Basel Committee
on Banking Supervision

Regulatory consistency assessment programme (RCAP) – Analysis of risk-weighted assets for credit risk in the banking book

April 2016

BANK FOR INTERNATIONAL SETTLEMENTS

This publication is available on the BIS website (www.bis.org).

ISBN 978-92-9197-461-0 (print)

ISBN 978-92-9197-460-3 (online)

Contents

Glossary

A/E	Ratio of actual (A) outcomes to default estimates (E)
AIRB	Advanced Internal Ratings-Based Approach (for credit risk)
AP	Asia-Pacific region
BCBS	Basel Committee on Banking Supervision
BIS	Bank for International Settlements
CCF	Credit conversion factor
CMG	Capital Monitoring Group
DTI	Debt-to-income (ratio)
EAD	Exposure at default
EL	Expected loss
EU	Europe region
FIRB	Foundation Internal Ratings-Based Approach (for credit risk)
G-SIB	Global Systemically Important Banks
HPE	Hypothetical Portfolio Exercise
IRB	Internal Ratings-Based Approach (for credit risk)
LGD	Loss-given-default
LTV	Loan-to-value
NA	North America region
OBS	Off-balance sheet
PD	Probability of default
PIT	Point in time
QRRE	Qualifying Revolving Retail Exposure
RCAP	Regulatory Consistency Assessment Programme
RWA	Risk-weighted assets
SME	Small and medium-sized enterprises
TTC	Through the cycle
ULF	Undrawn limit factor

Executive summary

Through its Regulatory Consistency Assessment Programme (RCAP), the Basel Committee on Banking Supervision monitors the timely adoption of regulations by its members, assesses the consistency of implementation with the Basel framework and analyses the quality of intended regulatory outcomes. Consistent implementation of the Basel framework is fundamental to raising the resilience of the global banking system, maintaining market confidence in regulatory ratios and providing a level playing field for international banking activities.

This is the Committee's second report on banking book risk-weighted assets (RWA) variation. The first report, published in 2013, focused on probability of default (PD) and loss-given-default (LGD) estimates for sovereign, bank and corporate exposures.[1] This report assesses regulatory outcomes by examining variability in the calculation of banking book RWAs among banks that have adopted the internal ratings-based models (IRB) approach for credit risk. It focuses on two aspects: RWA variability in retail and small and medium-sized enterprise (SME) banking book portfolios; and variability in estimates of exposure at the time of default (EAD) across the entire banking book.

The report's first objective is to identify the main drivers of RWA variation and evaluate their effects. The report distinguishes between risk-based drivers (ie those driven by underlying differences in risk) and practice-based ones (ie drivers that reflect differences in bank practices and regulatory environments). The second objective is to highlight where there is potential to modify current standards, either to reduce practice-based RWA variation or to simplify the IRB capital framework and increase its comparability.

The findings in this report are based primarily on two data collections carried out between September and October 2014. Data on Retail/SME exposures were received from 35 major internationally active banks across 13 jurisdictions. Data on EAD were received from 37 banks across 17 jurisdictions. Information was also gathered during meetings with representatives from a subgroup of banks that submitted data and via a survey of supervisors.

RWA variation in Retail/SME exposures

The retail/SME exposures reviewed in this study comprise four asset types: residential mortgages; SME Corporate; SME Other Retail; and Qualifying Revolving Retail Exposures (QRRE).[2] The analysis of Retail/SME exposures is based on a comparison of PD, LGD and EAD estimates to actual default and loss outcomes. The comparison for a given bank and portfolio is captured in the ratio of actuals (A) to estimates (E) or the A/E ratio. The extent to which IRB estimates align with actual outcomes provides indirect evidence that differences in RWAs are based on differences in risk rather than differing estimation practices. This approach, a "backtesting" exercise, differs from the 2013 exercise for wholesale exposures in the banking book, which directly evaluated differences in RWAs using hypothetical portfolios of exposures that were common across participating banks.

For PD, averaged across all banks for all portfolios, there is a close alignment of actual default experience and the PD estimates assigned by banks for retail and SME portfolios. The ratio of actual defaults to estimated defaults ranges from 0.87 to 0.97, on average, depending on the portfolio. That is,

1 See www.bis.org/publ/bcbs256.pdf.

2 See Basel II paragraphs 328, 273, 330 and 329 for definitions of residential mortgages, SME Corporate, SME Other Retail and QRRE respectively.

on average, banks experience fewer defaults than estimated across the sample period. The ratio however varies considerably across banks – across all portfolios the ratio for some banks is below 0.5, while for others it exceeds 1.5.

The study does not suggest similar alignment between average LGDs and loss rates. Results are mixed between average EADs and loss outcomes. Except for the QRRE asset class, the average A/E ratios for LGD and EAD are typically much lower than for PD. For the non-QRRE asset classes, the average A/E range for LGD is 0.47 to 0.60, and 0.37 to 0.59 for EAD. The lower values of the A/E ratios for LGD and EAD are not surprising. This is because, under the IRB approach, both LGD and EAD estimates should represent what happens in a downturn and so should ideally be compared with losses under downturn conditions, rather than the average year. The analysis would be more robust if longer time series had been available for all participating banks, in order to capture the influence of downturn conditions and full loss recognition on actual loss outcomes. However, few banks provided data going back more than five years. The dispersion in LGD A/E ratios and EAD A/E ratios is comparable to that observed for PD. That is, the A/E range varies by a factor of 3 to 4 across banks, ranging from well below 0.5 in some cases to much greater than 1 in other cases.

The "localness" of retail and SME exposures affects the interpretation of the observed differences in risk parameters and subsequently variation of RWA. This includes (but is not limited to) whether the sample period for a given bank incorporates a period of stress in the credit markets in which it operates. Clearly, if a bank has experienced stress during the sample period, its actual defaults or losses will be higher than if a bank has not experienced stress. This is borne out in the data. However, even when controlling for whether or not a bank experienced stress, there remains substantial variation in A/E ratios across banks (see Annex 4).

The main sources of practice-based variation include differences in: (a) interpretations of what is meant by "long-run averages" with respect to PD estimation; (b) methodologies for applying cyclicality adjustments to PD estimates (ie rating philosophy); (c) approaches to capturing downturn loss observations in LGD and EAD estimates; and (d) approaches to unresolved cases when modelling recoveries to develop LGD estimates. These and other areas where there is potential to improve policy are highlighted in the report.

As in the 2013 report, this study includes an analysis that shows how capital ratios would change from a 10% benchmark if banks adjusted their RWAs to the average risk-weighting behaviour of participating banks. RWAs are adjusted so that the A/E ratios for RWAs[3] are the same for all banks. The size of this adjustment indicates the deviation from the average risk-weighting behaviour (represented by the average A/E ratio for RWA) expressed in terms of bank capital ratios. If the actual outcomes reported by banks were a perfect proxy for risk, a strong assumption made throughout this report, this variation could be attributed to differing practices. This analysis shows that the capital ratios of most participating banks would lie within a 2 percentage point range (ie a 9% to 11% risk-based capital ratio, or 1 percentage point either side of the 10% benchmark). At the extremes, the variation in behaviour could move a 10% capital ratio by 2 to 3 percentage points (implying a range of 4 to 5 percentage points). Outliers identified in this analysis are often banks that experienced stress conditions in their retail or SME portfolios. Although this analysis is based on strong assumptions, the model performance of these banks would naturally be the focus of further review.

[3] The "A" in this case is derived by combining actual default rates, loss rates and defaulted exposure amounts in the same way that PD, LGD and EAD are combined to produce RWAs under the IRB framework.

Variation in EAD estimates

This study builds on the Committee's 2012 survey of EAD estimates for corporate general-purpose revolving facilities to include an evaluation of variability in EAD estimates for other IRB asset classes and other types of undrawn exposure. Rather than providing detailed information about specific EAD models as in the 2012 survey, participants submitted a range of portfolio-level data from which implied average EAD credit conversion factors (CCFs) could be calculated. The analysis focused on the extent of undrawn exposures within banks' IRB portfolios and the variation in implied average CCFs assigned to a given portfolio or type of exposure.

Advised undrawn exposure (ie amounts where the client has been informed or advised of a credit facility) accounts for a substantial proportion – nearly 30% – of total IRB exposure. Undrawn lending limits make up a little over 80% of that advised undrawn exposure, with issued off-balance sheet (OBS) items, such as trade letters of credit or performance guarantees, accounting for the remainder. Undrawn exposures are most significant for the corporate asset class. The QRRE (mostly credit cards) and residential mortgage asset classes also contain material undrawn exposures.

The study shows there is wide variation in implied average CCFs across all IRB asset classes. This variation is affected by estimation practices. Most significantly, with respect to what is permissible in the area of unconditionally cancellable commitments, a wide spectrum of views was evident. In particular, a number of banks apply zero CCFs to substantial portions of exposures despite the lack of empirical support for doing so. There are varying interpretations concerning how supervisory CCFs should be applied, which likely extend more generally to Foundation IRB (FIRB) and standardised banks. Another key finding is the common practice of using SME and mid-market data to derive CCF estimates for low-default exposures such as sovereign and bank obligors. The study suggests that the CCFs underlying EAD estimates can contribute materially to overall RWA variability, depending on the types of exposure held by a bank.

Several estimation issues were identified that are likely to contribute materially to differences in CCFs and therefore EAD estimates among IRB banks. These include: (a) potential estimation biases resulting from including restructured products, known problem obligors or facilities that are close to fully drawn in the reference data without adjustment; (b) the use by different banks, and within the same bank, of different estimation approaches; (c) varying practices by banks to adjust CCF observations when developing and validating EAD models (eg applying caps or floors); and (d) the use by some banks of estimators other than the mean. These and other areas where there is potential to improve policy are highlighted in the report.

Sound practices on model validation

The Basel II framework requires banks to have in place a robust system to validate the accuracy and consistency of rating systems, processes and the estimates of all relevant risk components in IRB models. A survey of supervisors identified some fundamental components of robust independent validation functions in banks, as recognised to be sound practices by the Committee's members.

The sound practices described in this report can help support a model validation framework that encourages and promotes effective challenge, which is an important focus for supervisors in reviewing banks' independent validation functions. This evidence can take different forms, such as documentation for the rejection of inadequate models or the imposition of more prudent calibrations. A low volume of critical findings made by supervisors after the models have been approved by the independent validation function can also serve as evidence of effective challenge. The practices covered in this paper are grouped into three categories: governance; methodology and scope; and interaction across different phases of model development and implementation.

Although there may be legitimate reasons for validation practices to differ from bank to bank, the Committee believes there is a set of sound practices for model validation that are both consistent with the Basel framework and, where properly applied, may reduce some of the variation in practice that otherwise might exist. The adoption of sound model validation practices can be expected to promote more robust IRB estimates. Therefore, harmonisation in the area of model validation could ultimately lead to reductions in practice-based RWA variation.

Chapter 1: RWA variation in retail and SME exposures

1.1 Introduction

This chapter describes the part of the study that analyses RWA variation in retail and SME exposures. It relies on comparisons of actual default and loss outcomes, measured over multiple periods, to IRB estimates. Underlying this analytical approach is the assumption that actual outcomes are a good proxy for inherent risk. To the extent that this assumption holds, variation in the actual-to-expected (A/E) ratios of banks provides a means of identifying potential practice-based variation in RWAs across banks.

The objectives of the study were, first, to identify the main drivers of RWA variation and evaluate their effect; and, second, to point out areas where there is potential to modify current standards, either to reduce practice-based RWA variability or to increase the simplicity and comparability of the IRB capital framework.

Many factors can influence the behaviour of banks' retail and SME exposures. These include not only the risk characteristics of a borrower but also other relevant risk factors tied to the location of the transaction. IRB risk estimates that take these factors into account will vary across portfolios and across banks even if banks apply the same modelling choices to develop the estimates. Analysis revealed only weak relationships between observed actuals and information collected on characteristics that would be expected intuitively to be fundamental drivers of risk for the respective portfolios. As a result, it was not possible to use this information to control for risk when comparing RWAs across jurisdictions. Moreover, the study's efforts to compare actual outcomes to IRB estimates were impaired by (a) the brevity of time series provided by banks, most of which do not cover a full cycle; and (b) the limited degree of geographical overlap across banks that would allow comparisons of banks in the same market. Because of these data limitations, which are described further in Section 1.5, some of the objectives of the study were not achieved.

This study expands on previous analysis by the Committee of RWA variation in banking book credit exposures.[4] The retail and SME asset classes covered by this study comprise 26% of banks' total RWAs and 37% of RWAs in the banking book.[5] Retail (ie mortgages and QRRE) and SME exposures attract widely varying risk weights across the sample of banks that participated in the study. As shown in Table 1, the average risk weight for the residential mortgage portfolio ranges from 5% to 80%, and from 11% to 82% for QRRE. The variation in these portfolios tends to be larger than that for corporate exposures.

4 For a description of Phase 1 of the Committee's work on this topic, see BCBS, *Regulatory Consistency Assessment Programme (RCAP), Analysis of risk-weighted assets for credit risk in the banking book*, July 2013, www.bis.org/publ/bcbs256.pdf.

5 Together with the Phase 1 study for the wholesale part of the banking book, half of banks' RWAs and 72% of banking book RWAs have been subject to in-depth studies, if one extrapolates the evidence for the available subset to the respective portfolio. These studies have been complemented by top-down analyses of portfolio RWA variation as well as follow-up discussions with selected banks that participated in the data collections.

Average risk weights (RWA/EAD) for corporate and retail portfolios Table 1

	Number of banks	In per cent						
		Mean	Median	Minimum	25th Percentile	75th Percentile	Maximum	Range
IRB	**30**	**29.4**	**28.7**	**12.5**	**25.7**	**33.4**	**52.6**	**40.1**
Corporate	30	46.2	47.6	25.1	39.0	53.9	62.5	37.3
Large Corporate	20	44.8	46.5	25.1	38.2	49.6	61.0	35.9
SME Corporate	14	60.9	59.8	46.2	50.8	63.8	91.2	45.0
Retail	30	25.4	21.0	8.9	17.8	30.3	59.9	51.0
Mortgages	29	24.1	16.9	5.2	14.1	30.7	80.1	75.0
SME Retail	17	47.5	46.8	23.4	38.2	54.2	90.8	67.4
QRRE	25	34.5	33.2	11.2	23.1	39.8	82.5	71.3

Average risk weights calculated by dividing RWA for each asset class by the unweighted exposure of that asset class.

Data as at year-end 2013, collected by the BCBS's Capital Monitoring Group (CMG). Annex 1 shows resulting portfolio distributions based on the data submissions for this study of Retail/SME exposures.

Phase 1 of the Committee's work on RWA variation in the banking book concluded:

- Based on a top-down analysis, more than half of the variation in risk weights for credit risk was attributable to the relative shares of different asset classes held by banks (ie overall banking book risk weights depend on the composition or mix of portfolio exposures), but also showed evidence of practice-based variation that merited more in-depth bottom-up work;
- A bottom-up benchmarking, hypothetical portfolio exercise (HPE) revealed dispersion in the risk weights that banks would assign to commonly held senior unsecured wholesale exposures, suggesting material practice-based RWA variation. These results were attributed in large part to variations in LGD estimates.
- The risk weight variations observed from the HPE suggested that implied capital ratios could vary at the extremes by 1.5 to 2 percentage points around a 10% benchmark in both directions when observed risk weight estimates were extrapolated to other wholesale exposures.[6] However, 22 of 32 participating banks would lie within 1 percentage point of the benchmark (ie 9% to 11%).

