adjustment. Moreover, this number declines in line with the concentration of the portfolio. For a very granular portfolio (HHI of about 0.0009) the 1.4% largest exposures would be sufficient to reduce this deviation to 10%. For a portfolio that has a HHI close to 0.025 the necessary number of exposures is almost halved at about 0.8% of exposure.

4.2 Sector concentration

Violations of the "single systematic factor" assumption may be more difficult to discern, and also more difficult to address than imperfect granularity. Even within a large single market such as the United States or the European Union, macroeconomic performance tends not to be fully synchronised across different geographic regions. Exposures in foreign jurisdictions are additionally subject to country-specific risks, including transfer risk, social risk and legal risk. Similarly, different industries can experience different cycles. These realities suggest that distinct geographic regions and industries ought to be represented by distinct (though possibly correlated) systematic risk factors. In this case, a particular bank may be heavily concentrated in its exposure to some of these risk factors and lightly concentrated to others. The extent to which a single-factor model (and, by extension, the IRB risk weights) understates economic capital depends on both the degree to which the bank is unbalanced in its geographic and industry exposures and the extent to which geographic and industry risk factors are correlated with one another. This form of credit concentration risk is known as sectoral concentration.

To fix ideas, consider a simple example of a portfolio that consists of two types of credit with similar unconditional probabilities of default but each driven by a separate systematic factor. Assume for simplicity that the portfolio contains a large number of small exposures to each type of credit so that idiosyncratic risk is diversified away. Figure 1 depicts the economic capital associated with different mixtures of type A and type B credits and with different assumptions about the asset correlation between the two types.

Figure 1

Portfolio capital requirements

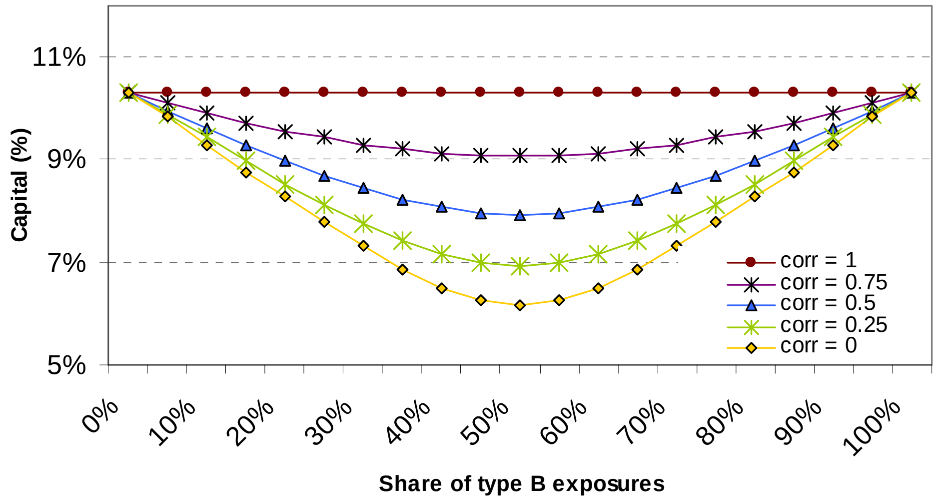

The following points are worth highlighting:

A portfolio that contains a mix of exposures from type A and type B requires a lower capital than a portfolio that is fully concentrated in one or the other exposure type (provided, of course, that the two risk factors are not perfectly correlated).

A single-factor model cannot be expected to capture all aspects of credit risk in a multi-risk-factor environment. Nevertheless, its parameters can be calibrated so that the model delivers economic capital estimates similar to those derived from a richer model for a certain portfolio (the “benchmark” portfolio). In this case, the calibrated asset correlation should be interpreted as an average correlation between two exposures each of which combines different degrees of sensitivity to the two underlying factors rather than the correlation of two exposures that have the same sensitivity to the risk factors.

This basic idea also underlies the IRB model which was calibrated to correspond to economic capital in large credit portfolios that are well-diversified across sectors (risk factors). Therefore, it is in principle possible that such a model produces a higher or lower capital figure than what would be appropriate for a real portfolio. This depends on the diversification of the real portfolio across sectors and on the correlations inside and between sectors.

Finally, note that as the correlation between the sectoral factors increases, the departure from a single-factor model becomes smaller. When the factor correlation is equal to unity then a single-factor model provides an accurate assessment of risk as in this case there is no conceptual difference between a “single-factor” world and a multifactor world with perfectly correlated factors. In the graph the economic capital for different portfolio weights is a straight line, and the economic capital of the portfolio is equal to the sum of the economic capital of the individual exposures evaluated in isolation.

The group has devoted a significant part of its effort to questions related to the topic of sector concentration. The work can be classified into two main categories. The first category focuses on gauging the magnitude of the potential gap between the ASRF model of IRB and “true” economic capital in the presence of credit risk that is fundamentally driven by more than one factor. The second category looks into practical approaches to bridge the gap by various methods that are model-based but remain tractable. Both categories are discussed below.

How important is the effect of sectoral concentration on economic capital?

A number of papers have looked at this question from different angles. The general result is that ignoring the impact of sectoral concentration can lead to a significantly different (higher and sometimes lower) assessment of economic capital.

Two papers produced by the group have looked at this issue using somewhat different data in measuring the co-movement between sectors. The general conclusion in both papers is that asset correlations vary significantly across sectors as well as over time and that, consequently, the magnitude of the concentration risk that is not captured by the ASRF model will tend to be significant and time-varying.

