BIS Working Papers
No 1125

Banks' credit loss forecasts: lessons from supervisory data

by Martin Birn, Renzo Corrias, Christian Schmieder and Nikola Tarashev

Monetary and Economic Department

September 2023

JEL classification: G21, G28, G32, G33, E44, P52.

Keywords: Expected loss forecasts; regulatory capital; portfolio credit risk.

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

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

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Banks' credit loss forecasts: lessons from supervisory data

By Martin Birn, Renzo Corrias, Christian Schmieder and Nikola Tarashev[1]

Abstract

Focusing on credit risk, we compare banks' expected loss (EL) rates, collected confidentially by the Basel Committee on Banking Supervision from 2009 to 2022, and the corresponding actual loss (AL) rates, as reported in vendor data. Consistent with the use of through-the-cycle risk estimates for regulatory purposes, EL rates rarely move in line with AL rates over time, which helps explain a large precautionary element in Basel III capital requirements. We also find that the rank-order of EL rates across banks matches closely that of the AL rates, in line with recent and forthcoming regulatory efforts to improve risk-measurement practices. EL rates are more likely to be excessively optimistic on the heels of higher bank profitability and financial overheating, as captured by the credit-to-GDP gap.

Keywords: Expected loss forecasts; regulatory capital; portfolio credit risk.

JEL classification: G21, G28, G32, G33, E44, P52.

[1] We thank Iñaki Aldasoro, Irina Barakova, Stijn Claessens, Sebastian Dörr, Rodrigo Coelho, Neil Esho, Marc Farag, Mikael Juselius, Will Kerry, Wolf Müller and Egon Zakrajsek for very helpful comments on an earlier draft, and Burcu Erik, Roberto Ottolini and Jhuvesh Sobrun for valuable research support. The views expressed in this paper are ours and do not necessarily reflect the views of the Bank for International Settlements.

Contents

1. Introduction

With borrowing and lending at the core of bank business models, credit assessments are essential for banks' risk management. Credit losses featured prominently in the great financial crisis (GFC) (eg Bernanke (2009), Claessens et al (2010)) and 60% to 80% of internationally active banks' capital requirements reflect credit risk (BCBS (2023b), p 50–51). Today, debt service costs seem poised to rise on the back of significant debt accumulation during the low-for-long era and recent interest rate hikes. Under plausible scenarios, the rise of these costs may drive credit losses up to GFC levels (BIS (2023)). Ultimately, accurate credit-loss forecasts and/or regulatory conservatism – ie larger precautionary elements in the mapping from the forecasts to capital requirements – would be needed to ensure enough resources for absorbing credit losses that exceed loan loss provisions.

We use a novel dataset on banks' credit loss forecasts – or expected loss (EL) rates – to assess their accuracy. We juxtapose these EL rates with banks' actual loss (AL) rates and identify drivers of the discrepancies. At the heart of the exercise are exclusive confidential supervisory data, collected by the Basel Committee on Banking Supervision (BCBS). These data contain one-year EL rates, as reported by 65 internationally active banks to supervisors from end-2008 to end-2022. Combining EL rates on non-defaulted exposures with vendor *accounting* data on AL rates, we answer two questions:[2] (i) *how well do EL rates capture the evolution of AL rates*? and (ii) *does the rank-order of EL rates across banks align with that of AL rates*? We also relate the discrepancies between EL and AL rates to bank-level characteristics and macro-financial variables.

EL rates perform differently along the time and cross-section dimensions. First, we find that they generally fail to capture the time profile of AL rates. As EL rates tend to miss both spikes and trends in AL rates, the correlation of year-to-year changes in the two series is statistically significant for only 15% of the banks. This implies that a conservative mapping from EL rates to capital requirements is needed to ensure enough loss-absorbing resources at each point in time. We estimate that – over our sample period – such conservatism would result in capital being (at least) twice as large as actual losses for three-quarters of the banks.[3] The data suggest similar or stronger conservatism in actual capital requirements. Our second finding is that banks fare well when it comes to signalling the riskiness of their credit portfolios *relative* to that of other banks' portfolios. Specifically, we find that, in each of the sample years, the rank-ordering of EL rates across banks closely matches that of the corresponding AL rates.[4]

Regression analysis confirms and expands these descriptive findings. We conduct this analysis in two steps. In the first step, we regress AL rates only on the corresponding EL rates (and a constant). We find that, on their own, EL rates can

2 The focus on non-defaulted exposures places the emphasis on estimates of probability of default, rather than loss-given-default.