1.2 Structure of the backtesting study for retail and SME exposures

The approach taken to exploring RWA variation in the retail and SME portfolios can be described as a backtesting (or benchmarking) study that attempts to establish whether bank IRB estimates (E) have a

[6] See Chart 12 in BCBS, *Regulatory Consistency Assessment Programme (RCAP), Analysis of risk-weighted assets for credit risk in the banking book*, July 2013, www.bis.org/publ/bcbs256.pdf, p 37). This analysis extrapolated the risk weight differences observed in the HPE to all sovereign, bank and corporate exposures, assuming that exposures to all counterparties are equally weighted, and assumed no variation for other assets. The variation at the extremes declines to around 1 to 1.5 percentage points around a 10% benchmark if the risk weight differences are weighted using the distribution of wholesale exposures by risk or PD-band (and assuming no variation for other assets), with a majority of the participating banks lying within 0.5 percentage points of the benchmark (ie 9.5% to 10.5%). This illustrates the sensitivity of results to differing assumptions, including portfolio mix.

reasonable relationship to actual default and loss outcomes (A). Assuming that actual default rates and loss outcomes are a reasonable proxy for risk over a sufficiently long time period, the dispersion in comparison outcomes could suggest possible differences in practices that can lead to RWA variability.

This approach was used because retail and SME markets and products were considered too diverse for an HPE-like study. Retail products in particular are adapted to local markets, local judicial processes and local historical experience, resulting in a diverse array of product structures whose risk parameters are not readily comparable from jurisdiction to jurisdiction.[7] Also, the local nature of retail and SME lending activities limits the geographic overlap of exposures that would be necessary to make more direct comparisons of risk estimates across similar types of portfolio exposure. Hence, the analytical approach chosen for this study is a less direct method of comparing IRB estimates across banks than that used in Phase 1. It relies on the strong assumption that actual default rate and loss outcomes are indicative of risk levels.

The central metric used in the study is the A/E ratio – that is, the ratio of actual outcomes to estimates. The A/E ratios were calculated for each parameter PD, LGD and EAD, as well as for RWAs to provide a composite effect at the asset type level and for retail/SME exposures as a whole. Applied at the parameter level, the A/E ratio indicates the extent to which a bank's actual outcomes exceed or fall below its IRB estimates. Significant dispersion in A/E ratios across banks would suggest differences in modelling practices so long as the central assumption underlying the approach holds: that actual outcomes are a perfect proxy for risk. The approach can also be applied at the level of RWAs by combining actual default rates, loss severity rates and defaulted exposure amounts in the same way that PD, LGD and EAD are combined to produce RWAs under the IRB framework.[8] Dispersion in the A/E ratio at the level of RWAs would suggest differences in RWAs that are driven by different practices (so long as the risk-proxy assumption holds).

The A/E ratio metric was calculated using time series averages for both the "A" and the "E" of the ratio. Alternatively, the calculations could have used differences in outcomes (absolute or percentage) or the variance in outcomes as a measure of "A". Largely due to the brevity of the time series provided by many banks, these alternative measures of "A" were not used. Correlations between A and E could also have been used as an alternative to the A/E metric. The A/E metric was chosen as a more straightforward way to present results. However, the A/E ratio analysis was supplemented by multi-regression analysis to determine the extent to which IRB estimates explained differences in outcomes. These regressions controlled for both country and year-of-observation effects.

Data on retail and SME portfolios were collected from 35 major internationally active banks from 13 countries. These banks covered the Asia-Pacific region (nine banks), Europe (19 banks) and North America (seven banks). Twenty-two G-SIBs participated in the current study, compared with 18 for the HPE study in Phase 1. A total of 25 banks participated in Phase 1 and Phase 2.

The study collected data on four asset types (also referred to as subportfolios): Retail Mortgages, SME Corporate, SME Other Retail and QRRE. Time series data on IRB risk parameters estimates and

7 For example, the European Banking Authority conducted a study in 2013–14 that found significant diversity in markets and retail products across jurisdictions and that banks do not use common sets of risk drivers to develop IRB parameter estimates for such products.

8 For the RWAs, the actuals and estimates were used as inputs to the respective risk weight formulas (eg actual default rates are input as PDs into the risk weight formula), assuming a maturity of 2.5 years. For corporate SME exposures, the size adjustment was assumed to be €27.5 million. Dividing the sum of actual RWAs by the sum of estimated RWAs for each asset class gives an A/E ratio. The aggregate RWA A/E ratio for the retail/SME portfolios is computed based on the respective RWA portion each bank has of these portfolios. Annex 3 contains a description of how A/E ratios were derived for each risk parameter.

outcomes, exposure size, RWAs and relevant risk drivers were collected. Banks were asked to deliver, if possible, estimates for both recent models (ie the models banks used in 2013, the latest reporting date of the study for their regulatory capital calculation, applied to the portfolio in the respective year) and models-in-use (ie the model banks used in the respective year for their regulatory capital calculation). For the PD estimates, 14 banks provided data for both models for at least one of the asset classes, 20 banks provided model-in-use data only, and five banks provided recent model data only.[9] For banks that provided longer time series and for those banks that were significantly affected by the global financial crisis, it was possible to identify periods of stress in the data submissions. Banks also provided qualitative information on modelling choices. For a detailed description of the data collection, see Annex 2.

In March 2015, the Basel Committee team responsible for this work held meetings with a sample of 12 banks participating in the study: three in North America, six in Europe and three in Asia-Pacific. Banks were selected on the basis of the length of time series data submitted, the level of parameter A/E ratios reported by the banks, whether the data suggested stressed conditions in actual outcomes and considerations of balance across regions. The discussions aimed to explore limitations in the ability of banks to provide long-term time series in their submission; to obtain information about economic conditions and the business environment; and to assess how the bank's modelling choices might help explain the results.

1.3 Findings of the study

1.3.1 Overview of findings

The main facts and findings of the study are as follows:

- For PD, averaged across all banks for all portfolios, there is a close alignment of actuals and estimates (see Section 1.3.2). The study does not suggest similar alignment between average LGDs and loss rates (Section 1.3.3) and results are mixed between average EADs and loss outcomes (Section 1.3.4), owing largely to the downturn nature of these two parameters.[10]
- Average PD estimates appear generally higher than the actual long-run average default rates in about two thirds of banks. About 90% of banks' average LGD and EAD estimates are higher than the average actual loss rate and defaulted exposure outcomes.[11]
- Several of the banks with actual outcomes that exceeded the Advanced IRB (AIRB) estimates are from countries that experienced severe downturn conditions, as evidenced by higher than average default rates, in more recent years (see Section 1.3.5). Although the persistence of higher outcomes relative to estimates would be a cause for concern, data limitations generally limited the ability to gauge the extent of this persistence.[12]

[9] Almost all banks reported difficulties in re-estimating parameters from historic data according to the most recent model, except in cases when the difference between the models is minor, such as in a recalibration of the master scale.

[10] Because LGD and EAD are based on downturn estimates, one would ideally accumulate estimates over a long time frame to evaluate the relationship between actual outcomes and estimates. Banks tended to provide a short time series of estimates and outcomes (few banks were able to provide data going back more than five or six years and several banks were unable to provide more than three or four years of data), and many banks provided data over periods that do not exhibit outcomes indicative of downturn conditions.

[11] Averages are calculated by bank and portfolio across all years of data submitted using estimates from the model-in-use at the time, if available, or otherwise estimates from the most recent model for the current and historical periods.

[12] In addition to the shortness of time series provided by many banks, a significant number of banks were unable to provide historical IRB estimates that would have been produced by the most recent models used for IRB estimates (recent model). In other words, banks may have recently recalibrated their models but the result of that recalibration cannot be demonstrated relative to historical periods.

- Unlike low-default portfolios, where differences in AIRB estimates might be expected due to a paucity of loss data, differences in IRB estimates for retail and SME exposures can be influenced by modelling choices that relate to different levels of margins of conservatism and adjustments, differences in the length of data available and included in the calibration (full cycle or part of the cycle), and assumptions underlying recovery estimates. Modelling choices may also be heavily influenced by local conditions and practices (eg business conditions and requirements or expectations of national supervisors).
- Discussions with selected banks that participated in the study revealed the following key practice-based drivers for retail and SME AIRB estimates: different approaches for making PD adjustments to reflect long-run averages (often referred to as cyclicality adjustments); different recovery estimation practices to include varying treatments for unresolved workouts and discount rates; different approaches to calculating downturn LGDs; and differences in the degree and type of data censoring applied to develop EAD/CCF estimates.
- Retail lending is found to be a predominantly local business; very few banks provided data for exposures covering portfolios in more than one country (see Section 1.5). Differences in IRB parameters are to some degree attributable to structural differences between countries such as the characteristics of local markets and the local legal environment.
- The study found no strong relationship between the risk driver information submitted by banks, such as loan-to-value (LTV) ratios, and observed PDs and LGDs. Accordingly, it was not possible to use these risk drivers to control for risk when comparing A/E ratios across jurisdictions.
- Variations from a 10% capital ratio benchmark are determined by adjusting RWAs so that the A/E ratios for RWAs are the same for all banks (see Section 1.3.6). If the actual outcomes reported by banks were a perfect proxy for risk, this variation could be attributed to differing practices. The results of this analysis show that the capital ratios of over half of the banks that participated in the exercise would lie within 1 percentage point of a 10% risk-based capital ratio benchmark. At the extremes, the variation in capital ratios is 2 to 3 percentage points relative to the 10% benchmark. Outliers identified in this analysis are often associated with banks that have experienced stress conditions in their retail or SME portfolios. Although this analysis is based on strong assumptions, the model performance of these banks would naturally be the focus of further review.

1.3.2 PD A/E comparisons

Chart 1 shows the distribution of actual default rates by asset classes, based on the universe of the observed default rates by bank, year and asset class.

Default rates by asset class

Chart 1

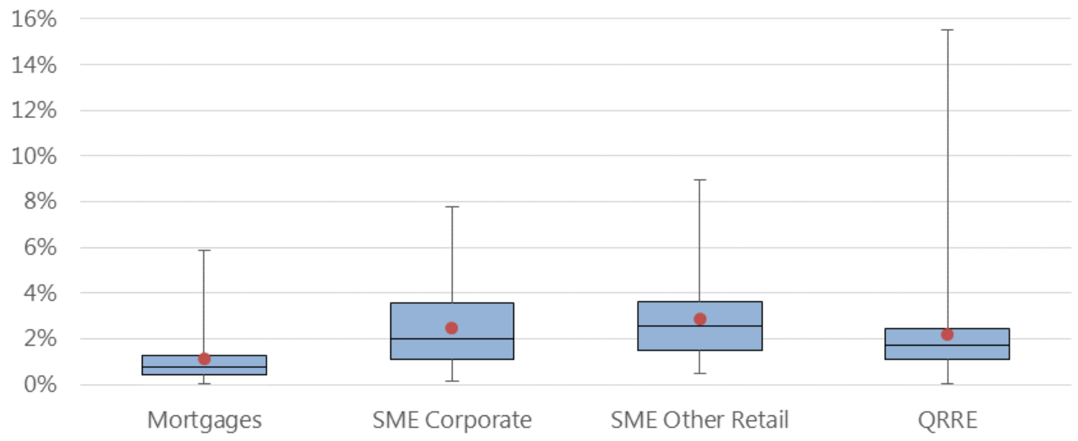

The chart shows the distribution of observed annual default rates in participating banks, by asset class and weighted by performing obligors (rather than EAD). Bars are the maximum and minimum, the box represents interquartile range and the red dot is the mean.

Banks' models appear, on average, to produce PD estimates that are reasonably consistent with outcomes (see Chart 2). The tightest correspondence between actuals and estimates occurs in QRRE. SME treated as Other Retail shows the weakest relationship, at least in more recent years. After aggregating over years and banks, the A/E ratios for low-risk ratings are similar to the A/E ratios for high-risk ratings.

Time series of actual default rates and estimated PD, averaged across banks, by asset class

Chart 2

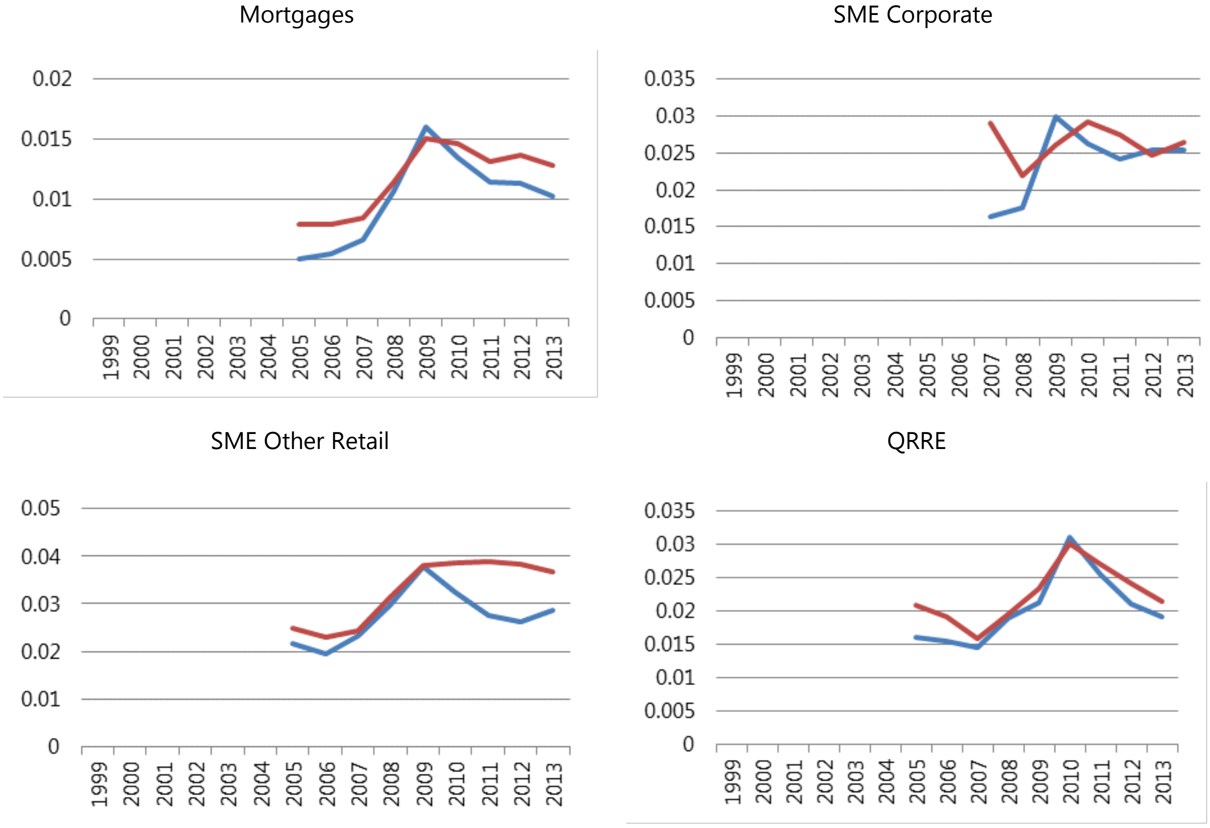

The blue lines show the average default rate (A) for banks that submitted data for the relevant asset class and year, while the red lines show the average estimated PDs (E) used in banks' models.

The number of observations varies across portfolios and over time. From 2008 to 2013, there are typically more than 20 observations per year, whereas prior to 2008 the data are more limited; typically less than 10 but not lower than three.

For most banks, on average, the PD parameters are higher than the actual average default rates. This is shown in Chart 3, where results below 1 indicate that a bank's actual average default rates are lower than the average PD assigned by the bank. The proportion of banks with A/E PD ratios above 1 ranges from 28% for mortgages to 44% for SMEs treated as Corporates. The outliers that appear on Chart 3 – either on the high or low end – can often be explained by stress experience (high default rate outcomes) or benign conditions (low default rate outcomes), respectively.[13]

13 Where banks provided data for both models-in-use and recent models, it appears that some outlier banks have recalibrated the risk estimates using updated data, bringing these banks A/E results closer to average.