Duellmann and Masschelein (2006) measure the impact of various degrees of sector concentration on economic capital. As input to those calculations they compute equity return correlations for a number of business sectors using the corresponding MSCI indices. They use the Global Industry Classification Standard (GICS) to allocate borrowers to sectors. The aggregate sector distribution of loans to corporate, non-financial borrowers in the universe of loans in the German credit register is used as a benchmark. The aggregate sector distribution of corporate, non-financial exposures in this universe is quite similar to that in the Belgian, French and Spanish credit registers. This suggests a greater applicability of their conclusions to continental European bank portfolios. They create a number of more concentrated portfolios by successively increasing the share of portfolio exposures to a specific sector. In order to focus on the impact of sector concentration they assume an otherwise homogeneous portfolio by requiring that all other parameters are uniform across sectors. Their baseline portfolio assumes that the portfolio consists of exposures of equal size which have a uniform PD of 2% and an LGD of 45%. Their analysis focuses on correlations calculated over a single one-year period (November 2003 to November 2004) and their estimates of inter-sector asset correlations range between 2.5 and 23%. Since they assume a uniform sector factor loading equal to 50%, the implied intra-sector correlations are fixed at 25%.

They find that economic capital increases from 7.8% in the case of the most diversified benchmark portfolio (which corresponds to the composition of the universe of loans in the registry) to 11.7% for the portfolio that is concentrated in a single sector. Two less extreme portfolios are of greater practical relevance as they correspond more closely to the characteristics of mid-sized banks and regional banks. The economic capital is equal to 9.5% for the less concentrated of the two portfolios and 10.7% for the more concentrated one. Finally, they find that for different patterns of the dependence structure between exposures the impact of increased concentration on economic capital might actually be stronger than the numbers above indicate. Clearly, focusing only on exposures to corporate sectors ignores diversification benefits which can arise from a quite often substantial share of retail exposures in banks' portfolios. But still the observed impact of sector concentration in certain real bank portfolios is substantial and typically higher than the impact of coarse granularity, even for mid-sized banks.

The paper by Duellmann, Scheicher and Schmieder (2006) evaluates the impact of sector concentration and granularity. The asset correlations used in their credit risk models were estimated from time series of asset value returns of 2,000 European corporate names. The asset values are based on the Merton-type asset value model of Moody's KMV and extracted from their database for an eight-year period.

The authors first analyse the pattern of asset correlations in their universe of firms. They calculate intra-sector correlations that average about 10% for the market model, and about 12% for the firms in the same sector in the sector model. For the latter model they can also calculate inter-sector correlations of about 67% between the eight sectoral indices. In addition to analysing the correlations over the whole sample, they examine a series of "sliding" two-year sample periods. This approach is useful because time variation in correlations is an issue that has been highlighted by practitioners as potentially very important. The authors find substantial time variation in asset correlations in the range between 4 and 16% for the market model and an even wider range for the sectoral model. They observe that time patterns of asset correlations tend to be disjoint from patterns in PDs supplied by the vendor.

The authors then proceed to simulate portfolio losses on the basis of the estimated asset return correlation, the PD from the database and an LGD assumption of 45%. The value-at-risk is calculated in two portfolio risk models: Firstly, in a *market model* in which systematic risk is captured by a single "market factor" calculated as the value-weighted average asset return of all firms in the sample. Secondly, in a *sector model* in which the "industry factors" are determined by the same method but using sector-dependent instead of borrower-dependent asset correlations and PDs. Both models are compared with the IRB model for a one-year maturity.

The authors find that the market model produces an estimate of economic capital that is 10% to 90% higher than the sector model, and that this difference substantially varies over time, influenced by variation in correlations and an upward drift in average PD over the particular sample period. This difference can be explained (at least partly) by the empirical observation that asset correlations increase with firm size. This stylised fact is better captured in the market model where borrower-specific correlations are calculated than in the sector model in which asset correlations are averaged across all names in the same sector. Economic capital in the sector model is lower than in the IRB model over the entire sample period. The same is true for the market model for the early part of the sample period. Towards the end of the period, however, when all measures of credit risk appear elevated, the market model produces a higher capital figure than the IRB model. In summary, these empirical results highlight the important role played by correlation estimates and model structure in the measurement of portfolio concentration risk. The potential variability of asset correlation patterns over time is an issue of particular importance. Finally, the authors measure the impact of granularity by comparing a portfolio in which each exposure weight is defined by the firm's total debt, taken from the Moody's KMV database, with a benchmark portfolio in which each exposure is set to one Euro. This rough benchmark for a highly granular portfolio is reportedly sometimes also used by banks. For the point in time with the highest value-at-risk in the sample period the granularity effect in the market model amounts to an increase in value-at-risk of 14% relative to the highly-granular portfolio.

It is also useful to compare the estimates of correlations reported above with estimates reported in the analytical papers published by rating agencies. Moody's and Fitch Ratings have recently reported estimates of intra- and inter-sector correlations for a large number of corporate names used in conjunction with the pricing of collateralised debt obligation (CDO) structures.[12] The estimates vary considerably depending on the method employed. Correlations are generally lower when derived from information contained in ratings transitions. In this case, the reported average figures are 12 and 8% for intra-sector correlations and lower still for inter-sector correlations. When correlations are estimated on the basis of equity market valuations, the resulting estimates are considerably higher. Average intra-sector correlations are reported in the 15–24% range, whereas asset correlations range between 13 and 21% for companies belonging to different business sectors.