3 We do not have information to shed light on whether greater conservatism would have been needed if banks had used pro-cyclical ("point in time") PD estimates.

4 With the exception of 2021, where the correlation between EL and AL rate is somewhat disrupted, in line with the exceptional nature of the Covid-19 pandemic and related support measures.

explain less than 5% of the volatility in AL rates over time. By contrast, EL rates explain almost 70% of the dispersion of AL rates across banks.

In a second step, we study drivers of the portion of AL rates that is not explained by EL rates, ie "step-one errors". In a sign that EL rates do not account fully for persistence in credit losses, we find that lagged AL rates help explain step-one errors. We also find robust statistically significant linear relationships between step-one errors, on the one hand, and country-level credit-to-GDP gaps and bank-level return-on-assets (RoA) and price-to-book ratios (PtB), on the other. Consistent with banks abstracting from macro indicators of overheating when forecasting credit losses, a one-standard-deviation rise in the credit-to-GDP gap accounts for a 0.12 standard deviation rise in step-one errors two years down the road. In turn, a one-standard deviation increase in RoA accounts for a 0.35 standard deviation rise in the discrepancy between AL and EL rates one year later. This is consistent with higher profitability introducing excessive optimism in risk measurement. Finally, a one-standard deviation increase in the PtB ratio accounts for a 0.17 standard deviation decline in the step-one error one year later, consistent with higher valuations enabling banks to afford greater conservatism in their EL rates.

Throughout, we are conscious of potential inconsistencies between the supervisory data behind EL rates and the vendor accounting data behind AL rates. While the former comprise on- and off-balance sheet exposures to credit risk, the latter cover systematically only balance sheet positions and may conflate securities exposed to credit risk with those exposed only to market risk. Moreover, while the supervisory data draw a clear line between defaulted and non-defaulted exposures, accounting practices blur this separation in the vendor data.[5] We thus check – and confirm – the robustness of our results to the exclusion of observations for which measurable inconsistencies between the two data sets – ie a wedge between the credit risk exposures underpinning AL and EL rates – exceed a particular threshold. In addition, we control for potential differences between the defaults that affect AL rates and those underpinning the losses that EL rates are supposed to forecast. Ultimately, our exercise is an evaluation of regulatory credit risk estimates on the basis of the best available cross-jurisdictional data on large internationally active banks.[6]

Our findings are related to the effects of post-GFC regulatory initiatives. For one, authorities sought to mitigate the tendency of risk assessments to be overly optimistic in tranquil times and spike in stress – ie to mitigate pro-cyclicality (BCBS (2021)). These efforts underpin banks' use of "through the cycle" estimates of probabilities of default (PDs).[7] Anchored in long-term historical default rates, such estimates tend to be stable and it is thus hardly surprising that they lead to EL rates that miss the evolution of AL rates.[8] In addition, recent policy efforts have sought to

5 Defaulted exposures remain in the supervisory data until they have been written off, recovered, or sold off to a separate legal entity, at which point they are removed from the data.

6 Ong et al (2023) present an attempt at collecting *publicly* available data on expected and actual losses across jurisdictions. They do not cover the supervisory data on banks' expected losses, which we use.

7 Related is banks' use of "downturn" loss-given-default (LGD), which reflect periods of exceptionally high credit losses.

8 The countercyclical buffer requirement could in principle correct (partly) for the failure of banks' estimates to capture the evolution of AL (see also Herz and Keller (2023)). Being another outcome of policy efforts to combat procyclicality, this buffer seeks to enhance loss-absorbing resources when banks ignore the build-up of risks.

ensure that differences in regulatory capital reflect genuine differences in underlying risks, rather than differences in risk-measurement practices across jurisdictions and entities (BCBS (2017)). The performance of EL rates in explaining the dispersion of AL rates across banks is in line with such efforts.

This paper complements previous assessments of banks' regulatory credit risk estimates. In the aftermath of the GFC, the BCBS explored the variability of such estimates across banks for the same hypothetical portfolios (BCBS (2016)), which later informed constraints on the design of models in the final Basel III package (BCBS (2017)). Hardy and Schmieder (2013) and Lewrick et al (2021) examined whether Basel III capital requirements and voluntary (or "management") capital buffers were sufficient to absorb losses at different levels of potential stress post-GFC. Most recently, BCBS (2022) concluded that the Basel III reforms increased the resilience of large internationally active banks and reduced systemic stress. Our analysis suggests that this resilience rests on accurate reporting of relative riskiness and conservative mapping from risk estimates to capital requirements.