Ratios of actual default rates to estimated PD

Chart 3

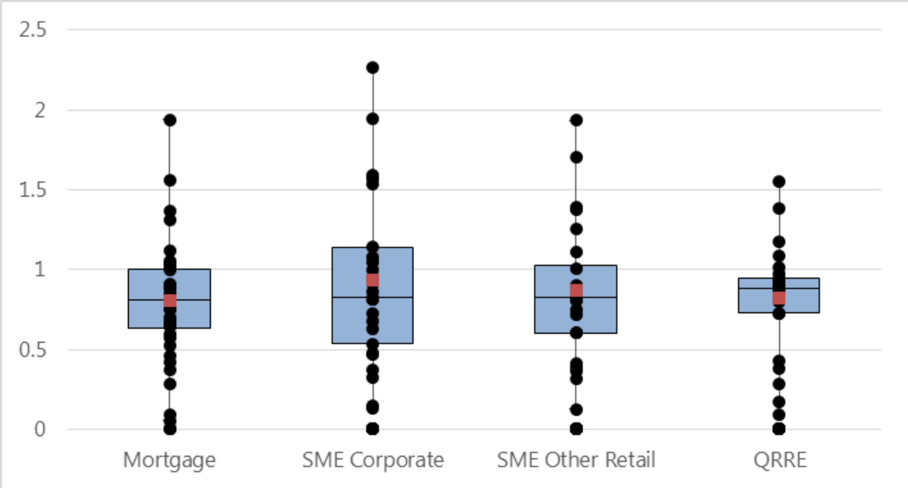

The black dots show the ratios of each bank and the red dot is the mean. The box represents the interquartile range.

Results are weighted by the number of obligors and derived from model-in-use observations unless only recent model observations are available. Weighting results by EAD does not materially alter the general observation. Using only recent model observations captures too few observations for some portfolios to be conclusive.

1.3.3 LGD A/E comparisons

Although PDs appear, on average, to align with default rates, the study does not show similar alignment between average LGD and loss rate outcomes (see Chart 4). The distribution of loss rates realised on defaulted exposures is shown in Chart 5, which shows significantly higher variation than that for observed default rates.

Time series of actual losses on default and estimated LGD, averaged across banks, by asset class

Chart 4

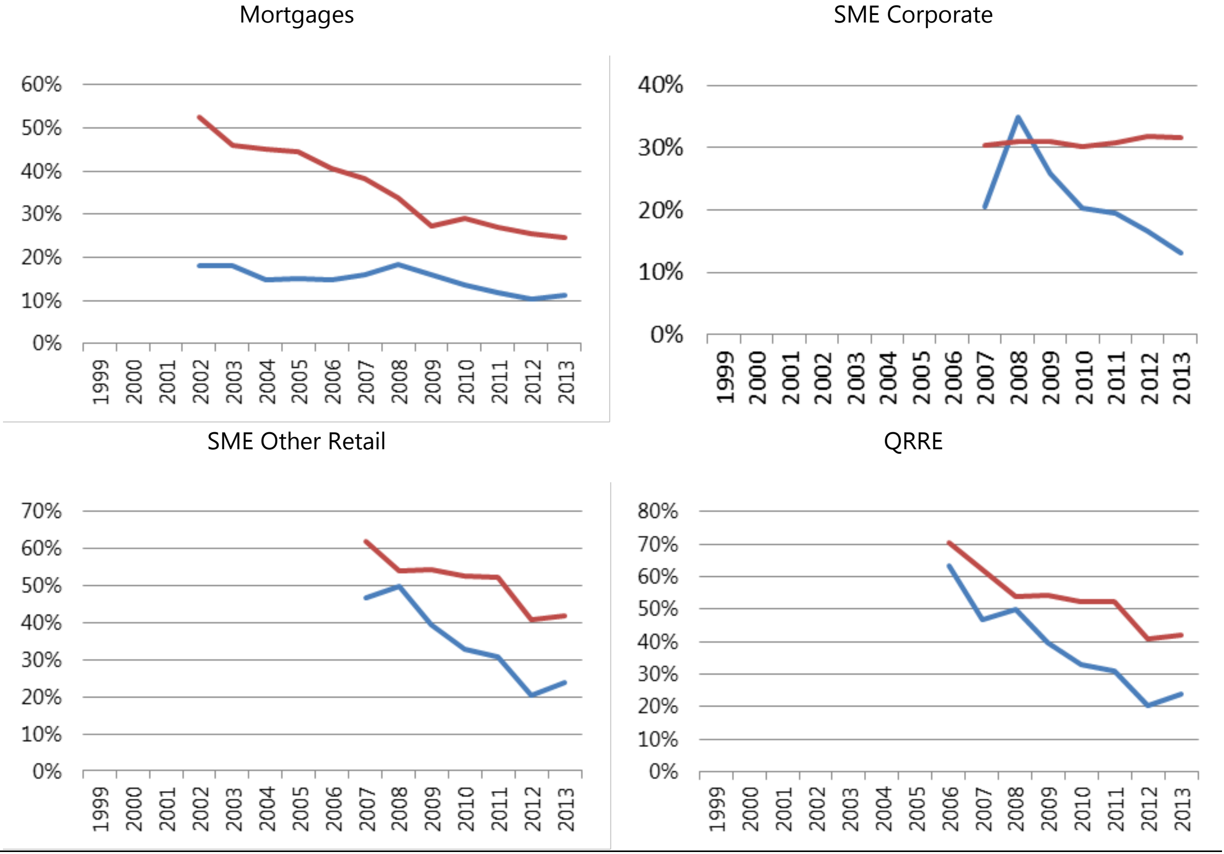

The blue lines show the average realised losses (A) in the case of default for banks that submitted data for the relevant asset class and year, while the red lines show the average estimated LGDs (E) used in banks' models.

The number of observations varies across portfolios and over time. From 2008 to 2013, there are typically more than 20 observations per year, whereas prior to 2008 the data are more limited; typically less than 10 but not lower than three.

Distribution of loss rates realised on defaulted exposures, by asset class

Chart 5

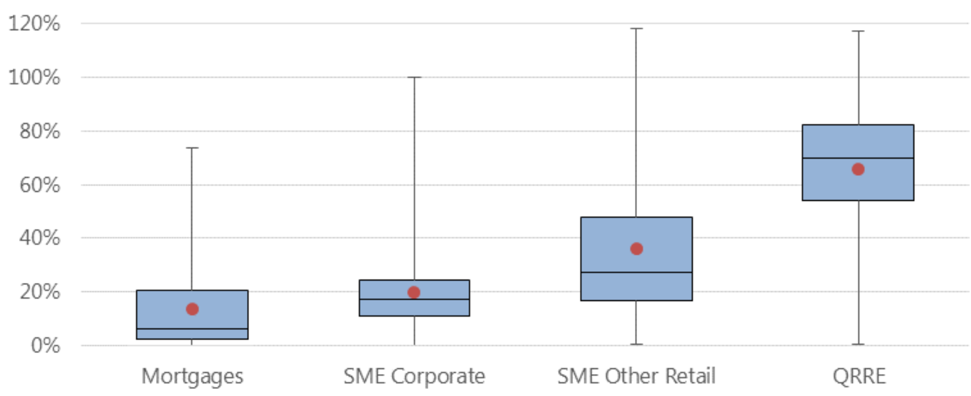

The chart shows the distribution of annual average loss rates realised on default at participating banks, by asset class and weighted by EAD. Realised losses are those applied to fully resolved defaults and defaults with sufficient loss information. Bars are the maximum and minimum, the box represents interquartile range and the red dot is the mean.

Realised losses within a given jurisdiction are likely to be heavily influenced by legal factors pertaining to foreclosure and debt collection practices. These differences should also be reflected in a wide range of LGD estimates, which is in fact what the data show. There is no strong alignment between actual loss outcomes and LGD estimates as there is with average PD actuals and estimates. Following are some possible explanations for this outcome:

- LGD estimates should be based on loss rates during downturn periods of the business cycle.[14] Absent many downturn loss outcomes, LGDs consequently exceed actual loss rates in the vast majority of cases, resulting in lower A/E ratios relative to the PD A/E ratios, especially for mortgage lending (see Chart 6). In discussions with banks, most highlighted the lack of guidance in the definition of and inclusion of downturn observations in the estimates. Most banks also highlighted the application of various techniques that result in higher estimates than a simple average would imply. The main methods conveyed include applying a high percentile of a historical distribution, applying a "stressed" discount rate to recoveries, or applying macro-analyses and expert judgment adjustments to collateral haircuts and recoveries.
- Banks tended to provide a short time series of estimates and outcomes, and many banks provided data over periods that do not exhibit outcomes indicative of downturn conditions. Because LGD (and EAD) parameters are based on downturn estimates, it is necessary to accumulate estimates over a long time frame to evaluate the relationship between actual outcomes and estimates.
- In some countries, banks' actual loss experience falls well below the 10% LGD floor for mortgages, resulting in a comparably low A/E RWA ratio for mortgages.
- For some portfolios and some jurisdictions, the workout period for a defaulted exposure is long, covering several years or more. As incomplete workouts from recent years (usually the ones with the highest actual LGDs) are not yet part of the recent "A" of the A/E ratio, the ratio could be biased downwards relative to A/E ratios that would otherwise have been observed if a larger number of more distant observations had been obtained.

On a per-bank basis, the LGD parameters appear generally higher than actual losses. This is shown in Chart 6, where results below 1 indicate a bank's actual average loss rates are lower than the average LGD assigned by the bank. As with PDs, there is wide variation in A/E ratios for LGDs across banks. As with PDs, the outliers in Chart 6 can often be attributed to whether a bank experienced stress or not.

[14] Basel II, paragraph 468.

Ratio of Actual Loss Rates for defaulted exposures to estimated LGDs

Chart 6

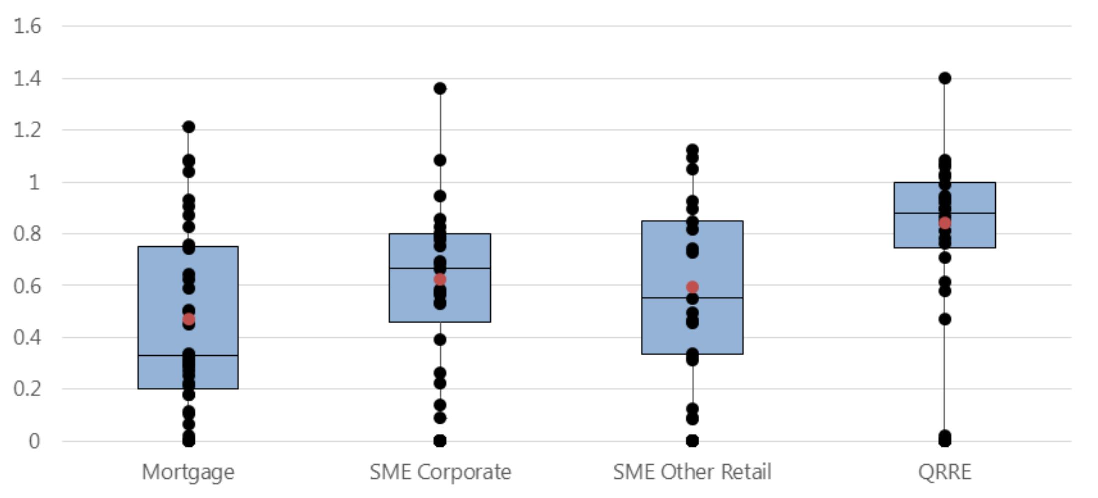

The black dots show the ratios of each bank and the red dot is the mean. The box represents the interquartile range.

Results are weighted by EAD and derived from model-in-use observations unless only recent model observations are available.

1.3.4 EAD A/E comparisons

The alignment between average EAD and loss outcomes is mixed (see Chart 7); mortgages and SME Other Retail show a reasonable alignment whereas SME Corporate and QRRE do not. For credit conversion factor (CCF) estimates, most banks reported estimates that exceeded defaulted exposure outcomes on average; however, there are significant outliers, particularly in the QRRE and mortgage portfolios (see Chart 8). As with LGD, the downturn nature of this parameter helps explain this result. However, it should be noted that the data quality was lower and far fewer observations were used in the analysis. Therefore, more evidence would be needed before firm conclusions could be drawn.[15]

[15] EAD estimation issues identified from two EAD surveys are discussed in Chapter 2 of this paper.

Time series of actual exposure at default and estimated EAD, averaged across banks, by asset class

Chart 7

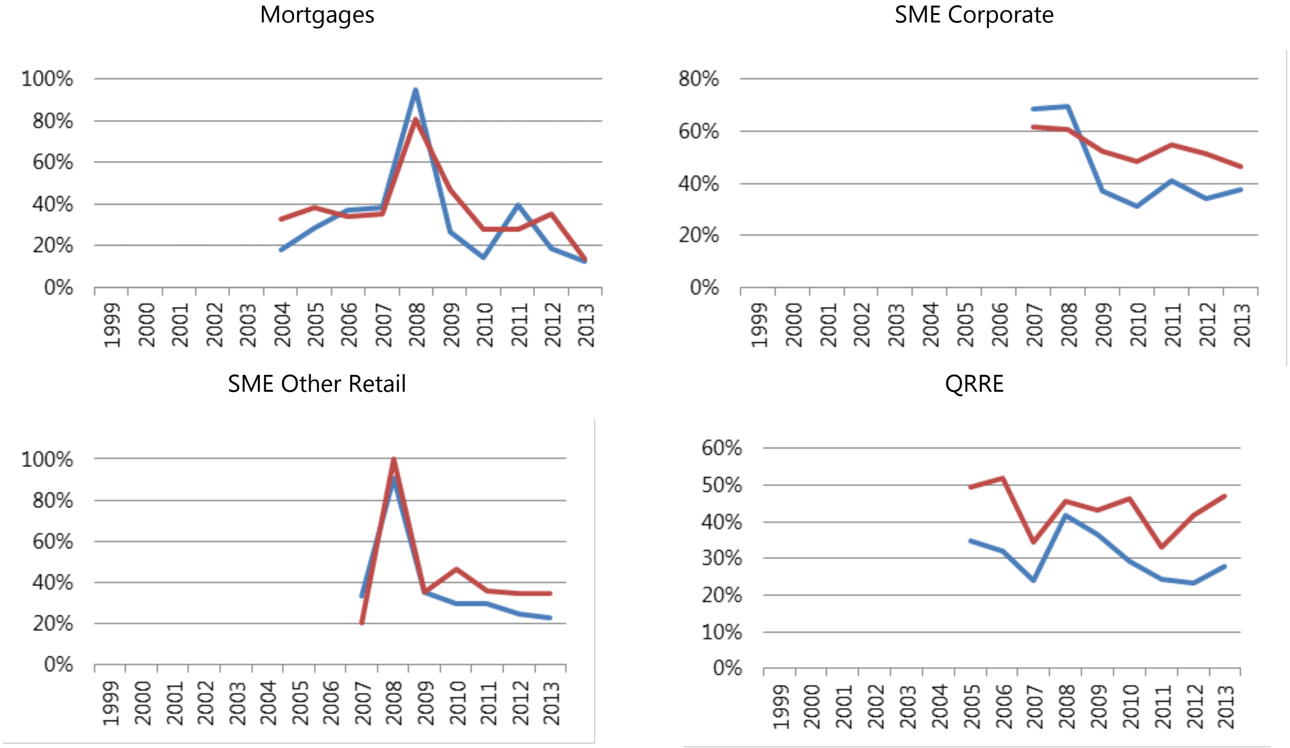

The blue lines show the average actual exposure at default (A) in the case of default for banks that submitted data for the relevant asset class and year, while the red lines show the average estimated EADs (E) used in banks' models.

The number of observations varies across portfolios and over time. From 2008 to 2013, there are typically more than 20 observations per year, whereas prior to 2008 the data are more limited; typically less than 10 but not lower than three.