Methodologies of dealing with sector concentration

There is a growing body of literature that deals with the question of measuring the role of sectoral concentration on credit risk assessment, either explicitly or implicitly through the analysis of multi-factor portfolio models. For the purposes of this note it is helpful to distinguish between two types of approaches. The first approach comes from the realisation that risk is inherently multi-dimensional and focuses on developing multi-factor models. The thrust of this approach is to find ways to overcome the reliance of the models on Monte-Carlo simulations that are portfolio specific and not easy to generalise and to validate. The starting point of the second approach is that the gap between the economic capital assessed through a multi-factor model and a more parsimonious framework is of second-order importance and can be bridged by adjustments to the economic capital figure obtained in closed-form for the simpler model using readily available inputs. Examples include: the binomial expansion technique, the infection model, and the diversity score model.

Multi-factor models

A multi-factor model is the theoretically correct and most general approach to deal with the potential shortcomings of the ASRF model. A major drawback is that most multi-factor

[12] For further details see Moody's Investors Service (2004) and Fitch Ratings (2005).

models typically do not admit a tractable, closed-form solution and require a numerical solution such as Monte-Carlo simulation. Simulations, however, have non-trivial computational requirements and their outcome is always inextricably related to characteristics of the particular portfolio used in the analysis. There is, therefore, value in developing techniques that may overcome these difficulties.

One such approach is proposed by Pykhtin (2004). He demonstrates that following a strategy similar to that of Gouriéroux et al (2000) and Martin and Wilde (2002), one can obtain a closed-form solution for a multi-factor model by accepting a few simplifying assumptions and an approximation to the full-blown solution. He models risk as driven by a number of "sectoral" factors which are common to all exposures within a sector, and an idiosyncratic component corresponding to each individual obligor.

The methodology approximates the economic capital calculated on the basis of a full-blown multi-factor model with two components that have analytical expressions. The first component is an extension of the economic capital as calculated through the ASRF model with one important difference: each exposure is allowed to have a different correlation with the (single) systematic factor. With this exception and ignoring the maturity adjustment, the calculation of economic capital for the portfolio proceeds from the bottom-up. The second component, referred to as the multi-factor adjustment, is more directly related to the fact that the underlying risk is driven by several factors. The required inputs for both components are the following: (i) the factor correlation matrix, (ii) the factor loadings for each exposure, (iii) the PD and expected LGD for each exposure, and (iv) the relative exposure size for each element in the portfolio.

The paper by Duellmann and Masschelein discussed above uses the Pykhtin model after making some important simplifications that greatly reduce the data and computational burden with only limited adverse impact on accuracy. In particular, they replace the borrower-specific data for PD, LGD, asset correlations and relative exposure by sectoral averages of these parameters. They then use sector distributions derived from the German central credit register to measure the relative performance of the (simplified) Pykhtin model on realistic bank portfolios. They evaluate the incremental improvement over the ASRF methodology in matching the multi-factor economic capital of the portfolio by employing the two components of the Pykhtin methodology separately. They find that for portfolios with relatively granular and homogeneous sectors the first component of the Pykhtin model, namely the ASRF model extended by allowing sector-specific correlations, provides a quite accurate estimate of the "true" economic capital (computed by simulations based on a multi-factor model). The incorporation of the multi-factor adjustment component further improves the approximation of the multi-sector simulation model, but its marginal contribution is smaller. This marginal contribution becomes more important for low factor correlations. These conclusions hold for portfolios with different patterns/levels of the sector concentration, the number of sectors, the level of average PD, as well as under various sector weight and correlation assumptions.

Finally, the authors analyse the impact of two assumptions that are arguably the most restricting ones: infinite granularity within each sector and a homogeneous PD for all exposures in the same sector. They find that a lower granularity leads to an underestimation of risk whereas neglecting PD heterogeneity causes an overestimation of risk. Their results indicate that for realistic parameter combinations the effect of PD heterogeneity is at least as strong as the impact of granularity, which implies that their model errs on the conservative side for practical applications. This result also holds for the first component of the Pykhtin model. If a higher accuracy is warranted this first component could easily be generalised to a calculation with a borrower-specific PD and exposure size. In this case, however, the estimate would be no longer conservative.

The overall conclusion is that this approach and possible refinements and extensions have the potential to offer a credible alternative to simulation-based assessments of economic capital, at least for diagnostic purposes.

Extensions of parsimonious, closed-form models

The binomial expansion technique (BET) model developed by Moody's in the context of credit risk analysis for CDOs exemplifies the practical orientation of this type of approach. It consists of a mapping of an actual portfolio with potentially complicated credit risk dependencies across individual exposures onto a hypothetical portfolio of homogeneous uncorrelated exposures that has similar properties for the purpose of assessing economic capital. The mapping is performed by calibration of two parameters in the hypothetical portfolio. The first is the (common) PD for the exposures, which is set equal to the average PD of the actual portfolio exposures. The second is called the diversity score and refers to the number of uncorrelated exposures (of equal size) necessary to form the hypothetical portfolio. This parameter is calibrated by equating the second moment of the loss distribution of the actual and hypothetical portfolios. The assessment of economic capital (EL plus UL) for the hypothetical portfolio is greatly simplified by virtue of its homogeneous structure and the assumption of independence of risks.

The infection model described by Duellmann (2006) extends the basic idea behind the BET by incorporating the suggestion of Davis and Lo (2001). This entails the introduction of a richer, but still tractable, structure of risk dependencies within the hypothetical portfolio. Namely it allows for the possibility that credit risk is correlated across exposures but restricts this correlation to be constant across all exposures. The calibration of the hypothetical portfolio requires the calibration of one additional parameter compared to the BET: the "infection" probability between exposures. The contribution of Duellmann (2006) is to devise a general correspondence between this parameter and a number of observable and easily measured characteristics of the actual portfolio: the HHI of the sectoral concentration of exposures, and the inter-sector and intra-sector asset correlations. This correspondence is performed by means of a (log-linear) regression using data that were generated from a series of portfolios with different underlying characteristics that span the range of values for the three parameters' normal expected range in real-life bank portfolios. The goodness-of-fit is very high auguring well for the applicability of this mapping in other portfolios.