Implicit justification for this conservatism and the attendant use of through-the-cycle PD estimates comes from another branch of the related literature. Namely, real-time (ie out-of-sample) forecasts of turning points in default-related losses have proven to be inherently difficult. A number of papers have been sceptical about the usefulness of such forecasts (Covas and Nelson (2018), Abad and Suárez (2017), Chae et al (2018), Krüger et al (2018), Goncharenko and Rauf (2020), and Loudis and Ranish (2019)). Such scepticism has underpinned the use of through-the-cycle credit-risk estimates, which are consistent with our findings about the performance of EL rates along the time dimension. Despite recent advances in econometric modelling and the selection of indicator variables (eg Lu and Nikolaev (2021) and Juselius and Tarashev (2020)), much uncertainty in the forecasts seems inevitable.

The rest of the paper is organised as follows. We outline the building blocks of regulatory capital requirements for credit risk and provide an illustrative example in Section 2. We present the data in Section 3. In Section 4, we juxtapose EL and AL rates, separately along the time and the cross-section dimensions. In that section, we also verify whether actual and forthcoming regulatory requirements exhibit the conservatism that our analysis calls for. In Section 5, we study the time and cross-section dimensions jointly, with panel regressions. Section 6 concludes.

2. Loss absorbing resources for credit risk: basics

2.1 Building blocks of regulatory capital

Banks' prudential requirements seek to ensure loss-absorbing resources that would be sufficient to cover actual credit losses with a high probability – 99.9% over one year. Regulation envisages two approaches to attain this objective. The so-called standardised approach (SA) is mostly used by smaller banks, but also by large banks in some cases,[9] often limited to certain low-default exposures. Under the SA, banks allocate exposures to *relative*-risk categories (in some jurisdictions, based on

[9] US authorities have issued a proposal that would require all banks in their jurisdiction to use the SA (Board of Governors of the Federal Reserve System (2023)).

assessment by credit rating agencies) and then a supervisory formula maps each category to a capital requirement. Alternatively, banks may, subject to supervisory approval, apply the so-called internal ratings-based (IRB) approach, for which they need to estimate risk parameters, including the probability of default (PD) and loss-given-default (LGD).[10] Implementation of an output floor (be it the floor in the Basel II framework, a jurisdiction-specific floor or the Basel III "output floor" – with expected implementation from January 2023), will limit to the capital reduction that a bank may obtain from adopting the IRB approach instead of the SA.[11]

The rest of this section outlines the IRB approach, since it uses as inputs the risk parameter estimates that we analyse below.

The IRB approach is rooted in a stylised model of portfolio credit risk. In this model, the portfolio is "asymptotic", consisting of a very large number, n, of small exposures. In addition to an idiosyncratic risk factor, a single macro risk factor governs the default on each exposure. Each factor is normally distributed. Assuming also that LGD_i, one-year PD_i and loading on the macro factor (ρ_i) are known parameters for each exposure i, the model implies that the credit losses per unit of portfolio exposure will exceed the following value-at-risk (VaR) over a year with probability α:

$$\mathrm{VaR}\left(\alpha\right)=\sum_{i=1}^{n} LGD_i\ \Phi\left(\frac{\Phi^{-1}\left(PD_i\right)-\rho_i\ \Phi^{-1}\left(\alpha\right)}{\sqrt{1-\rho_i^2}}\right)$$

where Φ denotes the CDF of a standard normal variable. Setting loss-absorbing resources to VaR(0.1%) is consistent with a target probability of 99.9% that the bank will withstand one-year losses without inflicting losses on its debt holders.

In the IRB framework, the model is used as follows. Loss absorbing resources are split in two parts: provisions and capital. Provisions are equal to expected losses, $EL=\sum_{i=1}^{n} LGD_i\ PD_i$. Regulatory capital requirements (K) are set equal to a fraction (k) of risk-weighted assets (RWA), which amount to unexpected losses, VaR(α)–EL, adjusted for exposures' maturity (M).[12] For the determination of RWA, prudential regulation assumes that ρ_i is a known function of PD_i, which avoids inherent difficulties in estimating ρ_i.