Ratio of Actual Losses (Relative to Undrawn Defaulted Exposures) to estimated CCFs

Chart 8

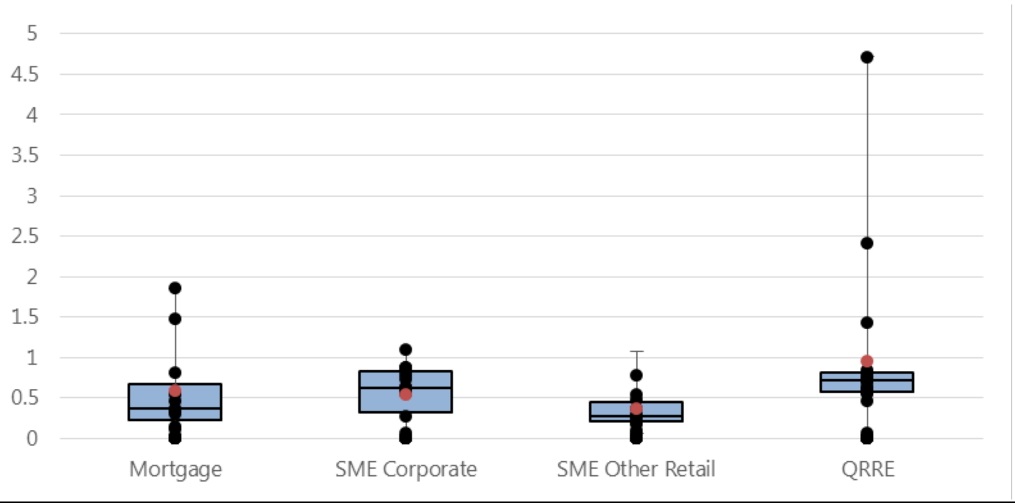

The black dots show the ratios of each bank and the red dot is the mean. The box represents the interquartile range.

Results derived from model-in-use observations unless only recent model observations are available.

1.3.5 Effect of stress conditions on observed A/E ratios

For banks that provided longer time series and for banks significantly affected by the global financial crisis, it was possible to identify periods of stress in the data submissions (see Table 2). These stress periods are defined by unusually high actuals that are higher than might be expected on average over a "normal" business cycle. For this subsample of banks with stress experience, defined as high percentiles in terms of all actuals experienced across banks, it is possible to study the impact of such observations on the A/E ratios.

For this exercise, periods of stress were assumed to happen roughly once in every seven years. Consequently, a threshold for identifying stress years was set at the 86th percentile of the respective empirical distribution of the default rates by asset class. This percentile of the distribution corresponds to a *worst-in-seven* year level for default rates, namely 1.7% for mortgages, 4.4% to 5% for the SME portfolios and 3% for QRRE exposure. Using these stress-level thresholds, banks subject to stress included those in the Asia-Pacific (AP) region for SME corporates and QRRE, those in the Europe (EU) region for all portfolios, and those in the North America (NA) region mainly for mortgages and QRRE. Around 25% of banks are affected by stress across the asset classes, except for QRREs, where 40% of banks are affected.

Number of banks affected by stress by region — Table 2

	Years with observed stress level			
	Mortgages	SME Corporates	SME Other Retail	QRRE
Number of banks: total	8 of 36	7 of 25	5 of 24	10 of 25
AP banks	0 of 7	1 of 8	0 of 7	2 of 6
EU banks	4 of 22	5 of 14	4 of 13	5 of 12
NA banks	4 of 7	1 of 3	1 of 4	3 of 7
Default rate at 86th percentile	1.7%	4.4%	5.0%	3.0%

Table 3 shows the average A/E ratios by portfolio and risk parameter as well as the number of banks with A/E ratios above 1 and the subsample of those banks that experienced elevated default rates. To illustrate, 10 banks exhibited mortgage A/E ratios above 1, of which four experienced stress. This implies that half of the eight banks (from Table 2) that experienced stress exhibited A/E ratios above 1. For the SME asset classes, this portion is even higher (six out of seven for SME Corporate and four out of five for SME Other Retail). For QRRE, the proportion is 30% (three out of 10). In sum, the outcome suggests that a significant portion of the elevated A/E ratios (especially for PDs) are driven by banks that experienced stress. Annex 4, which shows a comparison of A/E ratios for banks with and without stress, further demonstrates this finding.

Summary of A/E ratios by asset class

Table 3

Risk parameter	Portfolio	Average A/E	Number of banks	Of which: A/E > 1	Of which: experienced stress on subportfolio level
PD	Mortgages	0.86	36	10	4
	SME Corporate	0.97	25	11	6
	SME Other Retail	0.91	24	8	4
	QRRE	0.87	25	8	3
LGD	Mortgages	0.47	35	4	2 (stress for PDs)
	SME Corporate	0.62	23	3	1 (stress for PDs)
	SME Other Retail	0.60	21	6	0 (stress for PDs)
	QRRE	0.84	24	3	2 (stress for PDs)
CCF	Mortgages	0.59	11	2	1 (stress for PDs)
	SME Corporate	0.55	16	1	0 (stress for PDs)
	SME Other Retail	0.37	14	1	0 (stress for PDs)
	QRRE	0.95	19	4	1 (stress for PDs)

The A/E ratio per risk parameter is first averaged over the time series for a specific bank and asset class and then averaged over all banks. Stress for PDs means worst-in-seven year level for default rates: 1.7% for mortgages, 4.4% to 5% for the SME portfolios and 3% for QRRE exposure.

1.3.6 Variation in capital ratios implied by observed differences in RWA calculation

This study uses the variability in RWA "actuals" to RWA estimates to determine the range of capital ratios implied by the average risk-weighting behaviour of participating banks. This extends previous analysis contained in Phase 1 of the Committee's work on banking book RWA variation.

These calculations are based on the difference between a bank's actual RWA A/E ratio and the benchmark (average) RWA A/E ratio across all banks. First, a 10% capital ratio for the benchmark is assumed. Next, deviations from that benchmark are calculated by comparing each bank's RWA A/E ratio to the average RWA A/E ratio across all banks. Since exposure information was collected, it is possible to calculate the impact without making assumptions about exposure distributions.

On the assumption that actual default and loss experience is a good proxy for risk, Chart 10 shows that the capital ratios of over half of the banks that participated in the exercise would lie within 1 percentage point of the 10% risk-based capital ratio benchmark. At the extremes, the variation in capital ratios is 2 to 3 percentage points relative to the 10% benchmark.[16]

As indicated in the chart, many of the outliers identified using this analytical approach are associated with banks that experienced stress conditions in their retail or SME portfolios, which suggests the influence of cyclical economic factors (eg high default rates) on the variability shown. Although this analysis is based on a number of strong assumptions, the model performance of these banks would naturally be the focus of further review to ensure risk estimates are appropriately calibrated to withstand persistent stress conditions.

16 Despite the very different methodology adopted in this study, this range of variation is similar to that found during Phase 1 for sovereign, bank and corporate exposures. See Chart 1 in BCBS, *Regulatory Consistency Assessment Programme (RCAP), Analysis of risk-weighted assets for credit risk in the banking book*, July 2013, www.bis.org/publ/bcbs256.pdf.

Implied impact of RWA variations from benchmark on capital ratios — Chart 10

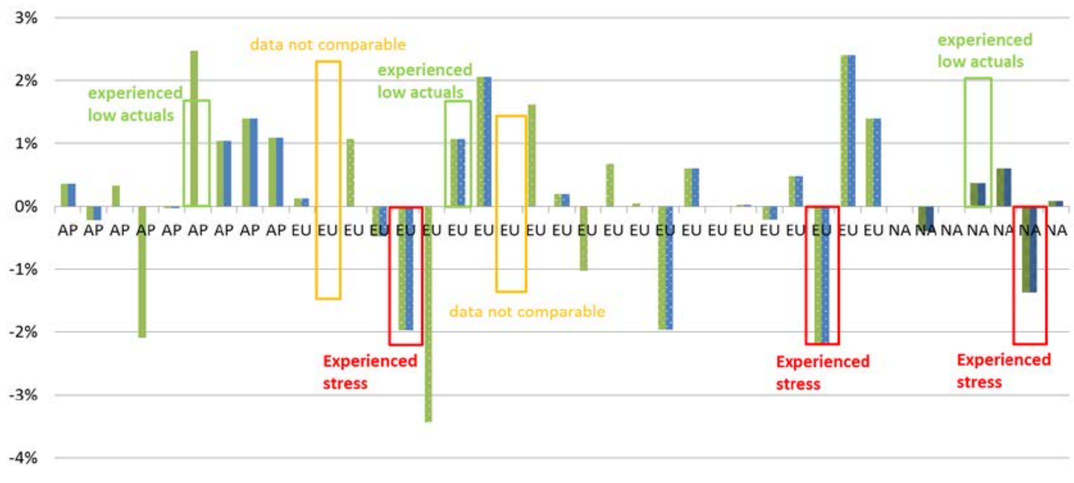

The green bars show the results for each bank based on the RWA A/E data for all years provided. The blue bars use the same data, but only for banks that provided at least four years of data. The horizontal axis shows the region of each bank.

The implied impact on capital ratios is calculated in four steps. (a) Derive actual RWAs by combining actual default rates, loss rates, and defaulted exposure amounts in the same way that PD, LGD and EAD are combined to produce RWAs under the IRB framework. Where applicable, maturity of 2.5 years is assumed. (b) Calculate, for each bank, an A/E ratio for RWAs (actual RWAs as computed in (a) divided by RWAs produced by bank models). (c) Associate the average RWA A/E ratio across all banks with a 10% capital ratio. (d) Determine how a bank's current capital ratio would change if its RWAs were adjusted so its A/E ratio was equivalent to that of the 10% benchmark. The impact on capital ratios is weighted by the RWA portion covered by the exercise.

For example, assume that the benchmark A/E ratio is 0.5 and a given bank's A/E ratio 0.4, with Retail/SME exposures comprising 20% of RWA. The implied change in capital ratio is +0.25% (0.5 / 0.4, x 0.2).

Banks with very low default rates for only one asset class were removed (two banks); banks with missing PDs or LGDs considered based on assumptions; CCFs A/E assumed 1 if missing. See also Annex 3.

The banks labelled with "experienced stress" had A/E ratios for RWAs that exceed 1. Because the risk weight function transforms expected PD into a "stressed PD" at the 99.9% confidence level, an A/E ratio that exceeds 1 does not necessarily imply that capital is less than losses. See www.bis.org/bcbs/irbriskweight.pdf.

1.3.7 Findings from meetings with selected banks

To complete the study, meetings were held during March 2015 with a sample of 12 participating banks.[17] The discussions aimed to explore limitations in the ability of banks to provide long-term time series in their submission; to obtain information about economic conditions and the business environment; and to assess how the bank's modelling choices might help explain the results. These objectives were met and the discussions yielded insights into differences in modelling practice that would likely result in RWA differences irrespective of the degree of similarity in risk levels across banks' portfolios.

As an example, the discussions confirmed the expectation that banks use a variety of methods to calibrate PDs to long-term averages. Most banks define a central tendency for each retail/SME segment and calibrate the master scale and/or the mapping to the master scale so that the segment's PD is equal to the defined central tendency. However, banks use many variants for defining the central tendency.

[17] There were discussions with three banks in North America, six banks in Europe and three banks in Asia-Pacific.

These include using a rolling five-year window of historical results; referencing the latest completed cycle and supplementing data gaps with external data;[18] or using the higher of observed default rates through a full economic cycle or forecasted stress condition default rates.

The range of PD modelling practices discussed by banks could produce material variation in PD estimates, as illustrated in Chart 11 below. Assuming a yearly recalibration is done, banks that use a more point-in-time (PIT)-oriented rating system (ie PDs estimated from a short-term history of default rates) would have capital requirements that react more rapidly to worsening economic conditions than a through-the-cycle (TTC)-oriented rating system (ie PDs estimated over a longer-run average). At the height of stress periods, the TTC PD could produce capital requirements that are well below PIT PDs; whereas TTC PDs remain high as conditions improve but PIT PDs decline. Another relevant dimension in this context is whether or not the estimates are based on a full cycle. Where estimates are not based on a full cycle, some banks adjust the level of the parameters (using past experience from the same country, for example), while other banks do not make such adjustments.

Hypothetical example for capital requirements under a PIT and TTC regime — Chart 11

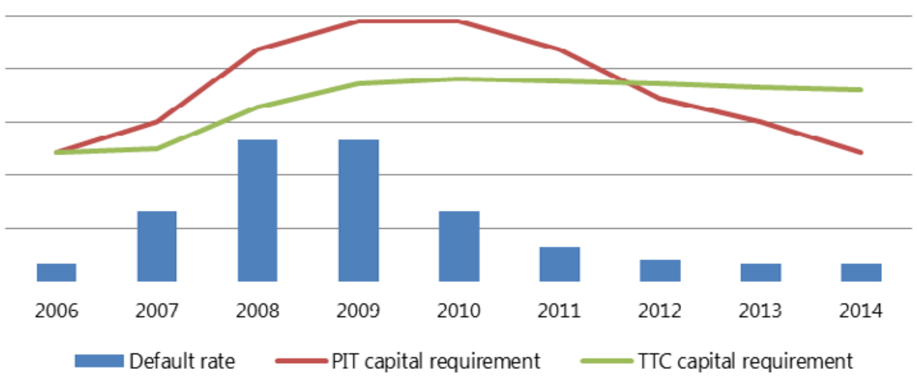

This illustrative chart shows the difference in capital requirements between PIT and TTC PD modelling practices during a period of stress. The capital requirements and default rates are not plotted on the same scale. The figures are illustrative only.

For LGD, all banks involved in the discussions use a “bottom-up approach” that entails estimating recovery cash flows to estimate LGD. However, banks use a wide range of assumptions and modelling choices concerning the recovery process, which include estimating the probability of cure, collateral haircuts, time-to-recovery, and discount rates. Some banks also adjust LGD to account for reductions in exposure from the reference date to the date of default, while others do not.

A number of banks noted the difficulty in performing backtesting LGD estimates due to the typical length of workout periods (in some cases spanning four to five years) and how to treat unresolved workouts. Some banks indicated that focusing on recent periods is particularly problematic as the recovery process has only just begun. For instance, a couple of banks with very low A/E ratios for LGD indicated that these ratios were likely not indicative of the portfolios’ loss performance since the resolved cases

[18] In some cases, a full cycle refers to observations that occurred quite some time ago (eg the 1990s).

included in the numerator include a large volume of cures and resolutions with little or no loss; the more significant losses were likely to occur much later in the workout process.

The two main approaches noted by banks for modelling EAD are the EAD factor approach (estimated EAD relative to credit limits) and the CCF approach (drawn EAD plus an estimated portion of undrawn exposures). The EAD factor approach is the predominant approach applied to retail exposures. Banks' choice of horizon methodologies varies with no clear dominant approach.[19] Banks apply a variety of rules for censoring reference data observations (eg flooring CCF at zero or some other level, applying caps, or removing observations with extreme values).

Finally, banks' validation and backtesting processes are typically performed at the modelled segmentation level defined by the bank as opposed to the regulatory asset classes used in this study.

1.4 Policy implications and recommendations

The nature of the analysis used in this study, combined with material data limitations, does not allow for a definitive measure of how much observed variation in retail and SME RWAs is driven by differing practices or differences in risk. Controlling for risk and other local factors that can influence A/E ratios proved challenging. Nevertheless, discussions with bank participants were informative and provided insights into the more significant sources of practice-based RWA variation in these portfolios.

The second objective of the study is to highlight where there is potential to modify current standards either to reduce practice-based RWA variation or to increase the simplicity and comparability of the IRB capital framework. Policy changes in the following areas may provide opportunities to reduce practice-based variation:

- Better definition around what is meant by "long-run average" with respect to PD estimation and acceptable methodologies and data with which to calibrate to these long-run averages;
- Providing banks with guidance on rating systems used in PD estimates (eg the use of PIT- or TTC-oriented systems);
- Better definition of what is considered a "downturn" with respect to LGD and EAD estimation and methodologies to identify a downturn (specifically what to do when stress period observations are not in the developmental reference data set). Supporting data should include suitably long time series of loss and exposure data and information on risk drivers;
- Exploration of treatments for unresolved cases in modelling recoveries within the LGD estimation process and within backtesting studies by banks where estimates are compared with actual losses;
- Consideration of options for reducing the range of discount rates banks apply to model recovery cash flows; and
- Further clarification around data censoring techniques applied to the estimation of EADs.