Duellmann (2006) proceeds to evaluate the ability of the BET and the infection model technique to match the economic capital requirement for a number of portfolios constructed drawing from the German credit register. He finds that the accuracy of the infection model is superior to that of the BET approach in terms of value-at-risk with an average error of the order of 5% compared to the full model as opposed to 30% for the BET. The performance of the infection model is less satisfactory when PDs are low and asset correlations are high.

The approach by Garcia Cespedes et al (2005) shares important characteristics of the closed-form approaches with characteristics of the approximations to multi-factor asset value models described in the previous subsection. It provides a closed-form solution of economic capital by introducing a scaling factor in the single-risk factor model which accounts for the sector distribution of the portfolio and correlations between sectors. This scaling factor is calibrated to the economic capital of a multi-factor model. Therefore, this model shares the idea of calibrating its parameters with Duellmann (2006). The authors' starting point is that the single-factor model fails to recognise the gains to sectoral diversification, and so provides a very conservative assessment of economic capital for many real-world portfolios. They construct a *diversification factor* as a multiplicative (downward) adjustment to the capital requirements of a single-factor model. This diversification factor depends on only two parameters and can approximate the outcome of a full blown multi-factor model without recourse to Monte-Carlo simulations. The two parameters they focus on are similar to those used by Duellmann: a capital diversification index which is a measure of the dispersion of exposures at the sectoral level and an average inter-sector asset correlation. However, both parameters are calculated slightly differently from Duellmann (2006). The dispersion is measured by the HHI applied to the economic capital of sectors obtained from a single-factor model and the average factor correlation is computed as a capital weighted average of sector factor correlations. These sector factor correlations are defined by sector-dependent correlations between sector factors and an economy-wide systematic risk factor. If all sector factor correlations equal one the model collapses to a single-factor model.

Through extensive numerical exercises, the authors estimate a parametric surface of the diversification factor as a simple function of the capital diversification index and the average factor correlation. They find that the out-of-sample fit for the estimated function is extremely good, even on portfolios with marked heterogeneity in factor correlations.

An important question regarding their results, however, concerns the relation between the single-factor part of their model and the IRB formulae. As mentioned earlier, the calibration of the IRB model is based on a well-diversified portfolio (in terms of sectoral and geographical composition) of the type that is typically associated with a large internationally active bank. In their model, the single-factor capital is by construction an upper limit to economic capital since it implicitly imposes perfect correlation across sectors. Consequently, the diversification factor they calculate is always smaller than unity. Therefore, to the extent that the calibration of the IRB model already captures diversification across sectors, the asset correlation parameters from the IRB formulae would need to be scaled upwards before they could be used in their single-factor model in order to avoid a "double-counting of diversification effects". This caveat, however, does not invalidate their general methodology, which points to a way to approximate the economic capital outcomes of a multifactor model by appropriately calibrated functions of a small number of readily available portfolio statistics. A careful calibration of their model with properly selected correlation parameters of the single-factor model could provide tabulated diversification factors dependent only on the capital diversification index and the average factor correlation. Such tables could present a convenient diagnostic tool.

Comparisons between various methodologies

This sub-section provides a brief comparison between the various approaches for dealing with sector concentration that were discussed above and that offer tractable solutions to the approximation of economic capital in a multi-factor setting.

The Pykhtin model is the only one that does not require parameter calibration through simulations. The relationships of the infection probability in the infection model and the diversification factor in the Garcia Cespedes et al model with the underlying observable characteristics of the portfolio (such as the average PD, HHI etc) require such calibration. However, in both cases to the extent that the calibration covers a reasonable range of values for those characteristics encountered in actual practice, it is a one-time exercise which does not need to be repeated. Though neither approach is entirely ready for application in practice, the papers by Duellmann and by Garcia Cespedes et al demonstrate how simple tools for economic capital assessment can be calibrated. Both methodologies have the potential to become more accurate, for example, in the first model by exploring alternative loss distributions.

By contrast, the Pykhtin model is more demanding in terms of its input requirements regarding asset correlations. It requires the complete structure of intra- and inter-sector correlations while the other two models rely only on average values of those correlations. The Garcia Cespedes et al model would need to be recalibrated since the asset correlations embedded in the IRB formula for wholesale exposures are arguably too low to reflect asset correlations in a single-factor/sector model. As an approximation to a multi-factor model their model closely resembles the approach by Duellmann and Masschelein. The main advantages would be the less complex formula for economic capital (and also marginal risk contributions) and the possibility of offering tabulated results. The approach by Duellmann and Masschelein has two relative advantages: Firstly, their model does not require and, therefore, does not depend on a numerical calibration. Secondly, it offers greater flexibility in terms of its input parameters, in particular the dependence structure between sectors. First tentative out-of-sample comparisons, which require further confirmation, reveal, however, a remarkable performance of the Garcia Cespedes et al model compared with the more flexible model of Duellmann and Masschelein.