Specifically, capital requirements – $K\left(PD_i, LGD_i, M_i\right)$ – for credit to borrower i, per unit of exposure to this borrower are equal to:

[10] Within the IRB approach, the Basel framework distinguishes between the Advanced IRB approach, which requires banks to estimate both PDs and LGDs, and the Foundation IRB framework, under which banks need to estimate only PDs and adopt fixed regulatory LGDs. Basel III rules agreed in December 2017, do not allow the Advanced IRB approach for exposures to banks and other financial institutions, large and mid-sized corporates, and equity (for the latter, the only option is the SA).

[11] The Basel III output floor works as follows: "banks' calculations of risk weighted assets generated by internal models [ie "the output"] cannot, in aggregate, fall below 72.5% of the risk-weighted assets computed by the standardised approaches. This limits the benefit a bank can gain from using internal models to 27.5%" (see BCBS (2017)).

[12] The forecast horizon underpinning provisions tends to be longer than the one-year horizon underpinning capital. This distinction is set aside, as it is not central for our arguments and because prudential regulation seeks to ensure that provisions are not lower than one-year expected losses.

$$K_i = K\left(PD_i, LGD_i, M_i\right) = k \cdot RWA\left(PD_i, LGD_i, M_i\right), \text{ where:}$$

$$RWA_i = 12.5\left(LGD_i\ \Phi\left(\frac{\Phi^{-1}\left(PD_i\right) - \rho\left(PD_i\right)\Phi^{-1}\left(0.1\%\right)}{\sqrt{1-\rho\left(PD_i\right)^2}}\right) - PD_i \cdot LGD_i \right)$$
$$\cdot \frac{1+\left(M_i - 2.5\right)b\left(PD_i\right)}{1-1.5\,b\left(PD_i\right)} \tag{1}$$
$$\rho\left(PD_i\right) = \rho_{low}\frac{1-e^{-S\cdot PD_i}}{1-e^{-S}} + \rho_{high}\left(1-\frac{1-e^{-S\cdot PD_i}}{1-e^{-S}}\right)$$
$$b\left(PD_i\right) = a - c\cdot\ln\left(PD_i\right)$$

and the parameters ρ_{low}, ρ_{high}, S, a and c are known positive constants. The first three differ with the exposure type: corporate, sovereign or bank; small or medium sized entity; specialised lending; residential mortgage; qualifying retail; other retail.[13]

The relative capital requirement (RCR) k determines regulatory capital per unit of RWA. If the IRB modelling assumptions match reality, the credit maturity is M_i=1 year, and the risk parameters are exactly estimated, then k=8% will result in loss-absorbing resources that are exceeded by losses over a one-year horizon with 0.1% probability. However, banks assign through-the-cycle PDs – often equal to the long-term default rate in internally determined *relative*-risk categories – and downturn LGDs. In addition, authorities are aware of inevitable mis-specifications and estimation errors. All this provides rationale for supervisory add-ons, on top of the value of k that any specific model would imply.

While regulation imposed only minimum requirements prior to the GFC, buffer requirements were added to regulatory minima post-GFC. Breaching a regulatory minimum should in principle involve "terminal" penalties, such as a license withdrawal or a declaration of (probable) failure, with supervisors seizing all available loss-absorbing resources to ensure an orderly recovery or resolution of the bank. By contrast, the breach of a buffer allows the bank to operate as a going concern, even though it does suffer penalties: eg restrictions on maximum distribution amounts, heightened supervisory monitoring, and an obligation to submit a plan for replenishing the buffer. Each minimum or buffer requirement is expressed in terms of a specific RCR, with the RCRs being additive within a bank.

Thus, the risk-based loss-absorbing resources for credit risk have the following key components: PD – which, together with LGD, delivers EL and together with LGD and M maps into RWA – and relevant regulatory RCRs.[14]

While loss absorbing resources are in practice determined by risk parameters at the exposure level, our data is on bank-level parameters (see below). Indeed, a key implication of the underlying model is that the contribution of an exposure to portfolio risk is set in isolation (expression (1)), by considering only risk parameters of

13 See BCBS (2023a), CRE31 for the exact parameterisations and adjustments to $\rho\left(PD_i\right)$ in the case of exposures to unregulated financial institutions and SMEs.

14 Regulation refers to several types of capital that differ with respect to their loss absorbency (ie Common Equity Tier 1 (CET1), Additional Tier 1 and Tier 2 capital). We will abstract from this distinction.

this particular exposure. Given our bank-level data, we assume that each bank's credit portfolio comprises homogeneous exposures.