Additional recommendations that could either indirectly encourage or lead to harmonisation in practices or improve understanding of differences in RWA variation include:

- Elaborate on expectations for validation. As examples, large and persistent gaps between actual defaults and loss experience and IRB estimates may suggest the need for recalibration; and

19 Approaches varied between (a) a fixed horizon approach, where defaults are assumed to occur 12 months from the observation date; (b) a cohort approach, where obligors are grouped and defaults could occur any time over the forthcoming 12 months (or another defined period); or (c) a multiple horizon (or variable period) approach, where exposure is considered at several different intervals over the horizon period.

expectations for validation include the existence of robust and flexible information systems. Sound practice observed by the Committee in model validation is set out in Chapter 3; and

- Consider requiring Pillar 3 disclosure of backtesting results performed at the level of regulatory asset classes to facilitate comparability across banks.

1.5 Limitations of the exercise

An important limitation to the study lies in the central assumption. Actual outcomes reported by banks may not always be good proxies for risk. This is especially the case for realised losses on defaulted exposures. This is due to the wide variation in loss recognition practices across jurisdictions, which affects both the timing and amount of recorded losses.

Banks generally did not submit long time series. Across regions, the average length of time series varies between four and eight years, being longest for North American banks. This result was similar for different asset classes (Chart 12). The explanations provided for the brevity of the time series included information technology challenges such as organising data from multiple entities or recent mergers; combining data from multiple systems (eg combining data from accounting and risk management systems); and the use of internal segmentations that do not align with the regulatory segments requested in the submission.

Two additional implications stem from the relative shortness of the time series of data collected. First, it is not possible to infer from this short time series whether specific banks have calibrated their IRB parameters appropriately, even though most observed A/E ratios fall below 1. Second, it is likewise not the case that A/E ratios exceeding 1 necessarily infer poor calibration as the framework requires calibration to long-run averages. A longer time series showing whether A/E ratios are persistently above or below 1, on average, would be necessary to make such judgments.[20]

Length of submitted PD time series by region and asset class — Chart 12

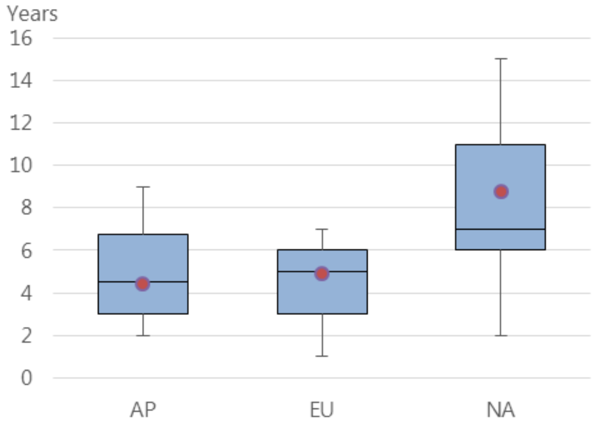

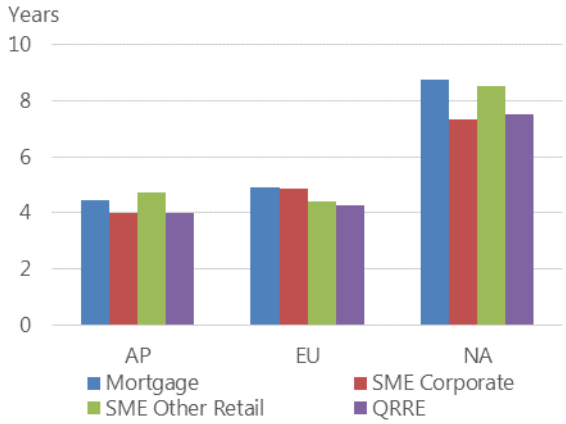

The left-hand chart illustrates the distribution of PD time series submitted by region, with each submitted asset class at each bank comprising one observation. Bars are the maximum and minimum, the box represents the interquartile range and the red dot is the mean.

Both charts are based on the final cleaned data set used for the computation of A/E ratios and the impact on banks' capital ratios. The entry for each asset class for each bank uses the longest series provided either for the model-in-use for that year or for the recent model.

[20] In addition, many banks were unable to provide IRB estimates that would have been produced by the most recent model used to develop IRB estimates (recent model). As a result, what is often represented in the "E" of the A/E ratio is an estimate based on the model in use for prior periods (model-in-use). In other words, banks may have recently recalibrated their models but the result of that recalibration cannot be demonstrated relative to historical periods.

A certain number of observations were omitted from the analyses due to data quality issues (eg identical line items, missing data elements, extreme values, or negative values). Some analyses focused on more recent observations (ie those after 2006) as few banks provided data covering periods prior to 2007. The PD parameter required the fewest omissions followed by LGD and EAD. See Annex 3 for additional details concerning data omissions.

Limitations in the data collected prevented analysis of the alignment of banks' IRB RWAs with risk outcomes over the course of an entire business cycle. There was also very little geographic overlap across banks (ie different banks operating in the same country), as shown in Chart 13, suggesting that banks' retail and SME lending activities are predominantly focused on domestic markets.

Number of countries covered by each bank's submission

Chart 13

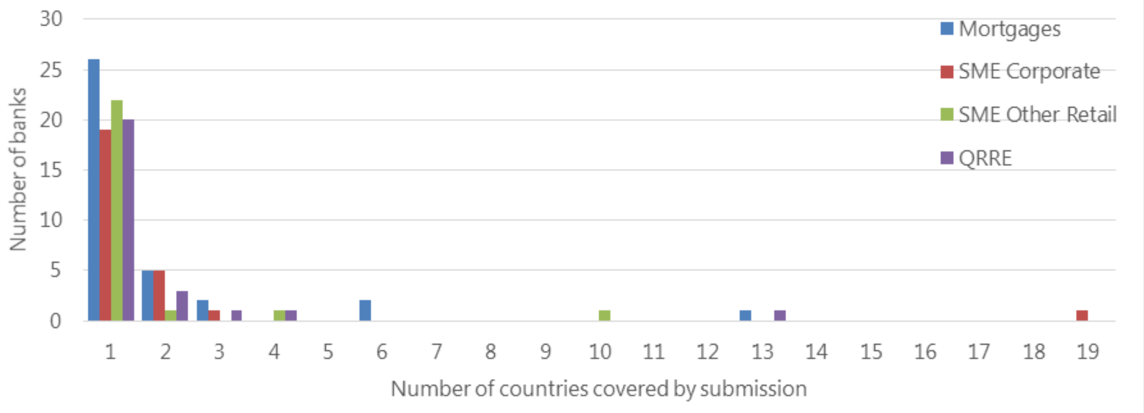

For example, one bank provided data for 19 BCBS countries for the SME Corporate asset classes, while the other banks provided data on SME Corporate for one to three countries only.

The "localness" of retail and SME exposures affects the interpretation of the observed differences in risk parameters and subsequently variation of RWA. That is, structural differences between countries, which might include the characteristics of the market and the local legal environment, will play a significant role in overall RWA variability. For example, the variability of RWA density (RWA divided by EAD) for the mortgages appears to be striking in light of the high level of product standardisation (see Chart 14). RWA densities are more homogeneous when observed at the regional level and appear to be driven largely by LGD variation. Again, LGD variation is likely to be closely linked to the bank's region to include factors such as market prices and characteristics as well as the local legal environment.

In the data submitted for this study, information on risk drivers could not be used to control for risk when comparing RWAs across jurisdictions. For example, preliminary analysis of this data set does not reveal a strong relationship, at the global level, between LTV and mortgage default rates or realised losses. There is evidence of a somewhat stronger relationship between risk drivers and outcomes when focusing at the level of an individual jurisdiction or multiple jurisdictions with similar economies and markets. Further analysis would be needed to draw robust conclusions about the relationship between risk drivers and risk weight variation.

Multivariate regressions of the PD and LGD parameters were used to control for country and year effects, to give an indication of the influence of local factors. After controlling for these effects, there were not statistically significant differences in A/E ratios for a majority of the participating banks. In other words, the differences in A/E ratios for a majority of the banks can be explained by local factors and the business cycle. This analysis does not explain the reasons for the remaining differences in ratios, which could be due to differing modelling practices, bank-specific or other factors.[21] Moreover, this analysis does not explain the reasons for most of the variation in both PD and LGD A/E ratios. It is not possible, based on this analysis, to determine whether or not the variation in A/E ratios is warranted.

Performing IRB Mortgage RWA density, compared with average LGD — Chart 14

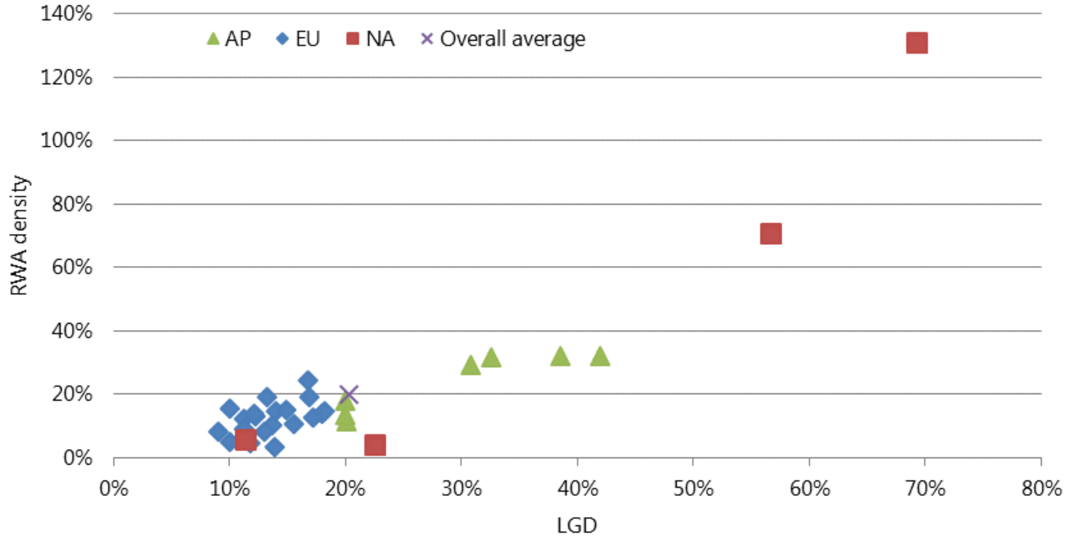

RWA density calculated for each bank as mortgage RWAs divided by mortgage EAD. Non-performing exposures excluded.
Source: CMG data as of year-end 2013.

[21] As examples, modelling choices for banks in a given region may be driven by expectations of the local supervisor, or local conventions may result in different problem loan workout, foreclosure and loss recognition practices.

Chapter 2: Review of EAD/CCF estimates

2.1 Introduction

EAD estimates seek to ensure that capital requirements cover the tendency for drawn exposures to increase in the lead-up to customer default. In Phase 1 of its work on banking book RWA variability, the Committee investigated EAD estimates for corporate general-purpose revolving facilities. This study extends that analysis to other IRB asset classes and other types of undrawn exposure. The objectives of the study were, first, to identify the main drivers of RWA variation and evaluate their effect; and, secondly, to point out areas where there is potential to modify current standards either to reduce practice-based RWA variability or to increase the simplicity and comparability of the IRB capital framework.

The Phase 1 survey, conducted in 2012, identified estimation issues attributable to low-default portfolios that contribute significantly to practice-based variation in CCF estimates for large corporate, bank and sovereign obligors. It also found compelling evidence of practice-based variation in estimates relating to the uneven application of supervisory CCF estimates under the standardised and IRB approaches that can materially influence overall variation in RWA. It also identified a range of other EAD estimation issues that exist to a greater or lesser extent more generally across IRB portfolios.

Reflecting its wider scope, participants in the Phase 2 (2014) survey submitted a range of portfolio-level data from which implied average EAD CCF estimates could be calculated. The ensuing analysis focused on the extent of undrawn exposure within banks' IRB portfolios and the variation in the implied average CCF estimates assigned to that exposure. The survey did not address counterparty risk exposures for derivatives, repo and repo-style transactions, securitisation exposures, settlement risk exposures or netting requirements.

The survey data distinguished "advised exposures", where banks had advised or informed the customer of the limit of the lending commitment or issued OBS item, from "unadvised exposures", where limits have been established within banks' systems but not formally advised to customers. The study found that most banks do not apply CCFs to unadvised exposures – that is they assume the exposure is zero. Consequently, these exposures represented a very small proportion of total banking book exposure. The remaining banking book exposures comprise "advised exposures". Of these IRB exposures, advised undrawn exposures comprise almost 30%. Variation in CCFs that are applied by banks could therefore materially affect total banking book exposure, and thereby have a material effect on overall RWA variation.

Of advised undrawn exposures:

- undrawn lending limits comprise a little over 80%, with issued OBS items (such as trade letters of credit or performance guarantees) accounting for the remainder;
- well over half of undrawn lending limits fall within the corporate asset class, followed by QRRE (mostly credit card exposures) and residential mortgages. These three asset classes comprise 86% of advised undrawn lending limits;
- undrawn lending limits, relative to advised exposures, are largest for the QRRE asset class, followed by Corporate, Retail SME, and other retail (see Chart 15);[22]
- in line with expectations, within the corporate asset class, larger obligors account for the larger share of exposures and tend to have relatively larger undrawn lending limits (see Table 4);
- the corporate asset class accounts for around 75% of issued OBS items; and

[22] As shown in Chart 15, there is some variation in the significance of undrawn lending limits by asset class across regions.

- EAD estimates are assigned to about 40% of issued OBS items (not all OBS items are eligible for EAD estimates).

Undrawn lending limits as a percentage of advised exposure, by asset class

Chart 15

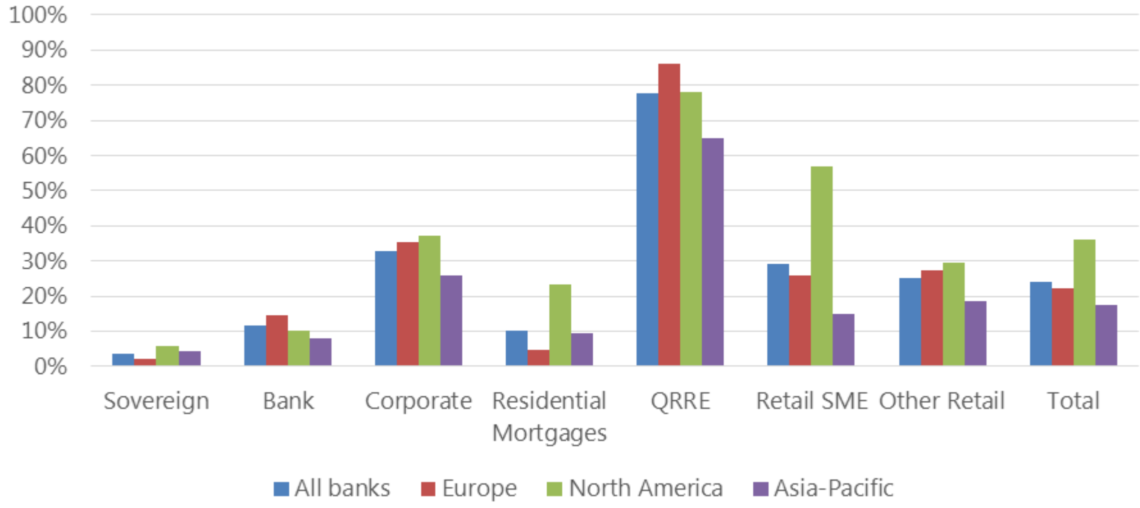

Corporate undrawn lending limits

Table 4

Exposure size	as a % of corporate advised exposure in each size category	as a % of corporate undrawn lending limits in each size category
0–€5 million	27	14
€5–50 million	33	30
> €50 million	41	56
All	33	100

Includes only those facilities where the client has been formally advised of the limit of the lending commitment.