In both models, comparisons with the IRB-implied capital requirements are, however, complicated by the fact that the IRB model framework and associated parameters are not clearly located within those models. The Garcia Cespedes et al model is similar in form to the model underpinning the IRB formula (though it should be noted that the former is default-mode while the latter is mark-to-market). However, the calibration of asset correlations in the two approaches may differ markedly. The Garcia Cespedes et al paper assumes that the single-factor asset correlation has been calibrated to the average intra-sector correlation, whereas the IRB formula has been calibrated to the overall average asset correlation in a large economy. In general, the IRB asset correlation will understate the average intra-sector correlation, and so the diversification factor of the Garcia Cespedes et al approach will overstate the capital relief due to an IRB bank from sectoral diversification. The question of proper "benchmarking" of concentration risk adjustments to Pillar 1 requirements is an important open issue which is briefly discussed in section 6 of this paper.

4.3 Contagion

A third possible source of concentration risk in bank portfolios is through exposures to independent obligors that exhibit default dependencies which exceed what one should expect on the basis of their sector affiliations. These dependencies might arise in the context of business inter-connections (such as supply chain links or counterparty exposures) which are atypical for the respective sector of these obligors. These links may lead to default contagion, or put differently, in the probability of an obligor's default conditional on another obligor defaulting being higher than the unconditional probability of default for the same obligor. Conceptually, contagion risk can be thought of as a half-way situation between name and sector concentration. The default dependency is driven by systematic links between two obligors, but these links are not captured by the overall sector structure. Alternatively, it can be thought as a systematic dependence of one obligor's default on another obligor's idiosyncratic risk. One could argue that contagion risk can provide a way of approximating a more complex sector structure with a simpler one, by allowing some of the residual co-movement to be accounted for in the form of this type of risk. This complexity might be due to a large number of sectoral factors that are not easy to identify with the relatively limited span of credit data that are typically available. Alternatively, it might be due to a fundamentally non-linear structure of risk factors that can be approximated by an extra set of parameters.

The issue of contagion in the context of credit risk has received only scant attention in the literature. Portfolio credit risk models used by the industry do not allow for contagion through business links. In particular, they regularly rely on the assumption of conditional independence; if this is violated and there remains dependence not captured by the model, the loss simulations will produce underestimated measures of economic capital. Recent academic research has raised doubts that the assumption of conditional independence describes the real world sufficiently well.[13] Although various recent and more general approaches have been put forward, this remains a field of ongoing research which is still far away from forming a common standard. Contagion models are one research stream in this area.

Contagion has drawn more attention in recent academic literature on dependence in credit risk. In terms of the empirical estimation of contagion effects the literature is very limited. Egloff et al (2004) use micro-structure information available in one bank to show that business dependence significantly increases the correlation between debtors and fattens the tail of the portfolio loss distribution. This issue is also discussed in the paper by Fiori, Foglia and Iannotti (2006) which associates historically observed sectoral default rates with macroeconomic variables. The paper finds that the explanatory power of macro factors for defaults is relatively limited, but that residual cross-section correlation of default rates suggests the presence of contagion effects from the impacts of sector-specific risk on the default rates of other sectors.

From a practical perspective, it is very difficult to see how contagion risk can be addressed in the context of actual bank portfolios. The required information on bilateral business links is not usually captured by existing information systems of banks. The necessary inputs require tapping what is often called “soft information” that typically exists at the level of individual loan officers and relationship managers. Even more problematic, it requires that business relationships of bank’s customers with other firms that may not be among the bank’s clientele are mapped onto the exposures of the bank in the corresponding business sectors. Nonetheless, these models offer useful insight and make clear that supervisors should assess how much of the actual correlation shown in the data is accounted for in banks’ credit risk models.

5. Stress testing

Stress testing is not yet as mature as other disciplines in risk management, and development of stress testing techniques in the industry is still ongoing. The term itself has no unambiguous agreed interpretation. To the contrary, recent industry studies on stress testing such as CGFS (2000, 2001, 2005) and other surveys such as Lopez (2005) show the wide range of applications and practices which are summarised under the term. Generally speaking, stress testing refers to the evaluation of the effects of extreme changes in input data on the object of interest, eg loss or risk of a portfolio. The probability of these extreme changes actually occurring is usually of second order importance. On the one hand, this makes stress tests less dependent on particular statistical assumptions for the input data. On the other hand, it can lead to implausible or unbelievable conclusions. But because stress testing closely links changes in input parameters to changes in results, it can clarify complex relationships and serve as a valuable communication tool.

It is important to differentiate at the outset between two types of stress tests, regular stress tests in which stress is incorporated in the model without changing its structure and stress tests to analyse “model stress”. As long as one does not want to fundamentally question the model, it is well advisable to choose stress scenarios which are consistent with the existing credit portfolio model. Otherwise stress testing results will have little relevance for risk

[13] See eg Das et al (2007) or Collin-Dufresne et al (2003).

management, or might even be misleading. This is not to say that model stress does not make any sense. Rather, one should be clear about whether one believes in the portfolio risk model, or whether the risk model itself is to be questioned. In the following the focus is on regular stress tests, not on “model stress”.

In the following the aim is to discuss some general problems and give an example of how they could be addressed in a stress testing concept for concentration risk. The work on stress testing by Bonti et al (2006) who focus on sector concentration risk is reviewed and name concentration issues are left aside. Sector concentration risk typically arises when a large percentage of the credits in the portfolio under consideration are closely linked, for example because they depend on a common risk factor or on a small set of highly correlated risk factors. A deterioration in these factors can trigger the default of a significant part of the portfolio and thus cause a material loss. These scenarios usually have a low probability of actually occurring. The combination of complex risk behaviour and a large credit portfolio which is driven by a large number of risk factors makes this an ideal candidate for the application of stress testing techniques. The paper by Bonti et al demonstrates how stress testing can help to clarify the impact of individual risk factors on the credit portfolio, identify those risk factors which contribute most to the overall loss distribution and therefore improve the understanding of how sector concentration risk influences the credit portfolio. In practice, it is very hard to separate concentration risk from credit risk – arguably, the largest contribution to credit risk comes from risk concentrations, either in names or sectors. Therefore, the methodology introduced in the paper has to be seen as an integral part of the more general stress testing methodology of credit risk.