2.2 Illustrative example with a hypothetical bank

We now illustrate the different types of regulatory capital requirements under the IRB approach. We work with a hypothetical portfolio of homogeneous exposures, assume that a maturity parameter M of 2.5 years and that LGDs increase in PDs levels, in line with evidence (Hardy and Schmieder (2013)). We use this portfolio to conduct comparative statics with respect to PD. Then we examine how the loss-absorbing resources implied by two alternative PD estimates compare to historical losses on investment grade (IG) and high yield (HY) bond exposures.

Graph 1 (left-hand panel) portrays expected losses and capital requirements for PDs ranging from the minimum one-year level that the IRB framework allows for (0.03%) and the highest observed default rate on a broad portfolio of corporate debt.[15] The mapping from PD to capital requirements is concave. The shape of the function reflects the assumption that asset correlations are lower at higher PDs (BCBS (2005), p 12). This calibration choice also dampens the increases in capital requirements as PDs spike up, thus mitigating procyclicality of the banking sector. This feature implies a declining ratio of required capital to EL rates: eg minimum capital requirements are 47 times larger than EL at the lowest PD but only 1.1 times larger at the highest. The total Pillar 1 requirements (ie minimum plus regulatory buffers) typically faced by global systemically important banks (G-SIBs) are twice as large as the Pillar 1 minimum requirements.

Next, we illustrate how a *hypothetical* bank would have absorbed the credit losses on a corporate portfolio over the past 100 years if it had reported a constant PD and had been subject to regulatory requirements under the Basel III IRB approach.[16] First, we consider an IG portfolio, assume that the bank sets the PD to the average historical default rate on that portfolio and compare its required capital with deviations of AL from EL (Graph 1, centre panel). We see that the regulatory buffers comfortably absorb IG portfolio losses in all years, keeping the bank far from breeching the minimum requirements. The takeaway for a representative HY portfolio is different (right-hand panel). For instance, the total Pillar 1 capital requirements would have been exhausted by the actual losses on this portfolio in 1932/1933. Absent additional capital, this implies a failure of the hypothetical bank and material losses to all its debt holders. In addition, much or all of the regulatory buffers would have been wiped out in 1970, 1990, 2001 and 2009. The breech of minimum requirements in 1990 and 2009 provide examples of instances in which today's authorities would have taken the hypothetical bank over from its management in order to protect its debt holders.[17]

15 As reported by Moody's (2022), for its universe of corporate exposures between 1920 and 2021. The highest default rate was 8.53% in 1933.

16 This stylised example assumes that management and Pillar 2 (ie supervisory) buffers are zero. It thus overstates the extent to which losses would deplete regulatory buffers.

17 This illustrative example notwithstanding, crises could be driven or deepened by losses on assets that only appeared to be of high quality and thus received low risk weights. This was the case of highly-rated securitisations in the GFC and, more generally, has been the case of housing-related busts.

Illustration of regulatory capital and its loss absorbency (hypothetical bank)