2.2 Structure of the study and survey participation

To assess variation in EAD/CCF estimates, the Committee undertook a survey in 2014. Participants were asked to submit a range of portfolio-level data from which implied average EAD credit conversion factors (CCFs) could be calculated. In addition to the portfolio-level data, survey participants responded to questions about certain EAD consistency issues: treatment of non-revolving facilities intended to be drawn down progressively over time, eg construction and development loans; when commitments are recognised as starting for regulatory capital purposes; treatment of multi-option facilities; and treatment of cash collateral.

The ensuing analysis focused on the extent of undrawn exposure within banks' IRB portfolios as well as the variation in the implied average CCFs that are assigned to each type of exposure. Variability was examined across the following dimensions: (a) among respondent banks; (b) across IRB asset classes; and (c) between banks' own estimates and equivalent supervisory CCFs under the FIRB and standardised approaches.

Thirty-seven banks from 17 countries participated in the EAD survey, including most G-SIB banks. However, complete data were provided by only 27 banks. One (of four) FIRB respondents completed only the qualitative part of the questionnaire. Of the remaining 36 banks, nine banks were excluded from the aggregate data analysis tables:

- one bank reported no undrawn limits, another reported no undrawn limits for its non-retail exposures or any drawn amounts, and a third reported no retail undrawn limits;
- two banks appear not to have followed fully the instructions for completing the templates; while
- four banks (including one FIRB bank) submitted corrupted data templates.

Although excluded from the general aggregate analysis tables, the above banks have been included in some of the analysis set out below based on the comments accompanying, and partial analysis of, submitted data.

Some other data quality issues are discussed within the body of the report where relevant to particular parts of the analysis. Nevertheless, the overall quality of the data submissions, together with comments provided by banks, was sufficiently high to support key findings. Often, the problematic aspects of data submissions appeared to affect only small exposure amounts. In other cases (eg where data were either not provided or not segregated into the requested categories), associated comments enabled qualitative observations.

2.3 Findings of the study

2.3.1 Overview of findings

Among respondent banks, implied average CCFs vary widely in all IRB asset classes. The overarching conclusion from the survey, consistent with the Phase 1 survey, is that this variation is significantly affected by estimation practices. This includes the uneven application of supervisory CCF estimates under the standardised and IRB approaches, which can materially influence overall variation in RWA. The influence of EAD-driven variation in RWAs is likely greatest in banks that focus on commercial lending activities.

The main findings of the survey are as follows:

- A substantial number of banks permitted to use their own EAD estimates apply zero CCFs, without empirical backing, to sometimes substantial proportions of their undrawn exposures.
- There are widely varying interpretations concerning how supervisory CCFs should be applied. These differing interpretations likely extend more generally to FIRB and standardised banks.
- CCF estimates applied by IRB banks generally do not differ between revolving and non-revolving undrawn exposures, despite supervisory expectations to the contrary.
- There is a wide range of practice concerning when to recognise the start date of a commitment that also likely extends to FIRB and standardised banks.
- There are a number of additional estimation issues that can contribute materially to differences in EAD estimates among IRB banks. These issues include potential estimation biases resulting from including restructured products, known problem obligors, or facilities that are close to fully drawn in the reference data without adjustment; the use by different banks, and within the same bank, of different estimation approaches (ie the fixed-horizon, multiple horizon, or cohort approaches); various applications by banks of caps or floors on the CCF observations used to develop and validate EAD models; and the use by some banks of estimators other than the mean.

2.3.2 Variation amongst non-zero CCFs

Where non-zero CCF estimates are assigned, there is considerable variation in average CCFs among respondent banks in all asset classes (see Table 5). Table 6 summarises this variability by showing the interquartile ranges of CCF estimates by asset class (the column labelled share of IRB undrawn lending limits gives a sense of the materiality of CCF estimates to RWA for a given asset class). Residential mortgages exhibit the highest variability, followed by retail SME. The "other retail" and non-retail asset classes are clustered together and QRRE shows the least variation based on the interquartile range. The reasonableness of this degree of variability is considered in the following paragraphs.

Implied average CCFs – Undrawn lending limits: CCF>0

Table 5

	(a) AIRB/Retail IRB banks – All IRB undrawn lending limits						(b) AIRB banks – "AIRB" undrawn lending limits					
			Percentile						Percentile			
Asset class	Av	Min	25th	Median	75th	Max	Av	Min	25th	Median	75th	Max
Sovereign	60	15	45	59	75	100	59	13	45	58	74	100
Bank	66	7	47	72	83	100	64	7	45	70	81	100
Corporate	62	29	44	57	75	100	59	27	44	55	70	100
Mortgages	72	2	50	75	100	102						
QRRE	53	0	39	55	68	100						
Retail SME	75	19	46	78	94	191						
Other retail	69	22	55	70	88	111						

Includes only those facilities where the client has been formally advised of the limit of the lending commitment.

CCF estimates of 100% pertain to banks within a single jurisdiction, where supervisors are currently reluctant to consider lower estimates.

Variation in implied average CCFs (Undrawn lending limits: CCF>0)

AIRB/Retail IRB Banks – All IRB undrawn lending limits where the client has been formally advised of the limit of the lending commitment.

Table 6

	Interquartile range (percentage points)	Share of IRB undrawn lending limits
Sovereign	30	2
Bank	35	4
Corporate	30	56
Residential mortgages	50	13
QRRE	28	17
Retail SME	49	2
Other retail	33	6

The corporate implied average CCFs in Table 5 are consistent with the Phase 1 EAD survey that covered only corporate general purpose revolving lending commitments. Excluding the respondent banks with 100% CCFs, the implied average CCF for corporate undrawn lending limits with non-zero CCFs is around 55%. This compares with CCFs averaging around 50% for general purpose revolving facilities (based on the Phase 1 survey's direct review of relevant EAD models).

For banks that participated in both surveys, Chart 16 shows 2014 implied average CCFs and 2012 CCFs for general purpose revolving commitments. Again, there is close correspondence between the two surveys (with the larger gaps likely reflecting model changes that have occurred within the intervening period).

Implied average CCFs for general purpose commitments, as estimated in 2012 and 2014 EAD surveys

Chart 16

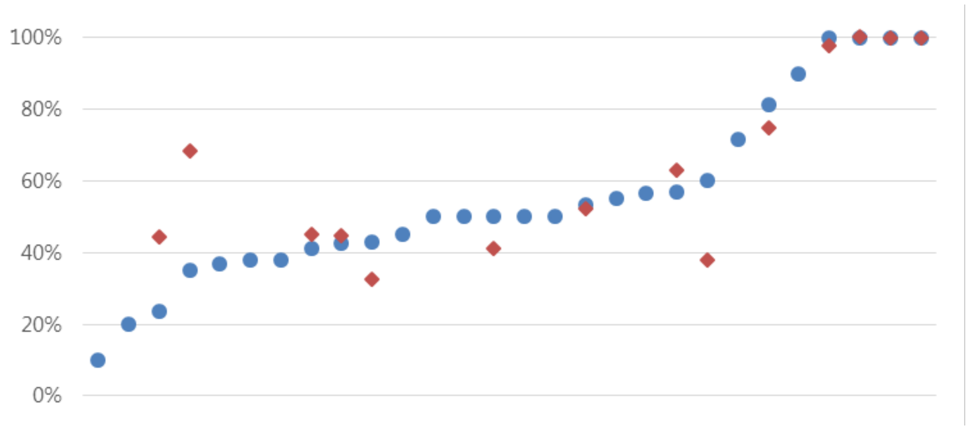

The blue dots represent the model CCFs reported in the Committee's 2012 survey of EAD practices. For those banks that participated in both surveys, the red dots represent the implied average CCF calculated based on data submitted in 2014.

Both surveys show considerable variation in the CCFs that are applied to corporate undrawn limits by IRB banks. However, an essential difference between the two sets of data is that the 2012 figures are actual CCFs that would be applied by different banks to the same or similar customers. In other words, the close correspondence of the 2012 CCFs with 2014 averages generally reflects the simple structure of most banks' CCF models rather than picking up relevant compositional differences in banks' portfolios.

An assessment of whether the variation observed in Chart 16 might nonetheless indirectly reflect mostly risk-based drivers, or is significantly affected by practice-based drivers, requires additional information not gathered by the current survey. The 2012 survey, however, sought such additional information and concluded there is a high probability that IRB EAD estimates are significantly affected by banks' estimation practices.

One of the main issues identified was the common practice of using SME and mid-market data to overcome low-default issues associated with large corporate obligors. Estimates based on such data are typically used also for sovereign and bank obligors. One unknown, however, is the extent to which existing variation in CCFs applied to large corporate, banks and sovereigns might be reduced if other estimation issues (discussed below) were better dealt with. In such circumstances, smaller corporates might ultimately look to be better proxies than they currently appear.

Other material estimation issues (with the first three considered most important) were as follows:

- Commonly occurring changes in loan characteristics or "product profile transformations" (eg a revolving loan that has been converted into a term loan or vice versa) between reference and default dates have high potential to introduce substantial arbitrariness and downward bias in banks' EAD estimates;
- Issues that arise when a transaction's realised CCF is driven by actions of the bank to constrain or reduce its exposure to the obligor. For example, when risk management is specially focused on

deteriorated borrowers and is in a position to successfully manage down limits (and outstandings) prior to default, lower realised CCFs and associated estimates might be the result. If such estimates are uniformly applied to all borrowers (including performing customers with no known issues), as is commonly the case, then the CCFs for lower-risk borrowers could be biased downwards. As performing customers comprise the vast bulk of banks' portfolios, these so-called "race-to-default" issues have potential to bias banks' overall EAD estimates;

- Developing estimates using observations from the region of instability associated with small undrawn limits can introduce substantial arbitrariness and possible downward bias, in banks' EAD estimates (applies to banks that use the undrawn limit factor approach for estimates);
- The use by different banks, and within the same bank, of fixed-horizon, multiple horizon or cohort estimation approaches.
- Various applications by banks of caps or floors on the CCF observations used to develop and validate EAD models.
- Use of estimators other than the mean.

As low-default issues tend to be less prevalent for retail portfolios, the issue of applying estimates to exposures that are based on reference data with possibly markedly different characteristics is likewise seen to be less prevalent than for large corporate, bank and sovereign exposures. Nevertheless, general EAD estimation approaches are not different for retail exposures; consequently, retail EAD estimates are similarly prone to the other practice-based differences set out in the section above. For example, as debt consolidation and other restructurings involving problematic retail exposures is common, issues relating to product profile transformation exist in retail as they do in non-retail portfolios. That said, residential mortgages are unlikely to be converted into other sorts of exposures, so that the issue of product profile transformation may be less important in the residential mortgage asset class than in other parts of banks' portfolios. Similarly, the revolving and unsecured nature of QRRE means that banks can be expected to respond more quickly and consistently to payment delinquencies. In so doing, banks are facilitated by the transactions-based nature of QRRE exposures, which typically provide a regular flow of information on their performance status. Not unexpectedly, therefore, the QRRE-implied average CCFs shown in Table 5 exhibit the least variability (although there are a few outliers).

2.3.3 Use of zero CCFs

The most significant additional finding of the study was the use of zero CCFs. The 2014 survey showed that a substantial number of banks permitted to use their own EAD estimates apply zero CCFs to sometimes substantial proportions of their undrawn exposures (see Chart 17). This result stems from a wide spectrum of views concerning what is permissible in the area of unconditionally cancellable commitments. These different interpretations are a driving factor behind material non-risk related differences in RWAs among IRB banks.

Undrawn lending limits with CCF = 0 Chart 17

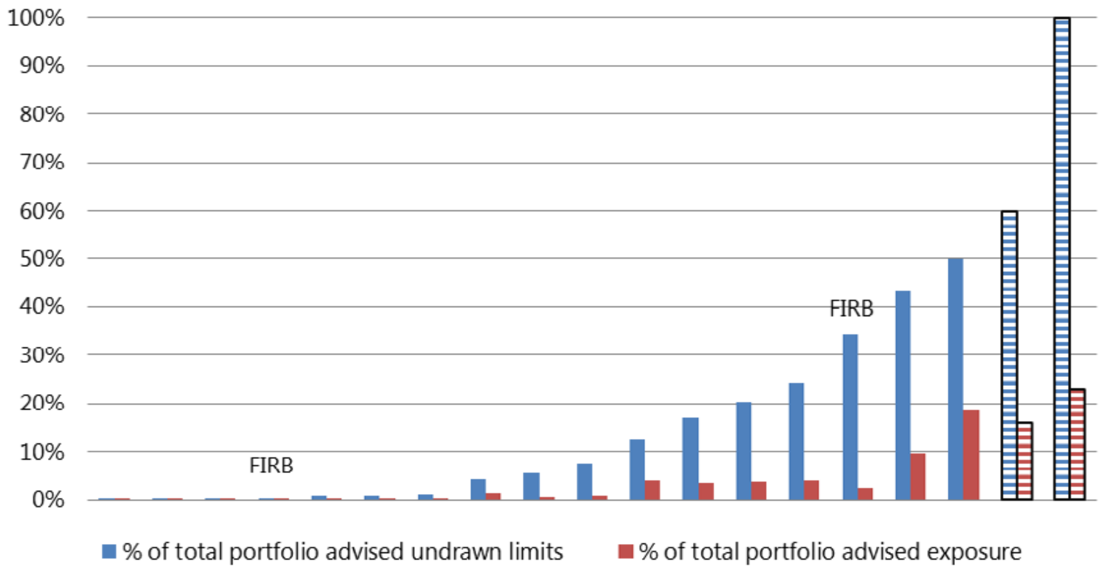

Each pair of bars relates to a participating bank. Two banks using the FIRB approaches are labelled as such in the chart.

Includes only those facilities where the client has been formally advised of the limit of the lending commitment.

The banded bars represent estimates for two banks that did not provide sufficient information on undrawn limits. For these banks, zero CCFs were assumed to apply to the unreported undrawn limits and the ratio of undrawn limits to total issued OBS and drawn exposure in each asset class was assumed to be equal to the reported average of other survey respondents.

Specifically, over a third of respondent banks that are permitted to use their own EAD estimates apply zero CCFs, estimates that are not based on empirical evidence, to significant and sometimes very substantial proportions of their undrawn lending limits. This outcome seems contrary to IRB requirements that, among other things, banks' EAD estimates must take into account all available information, be grounded in historical experience and empirical evidence, and not be based purely on subjective or judgmental considerations. It is likely that there are also differences in interpretations as to what constitutes an unconditionally cancellable commitment for regulatory capital purposes for banks that use the FIRB and standardised approaches, which currently permit zero CCFs for such exposures. For example, there is evidence that some banks do not assign CCFs (zero or otherwise) to unconditionally cancellable facilities (presumably on the basis that the facilities are not commitments), while some banks assign a zero CCF to commitments that are only conditionally cancellable.

The survey also found little, if any, empirical support for widespread application of zero CCFs or significant differences between CCFs for unconditionally cancellable and other commitments. Chart 18, for example, shows that where banks do not assign zero CCFs and distinguish between unconditionally cancellable and other facilities, CCFs usually do not differ materially between the two types of facility and are well above zero in both cases.

CCFs for unconditionally cancellable and other facilities

Chart 18

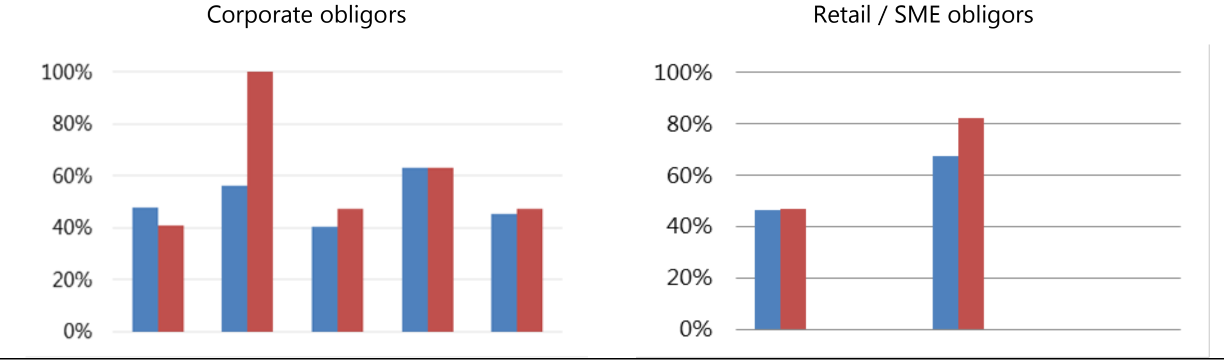

The red (blue) bars show the non-zero CCFs applied by individual banks to unconditionally cancellable (other) commitments, for both corporate and Retail / SME obligors.