5.1 Desirable properties of stress tests

In order to avoid pitfalls in the design of stress tests for sector concentration risk, the following three properties are desirable: Stress tests should be *plausible*, *consistent* with the existing model framework and *adapted to the portfolio and internal reporting*.

Plausibility implies credibility of the stress scenario which is necessary to have an impact on bank’s risk management. It requires that the stress scenario should be believable and have a certain probability of actually occurring. To this purpose a link between the real world and the model world is needed. This link is necessary, for example, in order to translate real world stress scenarios into stress scenarios for the risk factors of the credit portfolio model. Such a link makes it easier to devise plausible stress scenarios and improves communication about stress test results with credit risk management. The distance of stressed risk factors from current market conditions, for example, can give an indication for checking the plausibility of a stress scenario. As a counter example consider a scenario of a uniform increase in the PDs of all borrowers in a certain industry by a factor of 100. Such an event would surely create a large loss, but it does not seem to be a plausible scenario. No risk management actions will be taken based on implausible stress scenarios.

The second requirement is to use a consistent credit portfolio model or quantitative framework which captures and aggregates the relevant risks and serves as the basis for risk management actions such as hedging or exposure management to certain borrowers. Stress testing should give a reliable picture of how a credit portfolio would perform in a crisis situation. As risk management activities are usually based on a particular credit portfolio model, it is important to keep the portfolio model intact as far as possible. To this purpose the stress test should respect historical dependencies (correlations) of risk factors, although correlations can also be regarded as risk factors. All available information (including eg macroeconomic predictions) should be used. Consistency is often achieved when “internal” risk drivers are stressed. For a credit portfolio model this could mean stressing the systematic risk factors.

For a credit portfolio model, stress tests should be seen as a possibility to merge new information, such as risk management insights or economic predictions which are external to the credit portfolio model, with the assumptions and information contained within the existing credit portfolio model.[14] Sometimes stress tests are inconsistent, either with historic market experience or with the chosen risk management model. This applies, for example, to stress tests where certain model inputs are stressed in isolation, such as the PDs of borrowers in a particular sector (often called "sensitivity analysis"). Those situations are rarely observed, not only because the individual stress is unlikely, but also because other model inputs (namely, the PDs of borrowers in other sectors) would be expected to move as well, at least to some degree, and those moves are not captured in the stress test. While inconsistent stress tests can give an indication of the portfolio's sensitivity to a particular risk factor, they do not accurately capture the portfolio's behaviour in a realistic stress scenario. Especially for sector concentration risk it is crucial to take into consideration links such as correlation between different risk factors, because correlated risk factors in combination can generate losses which would not occur solely because of each individual factor on its own.

As a third requirement stress tests should be adapted to the portfolio and to internal reporting requirements. They should be adapted to the portfolio at hand and reflect certain portfolio characteristics. To achieve this requirement stress tests of concentration risk should eg focus on sectors with relatively high exposures and which are highly correlated with other sectors. Adapted for reporting means that the risk management should be able to translate the outcome into concrete actions or portfolio decisions. As an example, this can be achieved by identifying a small set of risk factors which have a high explanatory power. In the following example this is implemented by differentiating between core factors which are stressed and peripheral factors which also move in the stress event but only conditional on the core factors.

5.2 Example for a stress test methodology

One of many possible approaches[15] divides risk factors into "core" and "peripheral" factors, stresses the "core" factors and lets the "peripheral" factors move conditional on those "core" factors. The paper by Bonti et al represents a sample application of this approach. The selection of "core" risk factors for the concentration stress test reflects prior knowledge or guesses about sector concentrations in the portfolio. For example, if a risk manager is concerned about the exposure to the automobile industry, he can choose the corresponding risk factor, say the stock market sector index for the automobile industry, as core factor to be stressed. All other risk factors, such as the risk factor for the chemicals industry, will not be stressed directly, but will still be adversely affected due to their positive correlation with the core factor. In this paper, the authors do not assign just a single stressed value to the "core" factors. Rather, movement of the risk factors is only slightly constrained, depending on the restrictions that are placed on the "core" factors in order to specify the stress scenario and the modelled relationship between risk factors. Consequently, the output of such stress testing is not a single portfolio loss, but a whole portfolio loss distribution. In effect, the stress test shows what the portfolio loss distribution would look like in a "parallel" world where everything is the same except that the "core" risk factors have been stressed. Thus, the established key risk and performance measures such as expected loss, value-at-risk or expected shortfall can be computed just as in the original setup and compared for an easy, condensed and top-level indication of the effects of the stress scenario. Obviously, in order to

[14] See also Cherubini and Della Lunga (1999) and Berkowitz (1999).

[15] See eg Kupiec (1998), also Kim and Finger (2000).

gain the most information from this kind of stress testing the impact on the whole portfolio loss distribution should be considered. Mathematically, the stress test corresponds to calculating a conditional distribution, given the constraints on the "core" risk factors.

The methodology described above shows that it is possible to derive stress tests on the basis of a consistent model and a close link between the model and the real world, which have the desired properties: to be plausible, consistent with the credit portfolio model, adapted to the portfolio under consideration and reportable to senior management. The paper by Bonti et al (2006) gives a specific application to sector concentration risk. In that case, the risk factors of the credit portfolio model will correspond to sectors (countries or industries), so that stress scenarios for the risk factors can be used to identify risk concentrations in those sectors and assess the impact of stress events in sectors where the portfolio is already known to be concentrated.