Percent of credit exposure Graph 1

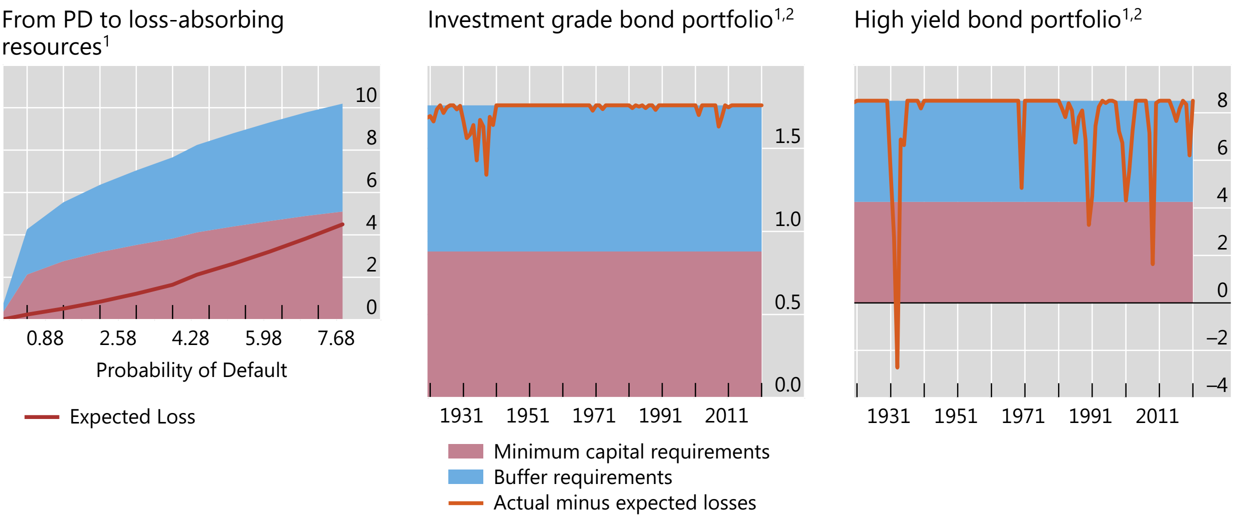

[1] Based on LGD = 3.38 · PD+0.2396, as estimated by Hardy and Schmieder (2013). Expected losses are set equal to PD (historical range of default rates observed from 1920-2021) times LGD. Minimum capital requirements are set to 4.5% of risk-weighted assets (RWA); buffer requirements to the sum of capital conservation buffer of 2.5%, counter-cyclical capital buffer of 1% and G/D-SIB buffer of 1% of RWA. [2] PD is equal to the average default rate of the IG (centre panel) and HY entities (right-hand panel) in Moody's rating universe from 1920 to 2021.

Source: Moody's.

3. Data and key metrics

3.1 Data

The main value of our analysis stems from confidential bank-level data from the Basel Committee's Basel III monitoring exercises. These data are collected semi-annually to provide detailed information on: banks' IRB credit risk parameters,[18] IRB and SA risk-weighted assets, corresponding Pillar 1 capital requirements and eligible regulatory capital as of end-June and end-December from 2008 to 2022.[19,20] The cross section expands over time, as the number of reporting banks increases from 31 in 2008 to 57 in 2022. Overall, 65 banks enter our sample, of which 36 from Europe, 18 from the Americas and 11 from the rest of the world. Of these entities, 26 are currently G-SIBs.[21] More generally, the sample includes about 57% of the banks that the BCBS

[18] Risk parameters are available separately for the corporate, bank, sovereign and retail asset classes. For many banks, there is an additional breakdown of the corporate asset class into small and medium-size enterprises, specialised lending exposures and large corporates. The retail asset class is often broken down into mortgages, qualifying revolving retail exposures and regulatory "other retail".

[19] Exceptions are Canadian and Japanese banks, for which the regulatory reporting dates are at the end of the corresponding financial year, which differs from that of the calendar year.

[20] Banks' Pillar 3 disclosures provide similar data points. The BCBS data we use improve on these disclosures by standardising the information across banks and over time. In addition, an automated collection process for the BCBS data allows for in-depth quality checking procedures.

[21] Chinese G-SIBs do not distinguish between defaulted and non-defaulted exposures and are therefore excluded from this analysis and another G-SIB, which does not report data.

classifies as large internationally active banks for the purpose of its monitoring reports (BCBS (2023b)). While a balanced sample of 65 banks over 14 years means 910 observations, the number of banks that reported to the BCBS with sufficient data quality varied over time, resulting in 643 observations in our sample.

The Basel III supervisory data is matched anonymously with bank-level financial statements data from Fitch.[22] While these accounting data are also available quarterly, the highest quality is at the yearly frequency, which we adopt for both datasets. Specifically, we use gross loans, net loans, total securities, loan impairment charges, securities and other credit impairment charges, return on assets (RoA), price to book ratios (PtB) and customer deposits-to-total funding (DepShare).

Finally, we also employ country-level data. These are: credit-to-GDP gaps and debt service ratios (provided by the BIS), and real GDP growth rates (provided by the IMF).

3.2 Key metrics

At the heart of our analysis are comparisons between a specific risk metric reported by banks in the supervisory data – the EL rate for their banking book portfolios – and the AL rates incurred by banks and reported in the Fitch data. We now explain the construction of these two metrics and reasons for discrepancies between them.