2.3.4 Revolving and non-revolving undrawn exposures

From the 2014 survey, an additional key finding is that CCF estimates applied by IRB banks generally do not differ materially between revolving and non-revolving undrawn exposures, despite supervisory expectations to the contrary (see Chart 19). Under revolving facilities, customers may generally draw down (ie borrow), repay drawn amounts and redraw previously repaid amounts on a fluctuating basis within a specified limit. Overdraft accounts are a common type of revolving facility. On the other hand, with non-revolving facilities, customers usually draw down and repay according to an arranged schedule with repayments generally not available to be redrawn. While the line of demarcation might not always be clear-cut, the latter characteristic is generally the key distinguishing feature. Undrawn non-revolving facilities include lending commitments, such as term loans, that have not yet been drawn down and partially drawn commitments that are meant to be drawn down progressively over time. This latter type of commitment is common for construction and development, project finance and agricultural loans, among others.

Implied average CCFs

Chart 19

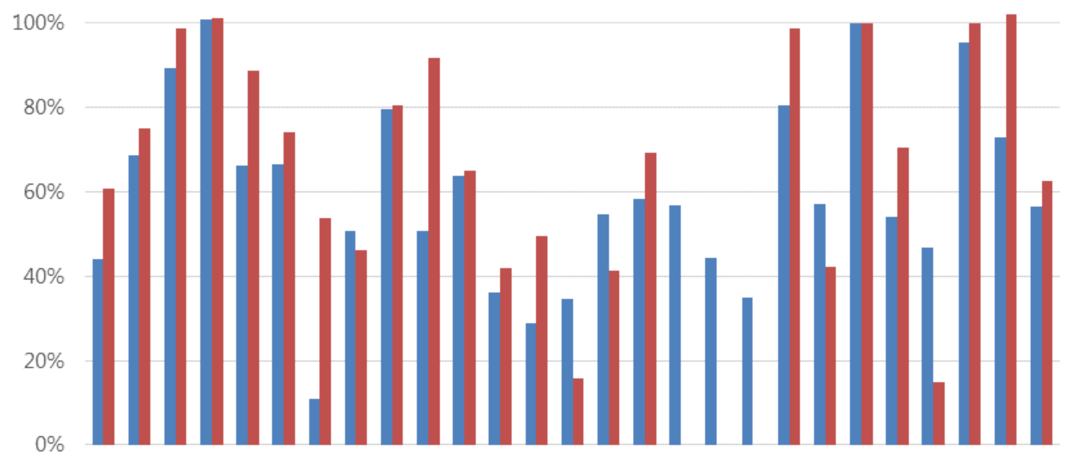

The blue bars represent revolving exposures and the red bars represent non-revolving exposures. Each pair of bars represents a participating bank. Only one bar is shown for the small number of banks that did not segregate non-revolving exposures.

The revolving/non-revolving distinction is important, as CCFs for these two types of exposure would generally be expected to be very different. Indeed, one reading of the Basel framework is that it generally sets the CCF for undrawn non-revolving commitments to 100%.[23] This observation extends to responses to the stylised examples contained in the qualitative part of the EAD survey that relate to CCFs applied to construction and development loans. The finding also extends to CCFs applied by IRB banks to non-revolving undrawn limits under the FIRB approach (and likely also extends to CCFs for non-revolving undrawn limits under the standardised approach).

2.3.5 Additional findings

The sections above describe the main findings of the EAD survey. Additional findings are summarised in brief below:

- The survey finds material variation in CCFs for trade-related advised facilities and undrawn lending limits linked to non-trade performance-related contingent items. Differing practices such as those mentioned above (broad application of zero CCFs, estimation issues associated with product profile transformation) could contribute to this variation. Perhaps because these types of exposure typically have few defaults, some IRB banks' CCF estimates for undrawn trade-related exposures appear to be the same as, or anchored to, equivalent supervisory estimates. However, different "equivalent" supervisory estimates often appear to be referenced by different banks. In a few cases, this might reflect a conservative choice but more generally it appears to reflect interpretation differences contributing to unwarranted RWA differences. Such interpretation differences almost certainly extend to banks using the FIRB and standardised approaches;
- It is possible that some banks are assigning own CCF estimates to ineligible issued OBS items, though this would need to be confirmed;
- A few banks' corporate EAD estimates are based essentially on the same broad benchmark studies as the 75% FIRB CCF but fall well below that benchmark;
- Survey responses showed that few banks go beyond advised exposures when developing CCFs to account for the potential for drawn exposure to increase in the lead-up to default;
- Survey responses showed that few banks assign CCFs to non-advised lines and other products not normally intended to result in credit exposures (eg deposit accounts with overdraft protection); and
- The survey showed a wide range of practice concerning when to recognise the start date of a commitment that also likely extends to FIRB and standardised banks.

[23] Under Basel II paragraph 316, for off-balance-sheet products, banks may use their own CCF estimates for exposures that are not subject to a 100% CCF under the FIRB approach. Note that this set of exposures is the same as under the standardised approach. Under the standardised approach, commitments with certain drawdown should be assigned a 100% CCF as opposed to: unconditionally cancellable commitments, which may be assigned a zero CCF; and "other" commitments, which may receive: a 20% CCF if the original maturity is under one year; a 50% CCF if the original maturity is one year or more; or a 75% CCF regardless of maturity under the FIRB approach. Commitments with certain drawdown would generally seem to include non-revolving facilities that are not unconditionally cancellable and either have not yet been drawn down or have not yet been fully drawn down because they are meant to be drawn down progressively over time. The CCFs for "other" commitments are generally meant for facilities with fluctuating balances, ie revolving facilities and other facilities where there is no prior expectation that the facility will become fully drawn as a matter of course.

2.4 Policy implications

The findings of this study suggest a number of possible policy solutions to reduce practice-based variation in CCF estimates. These solutions can, for the most part, be mapped to the following five principles:

1. Estimates of EAD/CCF should be empirically based.
2. EAD/CCF estimates should be based on reference data that reflect the customer, product and bank management practice characteristics of the exposures to which the estimates are to be applied.
3. Banks should ensure that their EAD/CCF estimates are effectively quarantined from the potential estimation issues under the undrawn limit factor (ULF) approach[24] when observations have small undrawn limits (referred to as the region of instability).
4. EAD reference data should not be capped at the principal amount outstanding or facility limits. Interest payments due and limit excesses should be included in EAD/CCF reference data.
5. Where a bank bases its estimates on an alternative measure (eg a median or other percentile of a distribution or using only "downturn" data), it should confirm that the basic downturn requirement of the framework is met, ie the bank's estimates do represent (conservative) estimates of the long-run default-weighted average EAD for similar facilities.

The following additional recommendations are put forward with the objective of harmonising practices in the area of EAD estimation:

- Consider the application of supervisory CCFs unless the banks possess strong evidence to move away from those estimates;
- Consider more specificity as to what constitutes an unconditionally cancellable commitment for purposes of applying zero CCFs under the standardised and FIRB approaches;
- Consider the benefits and possible undesirable consequences of requiring a single approach among the currently used approaches – ie the fixed-horizon, multiple horizon or cohort approach; and
- Consider adopting a unified rule for when banks should recognise the start date of a commitment. A variant of this recommendation would be to consider a unified rule for wholesale exposures and a separate unified rule for retail exposures.

2.5 Limitations of the exercise

This study finds wide variations in CCFs estimated by banks. An important limitation to this study is that data analysis cannot identify to what extent practice-based variations account for entire variations. However, the survey and discussions with selected banks reveal that there are wide ranges of practice in estimation of CCF, which imply that a substantial portion of the variations could be practice-based.

As noted in Section 2.2, complete data was only provided by 27 of the 36 banks that participated in the EAD survey. This limited the aggregate analysis that could be completed. Nevertheless, the overall quality of the data submissions, together with comments provided by banks, was sufficiently high to support key findings.

[24] The ULF is a specific type of CCF, expressed as a percentage of the undrawn limit that remains available to the obligor to draw down under the terms and conditions of a facility. ULFs are the most common type of CCF applied to undrawn credit limits.

Chapter 3: Sound practices in the independent validation of IRB models within banks

3.1 Introduction

Alongside the data collections described above, the Committee surveyed supervisors and collected sound practices observed in banks' independent IRB model validation processes. The practices covered in this paper are grouped into three categories: governance; methodology and scope; and interaction across different phases of model development and implementation. Although this survey was focused on IRB models, and subsequent supervisory responses addressed IRB model validation in particular, some of the practices identified could also be generalised and applied to Pillar 1 models beyond the IRB approaches for credit risk.

As noted above, large and persistent gaps between actual defaults and loss experience and IRB estimates may suggest the need for model recalibration or weaknesses in model validation. The application of sound practice in the area of model validation can reduce some of the variation in practice that otherwise might exist. Harmonisation in the area of model validation could ultimately lead to reductions in practice-based RWA variation. In particular, the sound practices described can help support a model validation framework that encourages and promotes effective challenge, which is an important focus for supervisors in reviewing banks' independent model validation functions.

Banks have wide discretion in the design of an independent model validation function and the methods it uses to ensure robust and accurate risk estimates. Accordingly, it should be expected that the model validation practices differ from bank to bank. Jurisdictional differences in the implementation of the IRB framework may also lead to variation in model validation practices. This report does not set forth new standards or guidelines for model validation. The Committee's formal requirements on model validation in the IRB approach are set out in Section III.H of Basel II.[25] Rather, the practices outlined in this paper should be viewed as useful referential benchmarks for both banks and supervisors.

3.2 Governance of the validation process

Paragraph 500 of the Basel II framework states that banks must have in place a robust system to validate the accuracy and consistency of rating systems, processes and the estimation of all relevant risk components. An important component of any robust system is the governance arrangements, which should delineate clear responsibilities for front-line model developers, model validators and internal audit. Often referred to as the "three lines of defence", these functions each play an important role in identifying and remediating model risk (ie incorrect estimation of IRB risk parameters). This paper focuses on sound practices pertaining to the model validation function (ie the second line of defence in that terminology).

Under the oversight of the Board of Directors, a bank's senior management is responsible for putting in place effective model validation governance processes.

Sound practice entails ensuring that the independent validation function is independent of the model developers and the commercial (ie business or risk-taking) functions. It should have sufficient stature in the organisational hierarchy to challenge effectively the modeller's work. Also, it is important that the remuneration of the staff in the independent validation function does not compromise their independence.

A key aspect of a sound independent validation function is that validation staff have adequate skills. The requisite skills would include knowledge of the business where the models are used, as well as

[25] See paragraphs 441, 442 and 500-5.

up-to-date knowledge of the mathematics and quantitative techniques used in the underlying models. Quantitative resources should be commensurate with the volume and complexity of tasks to be performed. Validation staff should avail themselves of new research, such as new tests or distributions.

Banks exhibiting sound practice avoid cross-validations, whereby two separate departments validate their respective models alternately. This practice maintains independence in the validation. Where validation activities have been delegated to external parties, the independent validation function would retain full and ultimate responsibility for the validation activities and results. This does not preclude any review of the independent validation function itself, which is subject to internal audit.

A sound independent validation function reports its results to a model acceptance or oversight committee that includes senior officers or staff with the power and ability to discuss the issues at stake and prepare the final decision for approval. The discussions of this committee are recorded. Its minutes document the challenges made to the model and, where applicable, the reasons why the independent validation function's concerns have not been addressed. Any conditions of model approval are also clearly stated in the minutes. If the model has significant weaknesses, then mitigation procedures are documented and applied (eg conservative adjustments). The Board of Directors is ultimately responsible for the approval of the models that will be used in the bank.

Model owners seek initial model approval by supervisors only after the model has been internally validated and has appropriately followed the bank's own governance process.

In some banks, the independent validation function may intervene while models are being developed. Sound practice exists where such early intervention during the development of models does not put the independence of the independent validation function into question. For instance, the independent validation function should not make modelling choices.

3.3 Methodology and scope of the validation function

Banks with sound independent validation functions generally establish a model governance and independent validation framework. This framework sets out policies and guidelines describing validation tasks, standards and methodologies and helps ensure that validation tasks are comprehensive and consistent. An effective independent validation framework includes the following core elements:

- an evaluation of the development guidelines for models;
- standards for the quality of documentation;
- an evaluation of model scope, together with a comprehensive risk analysis. This risk assessment includes input from the business, front-office or risk-taking functions, among others;
- an evaluation of the quality and relevance of input data, whether internal or external, and including the use of crisis data, analysis of the representativeness of the data, outliers and areas where data is missing;
- the governance arrangements for the data collection process and controls around it;
- an evaluation of model design. This would typically include, among other things, considerations of conceptual soundness, an analysis of risk drivers and an explanation of key model assumptions and limitations;
- sensitivity analysis of key model assumptions;

- performance analysis and backtesting of risk differentiation, the calibration and outcomes compared to expected values generated by models.[26] This analysis should take place at the aggregate model level as well as at more granular grade or segment levels;
- benchmarking to check the plausibility of model outputs;
- an evaluation of model risk (ie incorrect estimation of IRB risk parameters), together with an evaluation of the appropriateness of margins of conservatism to cope with model and data imperfections;
- an evaluation of the implementation of the model, which could include, for example, ensuring that the IT program implementing the model is consistent with the model documentation and reviewing whether the IT infrastructure is fit for purpose;
- clear criteria for determining follow-up actions, such as re-development or re-calibration, and a process for tracking recommendations or requested remediation activities;
- an evaluation of model use, such as whether there are limitations on input data, how overrides are documented, how model users are trained and feedback received from model users;
- an evaluation of the "use test". The Basel II framework recognises that the use of the model (or of its risk factors) in processes other than regulatory capital is an important indicator of management confidence in the model; and
- ensuring the adequacy of internal models with respect to regulatory requirements.

It is a sound practice for the scope of the work of the independent validation function to be defined clearly. This is supported by a bank having a policy that clearly defines a "model" and what constitutes model risks, a model inventory and an independent validation function that verifies the existence and quality of these documents. The model inventory should include a comprehensive list of models used by the bank, their scope (the portfolio coverage of the model), their materiality, a brief description of modelling methodology and approval conditions. It should not only include the main model (eg a loss-given-default, or LGD, model), but also any submodels that contribute to the main model (such as cure rate, realised LGD or downturn LGD). An independent validation function may also look for evidence that any uncertainty associated with model outputs is evaluated, if possible quantitatively (or, if not, at least qualitatively by expert judgment).

Sound practice entails independent validation at every tier within a banking group that uses the internal model (ie both at group and legal entity level), such that aspects of the local applicability of groups' models are also subject to independent validation. In this context, special attention is devoted to specific segments or activities. Independent validation encompasses not only the main models but also the submodels. It also covers structural choices made by the bank, including the choice of vendor models (developed outside the bank) and the design of databases (eg loss databases).

A key element of a sound independent validation function is the creation of a validation plan, which defines the validation activities to be performed. The function follows its plan and modifies it in response to findings to make sure that gaps are remedied. More material models are subject to greater scrutiny. The independent validation function follows standards for documenting the validation work.

Sound practice exists where the independent validation function performs its own tests (at least on representative samples) on all material issues, including model performance tests, the quality of databases used, data cleaning and quality assurance of computer code. These tests also cover tests already performed by the model developers, to check their reliability. Validation policies should contain (at least

[26] Paragraph 501 of the Basel II framework requires banks to compare regularly realised default rates with estimated PDs for each grade and to be able to demonstrate that the realised default rates are within the expected range for that grade. Banks using the advanced IRB approach must complete such an analysis for their estimates of LGD and EAD.

indicative) standards defining when a model or an estimate is not appropriate, and when backtesting results and model performance issues lead to a root-cause analysis, rather than recalibration.