6. Open technical issues in modelling concentration risk

This project has dealt with a range of theoretical and practical issues related to the measurement of credit risk. In the course of conducting its work, the group also identified a number of open issues of a technical nature which, despite their relevance, were deemed as being beyond the scope of this project. This section briefly lists those issues in order to motivate further analysis.

(i) What should be considered as an adequate sector scheme for the purpose of concentration risk assessment?

The definition of what constitutes a "sector" is a key question for the implementation of many of the techniques and methodologies discussed in this paper. Should exposures be thought as belonging to the same sector because of their similar characteristics or because of their close correlation of asset returns?

In an ideal situation, the modeller has already identified a fixed set of sector-specific factors and the classification of individual credits is based on the degree of systematic similarity between the particular exposure and each of these factors. In this case, the pattern of correlation of the obligor's asset returns with the identified factors would be a natural criterion to determine the sector classification.

In practice, however, the identification of the set of systematic sector factors is not unambiguous. It is typically based either on pure statistical criteria or follows the industrial sector classification of the borrower. The statistical approach builds the sector scheme endogenously on the basis of the pattern of bilateral asset correlations, but mechanically in the sense that it does not necessarily rely on a formal structural model that links credit risk to other economic quantities. It classifies exposures into groups in a way that maximises the similarities within a group (intra-sector correlations) and minimises the correlations across groups (inter-sector correlations). An often used alternative is to classify exposures either by their industrial sector affiliation or by their geographical location which are determined outside the risk model. The sectoral factors are then derived as the systematic drivers of risk in these various groups. Clearly, there is a certain degree of arbitrariness in this case. The classification of industrial conglomerates and multinational companies is a case in point.[16]

In either approach, the question arises as to where should the line be drawn between the degree of homogeneity within the sector and the overall number of sectors. The choice of the appropriate number of sectors represents a trade-off between accuracy (in the sense of recognising finer dependencies) and stability of correlation estimates. If the number of factors becomes too high relative to the number of firms it becomes difficult to reliably estimate the correlation with firms in other sectors.

The work of Morinaga and Shiina (2005) suggests that misclassification of borrowers into correctly specified sectors is more costly in terms of the miscalculation of economic capital compared to correctly classifying borrowers into mis-specified sectors. This suggests that the question of sector definition might not be as important as the consistent use of the scheme. However, more analysis is certainly warranted on this front.

(ii) Definition of a "benchmark" for concentration risk correction

In the case of name concentration a stylised "infinitely granular" portfolio offers a natural benchmark to measure this kind of concentration risk relative to the ASRF model. In the case of sector concentration, however, a comparably obvious or at least a common definition of a benchmark does not exist. This issue arises in the context of applying the various methodologies that account for concentration risk in a multi-factor framework. In particular, can the adjustment computed in this framework relative to a sector-wise well-diversified portfolio be considered as equivalent to the relative difference between the actual portfolio and the ASRF model? How meaningful is a direct comparison of capital figures from a multi-factor model with the ASRF model?

These are broader questions which touch upon a number of implementation issues as well as getting to the core of issues related to model structure and model calibration of the IRB model. The various methods represent different degrees of departure from the IRB model framework and in most cases these differences in framework are very difficult to reconcile. The granularity adjustment comes close with the use of an additive adjustment to the IRB model. By contrast, models that deal with sector concentration need to replace the one-factor structure which is at the core of the IRB model. Therefore, one can measure sector concentration as an economic capital figure but it is difficult to compare it with the capital figure from the IRB model.

It might actually not be either feasible or necessary to identify a benchmark against which concentration risk is defined. Not feasible because the degree of diversification used in the calibration of the IRB model, which would be necessary to identify the corresponding "benchmark model" in a different model setting, is not clearly defined. Not necessary because the tools presented already provide a ranking of credit portfolios in terms of economic capital for sector concentration which is already important information.

[16] The survey of practitioners indicated that another driver in the choice of classification scheme may be the choice of a specific model. For example, users of Moody's KMV tend to prefer the Global Industry Classification Standard (GICS) classification scheme.

(iii) Data-related issues

The group focused its efforts on assessing questions regarding the measurement and management of concentration risk for banks that rely on internal ratings and models. It did not explicitly deal with data-related issues. In any risk measurement application, however, questions related to the availability and quality of data are key parameters of success. This sub-section highlights a few of these questions and briefly discusses their relevance.

Data that relate to the risk parameters of exposures such as PDs and LGDs are important inputs to the techniques analysed by the group. IRB banks are more likely to fulfil the technical requirements regarding this type of data by virtue of possessing qualified internal credit risk models. At the same time, the more advanced IRB banks are likely have more sophisticated internal models than any specific tool considered in this project, which require more refined inputs and may be better at recognising certain exposure characteristics (eg specific collateral, hedging, optionalities etc). In this sense, the main usefulness for regulators of the methods discussed here would be diagnostic. They provide a means of comparison across banks because they offer a consistent framework whereas the internal models differ in various ways (see the summary of the bank survey).

The tools discussed in this paper may also be useful for less sophisticated banks and those implementing the standardised approach. Given that these institutions also tend to have smaller portfolios, concentration risk is more likely to be an important source of concern for them and the benefit from using a model-based tool rather than *ad hoc* approaches could be substantial. Data issues, however, are also likely to be more challenging for these banks. It is conceivable that inputs based on supervisory experience could fill in.