Expected loss rate

The EL rate on a credit exposure is the probability of default (PD) times the percentage loss-given-default (LGD). We obtain it at the portfolio level by dividing (i) the expected losses over **year *t*** on a bank's IRB credit risk exposures that are non-defaulted at the end of **year *t*–1**, by (ii) the stock of these exposures at the end of **year *t*–1**.[23] Concretely, based on information available at the end of year *t*–1:

$$EL\ rate_{t-1} = \frac{EL\ over\ year\ t\ on\ non\text{-}defaulted\ IRB\ exposures\ at\ end\text{-}year\ t-1}{non\text{-}defaulted\ IRB\ exposures\ at\ end\text{-}year\ t-1}$$

Actual loss rate

With the AL rate, we seek to measure the losses that the EL rate forecasts. Concretely, we divide the following two data fields: (i) the combined impairment charges for loans and securities over **year *t*** by (ii) the corresponding loan and securities exposure at the end of **year *t*–1**:

$$AL\ rate_{t-1} = \frac{loan,\ securities\ and\ other\ impairment\ charges\ over\ year\ t}{net\ loans + gross\ securities\ at\ end\text{-}year\ t-1}$$

[22] BIS staff without access to the confidential supervisory data provide the vendor data without bank names, but using a specific matching key. Different BIS staff match the anonymised vendor data with the confidential supervisory dataset.

[23] As regards a bank's total IRB credit-risk exposure, the data distinguish between non-defaulted exposures and defaulted exposures, ie exposures to borrowers that are past due for more than 90 days or who are unlikely to pay. See BCBS (2023a), paragraph CRE36.68. If the split between defaulted and non-defaulted exposure and EL is not available for all asset classes, we extrapolate based on data for other the asset classes, as long as these represent at least 50% of total exposure or EL.

where net loans are equal to gross minus previously impaired loans. We prefer this measure of AL rates to alternatives because it covers a broader set of assets subject to credit risk and because of the timeliness with which it reflects credit losses.[24]

First juxtapositions of EL and AL rates

For a first look at the EL and AL rates, we compare their medians and volatilities at the bank level (Graph 1). As regards medians, we find a close alignment – a correlation coefficient of 76% (left panel), even though the two rates can differ substantially in some cases. Turning to volatilities, that of AL rates is markedly higher than that of EL rates (right panel). This could be due to the fact that banks use through-the-cycle PDs and downturn LGDs, which tend to stabilise EL rates. In addition, inevitable deviations from expectations render AL rates more volatile than EL rates.

Unfortunately, some of the differences between AL and EL rates could be an artefact of the data. We discuss this issue and remedies next.

Comparison of expected-loss and actual-loss rates

In percent | Graph 2

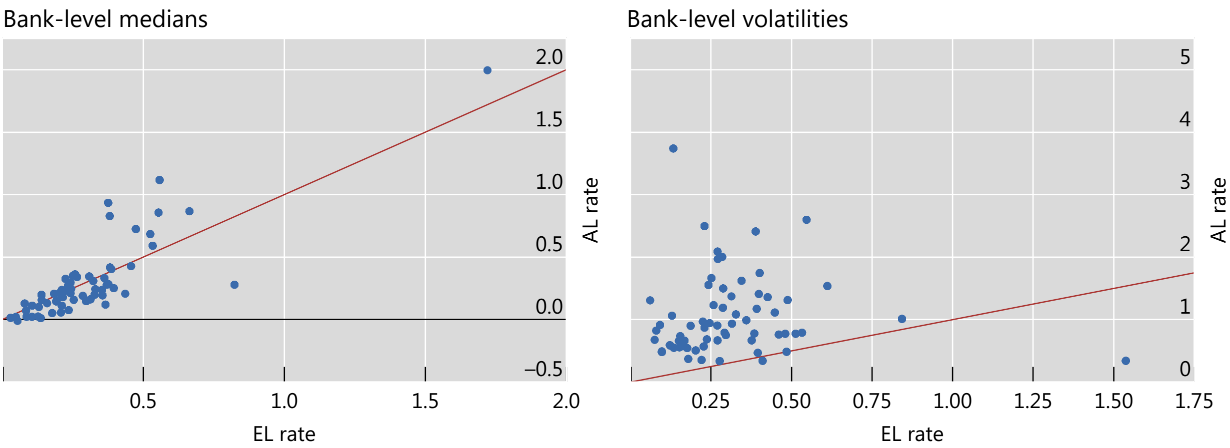

Note: Each dot corresponds to a bank: the median AL and EL rates over time (left-hand panel) and the corresponding coefficients of variation (ie the standard deviation divided by the mean, right-hand panel) over the 2009-2022 period.

Source: BCBS and Fitch data.