It is important that the analysis and assessment of the independent validation function leads, in written form, to an explicit opinion about the model's adequacy (approval, conditional approval or rejection of the model), including an executive summary with a discussion of its strengths and weaknesses. The validation report contains an assessment of model risk, for the model on its own and also in the light of the model's materiality. In response to the weaknesses identified by the independent validation function, model developers formulate action plans to remediate weaknesses. The independent validation function, possibly aided by an issues register, periodically follows up on remediation efforts and escalates issues that are not being promptly addressed.

3.4 Role of the validation function across different phases of model development and implementation

The role of the independent validation function varies at different points of the model development process.

A sound independent validation function evaluates all new models, including vendor models, before they are approved.

Banks exhibiting sound practice in independent model validation ensure that there is regular follow-up after a model has been approved. Many of the tests and principles used at first approval are included in the regular follow-up, including the use test. This leads to the independent validation function expressing an opinion on the results of the (at least) annual model review performed by the model owners. Models for which uncertainty is higher are subject to higher scrutiny. The opinion of the independent model validation function includes, but is not limited to:

- confirmation that the model is adequate, especially where it will be applied with a different scope than that originally planned (eg to new products, clients, countries, market conditions or IT systems);
- an evaluation of model overrides, including changes suggested by the users of the model; and
- any variation in its discriminatory power over time.

When models change, it is important that banks record these changes and the reasons for them in a model log. Independent validation functions review both the policy for making model changes and its application. Sound practice exists where significant model changes are subject to particular scrutiny from the independent validation function. Less significant model changes are also analysed with the aim of providing an overview of the model's service life. The independent validation function keeps track of the cumulative impact of less significant changes.

Annex 1: Portfolio risk weight distribution based on bank data submissions

This table shows the portfolio distributions of risk weights for retail and SME portfolios calculated from the data submitted in the retail/SME study described in Chapter 1. It is different from Table 1 because that table uses year-end 2013 data from the BCBS Capital Monitoring Group, which covers more IRB asset classes than those which were the focus of this study.

Risk weights (RWA/EAD) for retail and SME portfolios (percent, 2012–13) Table A1.1

Portfolio	Banks	Mean	Median	Min	25th Percentile	75th Percentile	Max	Range
Mortgages	35	20.7	13.7	3.5	9.7	19.0	131.1	127.6
SME Corporate	25	53.1	52.8	26.7	45.1	60.6	83.4	56.7
SME Other Retail	25	47.3	43.5	4.1	32.2	56.2	122.3	118.2
QRRE	22	36.5	31.9	11.9	26.5	43.1	92.4	80.5

Annex 2: Description of data for the retail/SME benchmarking analysis

The study covered SME corporate exposures as well as major subclasses of retail: retail mortgages, QRRE and retail SME exposures, and collected PDs, LGDs and EADs used for the IRB RWA computations. The study did not consider maturity or the size adjustment to RWAs relevant to SME corporate.

The collection was at a segment level, with segments defined as a distinct combination of year, country of obligor, (sub)category of asset class, risk band and the model used – either the model used for the year in question or the model used at the end of the study term (ie by 2013). As an example, a bank that could provide data for 2002 through 2013 would provide a segment for 2002 exposures, HELOC mortgages, third highest quality risk band, and model-in-use. Data associated with different attributes (perhaps a different year of exposure, or recent model) would be assigned to a different segment.

The three rows below illustrate the definition, showing three distinct segments.

2013	CA	SME Corporate	Band 4	Model in use
2012	CA	Mortgages - Revolving- not guaranteed	Band 11	Recent Model
2013	US	Qualified Revolving Exposures	Band 15	Recent Model

Although the segmentation scheme was the same for the PD, LGD and EAD collections, the data collected for each segment depended on the parameter, as shown in Table A2.1.[27]

Information collected by IRB risk parameter — Table A2.1

IRB Risk Parameter	Information collected
PD	• Borrower PD estimate before guarantee • PD estimate used to calculated RWA • Number of defaults • Volume of default • Average LGD used to calculate used to calculate RWA • RWA calculated for the segment • Expected loss for the segment
LGD	• LGD estimate before guarantee • LGD estimate used to calculate RWA • Number of defaults • Volume of defaults • Average LGD used to calculate RWA • Total EAD • RWA calculated for the segment • Expected loss
EAD	• CCF for undrawn portion of facility • EAD estimates for the undrawn • EAD outcome for the undrawn • Average PD for the undrawn • Average LGD for the undrawn • EAD weighted average PD for the undrawn • EAD weighted average LGD for the undrawn • Number of defaults undrawn • Amount defaulted of undrawn

27 For example, the number of performing obligors is relevant for PD studies, but not for LGD and EAD, which look at transactions that have defaulted.

For PD and LGD, additional data were collected for certain risk drivers, including LTV and debt to income (DTI) ratios for retail mortgages in the PD collection and LTV for retail mortgages in the LGD collection.

Along with the PD, LGD and EAD collections just described, the study also collected high-level data that maintained distinctions by asset subclass and year but did not distinguish model, risk band or country, referred to as "portfolio data". Generally, this collection could not be reconciled to the detailed collections, because the high-level data include all countries of operation, whereas banks submitted detailed data only for select countries, and not necessarily for all IRB models. However, the high-level data were used to give an idea of the coverage of the study and for robustness checks.

Banks also provided information about their modelling in a "qualitative" questionnaire. This covered several topics (Table A2.2).

Information collected in the qualitative questionnaire

Table A2.2

Dimension	Information collected (limited to core dimensions)
General questions about modelling methods and the submissions	• The use of long-run loss rates to infer PD or LGD per para 465 • The alignment of the years with calendar years • Coverage of submitted data • Rating philosophy – point in time or through the cycle
Definition of default	• Triggers for default • Materiality thresholds • Whether default is declared for a transaction or the obligor • Treatment of cures
PD	• Details of the time series used • Whether a downturn is included, and adjustments made to arrive at a long run average • Treatment of obligors that default more than once • Inclusion accounts that are inactive or exit during a year
LGD	• Policy about the inclusion of different years of data • Weighting of data from different years in the time series • Inclusion of a downturn in data • Choice of discount rate • Time to resolution and treatment of incomplete workouts
Risk drivers	• What drivers does the bank use and which are most important? • How levels of major drivers are distributed among obligors or transactions
Other	• Some questions about the development of EAD estimates • Some questions about margins of conservatism, aggregation and guarantees

Annex 3: Retail/SME data cleaning and aggregation

The raw data were edited and cleaned as explained below, resulting in the final observations used for the study as displayed in Tables A3.1–3.

PD analysis

- Step 0: Raw data
  - Use of most recent data (ie submission 1, 2 or 3)
  - Amounts in local currency are all converted to EUR using historical foreign exchange rates
- Step 1: Data cleaning
  - Drop "true" duplicates (ie lines which are identical)
  - Drop observations without PD estimates (Rule: use "PD Estimate used to calculate RWA"; otherwise, if available, use "Borrower PD Estimate (before guarantee")
  - Drop observations without actuals (either number-weighted default rates (= number of defaults/ performing obligors) or EAD-weighted default rates (= volume of default / EAD))
  - Drop observations with PD estimates=100%
  - Drop if number of defaults=0, but volume of defaults >0
  - Drop observations with 0 performing obligors
  - Drop observations with negative EAD
  - Drop observations from countries that are outside the BCBS members
- Step 2: Remove overlap resulting from data collection (mainly for mortgages, as data were collected on an aggregate level "Mortgages (All)" and sublevels
  - Rule for the removal of overlapping data: use data at the lowest level of aggregation, if available, otherwise use aggregate data
  - Note that we did define an aggregate category called "Mortgages (All)" and five subcategories – "Mortgages Senior Lien Amortising", "Mortgages Junior Lien Amortising", "Mortgages – Revolving", "Mortgages – Other", "Mortgages unclassified"; for each subcategory, we suggested a split into "guaranteed" and "non-guaranteed"
  - Data for "other retail" received from one bank (which also submitted "SME other retail") dropped
- Step 3: Final analysis
  - Use of the IRB scaling factor for all banks (including those that submitted RWAs without the scaling factor of 1.06)
  - Use of data based on model in use and/or recent model (if model in use not available or for banks that changed their models significantly)
  - Focus on observations after 2006

Methods for aggregation of PD data

- Data for A and E are based on aggregations of data submitted at the risk band level (ie the greatest level of granularity). A/E ratios are computed for each combination of bank, asset class, year and model type; the overall average per bank and asset class is the simple average over the years, using the model-in-use figures if available, otherwise the A/E ratios from the recent model are used
- Method 1: Number of obligors weighted (Default method used for this study)
  A = (default rate * number of obligors) / (number of obligors)
  = (number of defaults) / (number of obligors)
  where default rate = number of defaults / number of obligors
  E = (PD * number of obligors) / (number of obligors)
- Method 2: EAD weighted (for comparison)
  A = (default rate * EAD) / (EAD) ,
  E = (PD * EAD) / (EAD)
- Method 3: Default volume and EAD weighted (for robustness checks only)
  A = (default volume) / (EAD)
  note that default volume (default volume / EAD) * EAD , which means that the default rate is measured on a volume basis instead of number basis.
  E = (PD * EAD) / (EAD)

Table A3.1		Step 1	Step 2	Step 3
Bank	PD (raw)	PD (after cleaning)	PD (overlap removed)	PD (if data prior to 2007 were to be removed)
Total	31,748	29,231	20,948	18,605
Portion	100%	92%	66%	59%

LGD analysis

- Step 0: Raw data
  - Use of most recent data (ie submission 1, 2 or 3)[28]
  - Amounts in local currency are all converted to EUR using historical foreign exchange rates
  - For banks with missing data at granular level: take portfolio level data (bank 18 and 26)
- Step 1: Data cleaning
  - Drop "true" duplicates (ie lines which are identical)
  - Drop observations without LGD estimates

[28] For one bank, a subsequent submission amended previous data, so both submissions were kept.

- Drop observations with LGD estimates > 120 %
- Drop observations without actual LGDs
- Drop observations for which LGD estimates = LGD actuals
- Drop observations with number of defaults zero
- Drop observations with negative or missing EAD
- Drop observations from countries that are outside the BCBS members

- Step 2: Remove overlap resulting from data collection (mainly for mortgages, as data were collected on an aggregate level "Mortgages (All)" and sublevels
- Step 3:
  - Use of the IRB scaling factor for all banks (including those that submitted RWAs without the scaling factor of 1.06)
  - Use of data based on model in use and/or recent model (if model in use not available or for banks that changed their models significantly)
  - Focus on observations after 2006

Methods for aggregation of LGD data

- Data for A and E are based on aggregations of data submitted at the risk band level (ie the greatest level of granularity). The A/E ratio is computed for each combination of bank, asset class, year and model type; the overall average per bank and asset class is the simple average over the years, using the model-in-use figures if available, otherwise the A/E ratios from the recent model are used
- Method: Sum of actuals times EAD to sum of expected times EAD
A = actual loss rate times EAD
E = LGD outcome times EAD
- *Note: LGD estimates above 120% are dropped (for both A's and E's)*

Table A3.2		Step 1	Step 2	Step 3
Bank	LGD (raw)	LGD (after cleaning)	LGD (overlap removed)	LGD (final step, if data prior to 2007 were to be removed)
Total	17,016	11,458	8,362	7,623
Portion	100%	67%	49%	45%

EAD analysis

- Step 0: Raw data
  - Use of most recent data (ie submission 1, 2 or 3)
  - Amounts in local currency are all converted to EUR using historical foreign exchange rates
- Step 1: Data cleaning
  - Drop "true" duplicates (ie lines which are identical)
  - Drop observations without CCF estimates
  - Drop observations with CCF estimates < 0 & CCF estimates > 120 %
  - Drop observations without realised CCFs
  - Drop observations with negative total EAD
  - Drop observations with number of defaults <2 (due to high volatility of A/E ratio for number of defaults =1)
  - Drop observations with undrawn portion relative to the committed line > 5%[29]
  - Drop observations which are implausible (implied CCF from data does not match CCF or A/E ratio above 20 – and manual check)
  - Drop observations from countries that are outside the BCBS members
- Step 2: Remove overlap resulting from data collection (focus on revolving exposure only)
- Step 3:
  - Floor negative actuals (to put all banks at par – as some banks do this and others don't)
  - Focus on observations after 2006
  - Use of data based on model in use and/or recent model (if model in use not available or for banks that changed their models significantly)

Methods for aggregation of EAD data

- Data for A and E are based on aggregations of data submitted at the risk band level (ie the greatest level of granularity). The text below explains how the A/E ratio is computed for each combination of bank, asset class, year and model type; the overall average per bank and asset class is the simple average over the years, using the model-in-use figures if available, otherwise the A/E ratios from the recent model are used
- Method: Sum of actuals to sum of expected
  A = Defaulted exposure outcomes for "undrawn" portion
  E = EAD estimates for "undrawn" portion
- Note: EAD outcome are floored at zero (at the level of each single observation)

[29] Based on a sensitivity analysis, this cut-off level removed a significant portion of very high A/E ratios from the sample.

Table A3.3		Step 1	Step 2	Step 3
Bank	EAD (raw)	EAD (after cleaning)	EAD (overlap removed)	EAD (final step, if data prior to 2007 were to be removed)
Total	8,810	3,734	3,734	3,255
Portion	100%	42%	42%	37%

RWA analysis

- Data for A and E are based on aggregations of data submitted at the risk band level (ie the greatest level of granularity). The A/E ratio is computed for each combination of bank, asset class, year and model type; the overall average per bank and asset class is the simple average over the years, using the model-in-use figures if available, otherwise the A/E ratios from the recent model are used
- Method: Sum of actuals to sum of expected
  A = RWAs computed using actual default rates (PD), actual loss rates (LGD) and realised defaulted exposure amounts in place of PD, LGD and EAD, respectively, in the appropriate IRB risk weight formula
  E = RWA estimates computed using PD, LGD and EAD CCFs in the appropriate IRB risk weight formula
- Note:
  - RWA estimates (computed based on reported IRB parameters)
    - Use of cleaned PD data (from step 2)
    - Use corresponding LGDs reported by banks (if available, from PD tab) or average LGD (from LGD tab)
    - Use CCF estimates and portion of EAD subject to CCFs (from EAD analysis)
    - No use of floors (eg for PD)
    - Where applicable, use of maturity of 2.5 years and turnover of €27.5 million
  - RWA outcomes as reported by banks used as robustness check

Annex 4: Additional retail/SME results

Distribution of PD and LGD A/E ratios for banks with and without stress experience, by asset class

Table A4.1

A/E ratios for PDs		Mortgages	SME Corporate	SME Other Retail	QRRE
with stress	min	0.52	0.94	0.82	0.73
	p25	0.74	1.07	1.00	0.88
	Median	0.91	1.50	1.37	0.94
	p75	1.28	1.74	1.39	1.00
	max	1.94	2.27	1.71	1.55
	mean	1.04	1.48	1.26	0.97
w/o stress	min	0.09	0.15	0.31	0.09
	p25	0.63	0.49	0.63	0.58
	Median	0.82	0.76	0.82	0.85
	p75	0.91	1.00	0.90	1.04
	max	1.76	1.59	1.94	1.38
	mean	0.81	0.78	0.81	0.80

A/E ratios for LGDs		Mortgages	SME Corporate	SME Other Retail	QRRE
with stress	min	0.10	0.19	0.12	0.61
	p25	0.29	0.49	0.27	0.85
	Median	0.78	0.79	0.53	0.96
	p75	1.01	0.83	0.79	1.03
	max	1.08	1.09	0.93	1.08
	mean	0.66	0.67	0.53	0.91
w/o stress	min	0.00	0.09	0.08	0.47
	p25	0.22	0.40	0.43	0.75
	Median	0.32	0.58	0.56	0.83
	p75	0.66	0.77	0.86	0.90
	max	1.21	1.36	1.12	1.40
	mean	0.44	0.60	0.62	0.85