The key operational challenge for banks may be the need to aggregate exposures to risk entities. This issue is particularly relevant for the calculation of the adjustments for imperfect granularity. However, one may argue that this is a general requirement of sound risk management rather than an issue of data availability for the purpose of measuring concentration risk. Furthermore, such an aggregation is already required (at least in some jurisdictions) by the large exposure rules.

Finally, the issue of parameter accuracy and stability is of particular relevance in the case of asset correlation, a key input for tools dealing with the measurement of sector concentration. Correlation structure estimates are often derived on the basis of equity returns and are notoriously volatile. Larger sample sizes are the best way of reducing estimation error, and these will become available as banks and data vendors intensify their data collection efforts over time. Volatility of correlation estimates may also reflect changes in the nature of risk over time. It has been noted in various contexts that asset return correlations tend to be higher during periods of economic and financial stress compared to more tranquil times. Such systematic movement in asset correlations would also imply systematic shifts in the level of portfolio credit risk due to concentration in exposures, thus adding another layer of complexity to the problem.

References

Basel Committee on Banking Supervision (2001): *The new Basel capital accord, consultative document*, Basel, January.

——— (2003): *The new Basel capital accord, consultative document*, Basel, April.

——— (2004): *An explanatory note on the Basel II IRB risk weight functions,* Basel, October.

——— (2006): *International convergence of capital measurement and capital standards: a revised framework, comprehensive version*, Basel, June.

Berkowitz, J (1999): *A coherent framework for stress-testing*, Board of Governors of the Federal Reserve System, working paper, Washington, July.

Bonti, G, M Kalkbrener, C Lotz and G Stahl (2006): "Credit risk concentrations under stress", *Journal of Credit Risk,* vol 2, no 3, pp 115-136.

Cherubini, U and G Della Lunga (1999): *Stress testing techniques and value at risk measures: a unified approach*, working paper, July.

Collin-Dufresne, P, R S Goldstein and J Helwege (2003): *Is credit event risk priced? Modeling contagion via the updating of beliefs*, Carnegie Mellon University, working paper.

Committee on the Global Financial System (2000): *Stress testing by large financial institutions: current practice and aggregation issues*, Basel, April.

——— (2001): *A survey of stress tests and current practice at major financial institutions*, Basel, April.

——— (2005): *Stress testing at major financial institutions: survey results and practice*, Basel, January.

Das, S R, D Duffie, N Kapadia, L Saita (2007): "Common failings: How corporate defaults are correlated", *Journal of Finance*, vol 62, no 1, forthcoming.

Davis, M and V Lo (2001): "Infectious defaults", *Quantitative Finance*, no 1, pp 382-387.

Duellmann, K (2006): "Measuring business sector concentration by an infection model", *Deutsche Bundesbank Discussion Paper (series 2)*, no 3.

Duellmann, K and N Masschelein (2006): "Sector concentration risk in loan portfolios and economic capital", *Deutsche Bundesbank Discussion Paper (series 2)*, no 9 and *National Bank of Belgium Working Paper*, no 105.

Duellmann K, M Scheicher, and C Schmieder (2006): *Asset correlations and credit portfolio risk – an empirical analysis*, working paper.

Egloff, D, M Leippold, and P Vanini (2004): *A simple model of credit contagion*, University of Zurich, working paper, September.

Emmer, S and D Tasche (2003): "Calculating credit risk capital charges with the one-factor model", *Journal of Risk*, vol 7, no 2, pp 85-101.

Fitch Ratings (2005): *A comparative empirical study of asset correlations*, 6 June 2005.

Fiori, R, A Foglia and S Ianotti (2006): *Estimating macroeconomic credit risk and sectoral default rate correlations for the Italian economy*, working paper.

Garcia Cespedes, J C, J A de Juan Herrero, A Keinin and D Rosen (2005): "A simple multi-factor 'factor adjustment' for the treatment of diversification in credit capital rules", *Journal of Credit Risk*, vol 2, no 3, pp 57-85.

Gordy, M (2003): "A risk factor model foundation for ratings-based bank capital rules", *Journal of Financial Intermediation*, vol 12, pp 199-232.

——— (2004): "Granularity adjustment in portfolio credit risk measurement", in G Szegö (ed), *Risk measures for the 21st century*, Wiley.

Gordy, M and E Lütkebohmert (2006): *Granularity adjustment for Basel II*, working paper.

Gouriéroux, C, J P Laurent and O Scaillet (2000): "Sensitivity analysis of values at risk", *Journal of Empirical Finance*, vol 7, pp 225-245.

Kim, J and C C Finger (2000): "A stress test to incorporate correlation breakdown", *Journal of Risk*, vol 2, no 3, pp 5-19.

Kupiec, P (1998): "Stress testing in a value at risk framework", *Journal of Derivatives*, vol 24, pp 7-24.

Lopez, J A (2005): "Stress tests: useful complements to financial risk models", Federal Reserve Bank of San Francisco, *FRBSF Economic Letter 2005-14*, 24 June 2005.

Martin, R and T Wilde (2002): "Unsystematic credit risk", *Risk Magazine*, November, pp 123-128.

Moody's Investors Service (2004): *Moody's revisits its assumptions regarding corporate default (and asset) correlations for CDOs*, 30 November 2004.

Morinaga, S and Y Shiina (2005): *Underestimation of sector concentration risk by mis-assignment of borrowers*, working paper.

Pykhtin, M (2004): "Multi-factor adjustment", *Risk Magazine*, March, pp 85-90.

Vasicek, O A (2002): "Loan Portfolio Value", *Risk Magazine*, December, pp 160-162.