Inconsistencies between data on EL and AL rates

Differences between the data sources underpinning EL and AL rates may introduce a degree of inconsistency. The first three differences we list next imply that, all else the same, the EL rate would be larger than the AL rate.

1. The securities in the accounting data could include instruments that are subject only to *market risk* and are thus excluded from the supervisory data on credit exposure. This raises the denominator of the AL rate without affecting the EL rate.

[24] An alternative measure uses net charge-offs on on-balance sheet loans. While this measure would reflect more accurately the ultimate losses incurred on loans, it abstracts from credit losses on securities and, for some jurisdictions, signals losses with a longer delay. See also Ong et al (2023).

2. Likewise, the supervisory data behind EL rates refer to non-defaulted exposures only, whereas the denominator of AL rates could include also *non-performing securities*.

3. The EL rate reflects only banks' IRB credit exposures, whereas the AL rate also includes exposures under the SA. The latter exposures tend to carry lower credit risk (notably, sovereign paper).[25]

By contrast, the following two effects would have the opposite implications.

4. The supervisory data include *off-balance sheet credit exposures*, eg credit lines and counterparty credit risk, which are not included in the accounting data. Since these exposures tend to carry less credit risk than the typical balance sheet exposure, their inclusion would tend to reduce the EL rate below the AL rate.

5. Exposures booked after the end of year t–1 could become impaired and thus raise *AL rate*$_t$. Such losses are not meant to be forecasted by the *EL rate*$_{t-1}$.

Finally, the implications of the following are ambiguous:

6. *AL rate*$_t$ could incorporate reassessments of the degree of impairment of exposures that were impaired before year t. These reassessments could raise or lower *AL rate*$_t$ without being reflected in *EL rate*$_{t-1}$ as it refers only to exposures that were non-defaulted at the end of year t–1.

In order to reduce the impact of data inconsistencies on our findings, we consider two remedies in the regression analysis below.[26] First, we introduce a control variable that addresses the 6th source of inconsistencies. Second, we note that, while we cannot determine whether forecast errors or data inconsistencies drive differences between expected and actual loss *levels* – ie the numerators of the corresponding ratios – only the latter driver can create a wedge between the denominators. Thus, we run robustness checks after dropping bank-year observations for which the absolute difference of the *EL rate*$_{t-1}$ and *AL rate*$_t$ denominators is at least 50% apart.

4. Adequacy of loss estimates and capital requirements

4.1 EL vs AL: cross section

Next, we study the *relative* riskiness of credit portfolios, both in terms of banks' own assessments (as captured by EL rates) and in terms of realised losses (AL rates). We first consider ELs and ALs separately, analysing the stability of their rank ordering in the cross section over time. Then, we juxtapose the rank-ordering of EL and AL. A tight alignment between the two rank-orderings indicates that banks' credit risk assessments are accurate in relative terms and that any discrepancies in the two underlying datasets have a largely uniform effect across banks in each year.

The rank-ordering of EL rates is remarkably stable. We illustrate this in Graph 3 (left panel) where the close alignment of the dots along the 45-degree line indicates

[25] Regulatory EL are not calculated for SA exposures, and inferring EL rates for SA exposures requires a number of assumptions that would generate noise.

[26] The effects of these "remedies" do not change any of our findings in a material way. We report explicitly such effects only in the context of the regression analysis (Section 5).

that a relatively low (high) assessment of credit risk in 2017 went hand-in-hand with a similarly low (high) assessment in 2020. The underlying correlation is 91%. As regards all the rank correlations of ELs observed three years apart in our data, the range is from 65% to 96% and median is 84%.

There is similar, albeit weaker stability in terms of ALs. This is illustrated by the greater dispersion around the 45-degree line in the middle panel of Graph 3, where the underlying rank correlation is 75%. Overall, such correlations have a median value of 63% and range between 50% and 78%.[27]

Next, we assess whether the rank-ordering of EL is consistent with that of AL. The right panel of Graph 3 compares the ranks of ELs, as reported by banks at end-2019, and the corresponding AL rate over 2020. Again, we observe a close alignment along the 45-degree line; rank correlation is 70%. In the overall sample, this correlation has a median of 66% and ranges from 55% to 76% across years.[28] In sum, this is preliminary evidence that banks' risk estimates tend to capture accurately the *relative* riskiness of their credit portfolios.

Rank-ordering bank-level loss rates

Percentiles

Graph 3

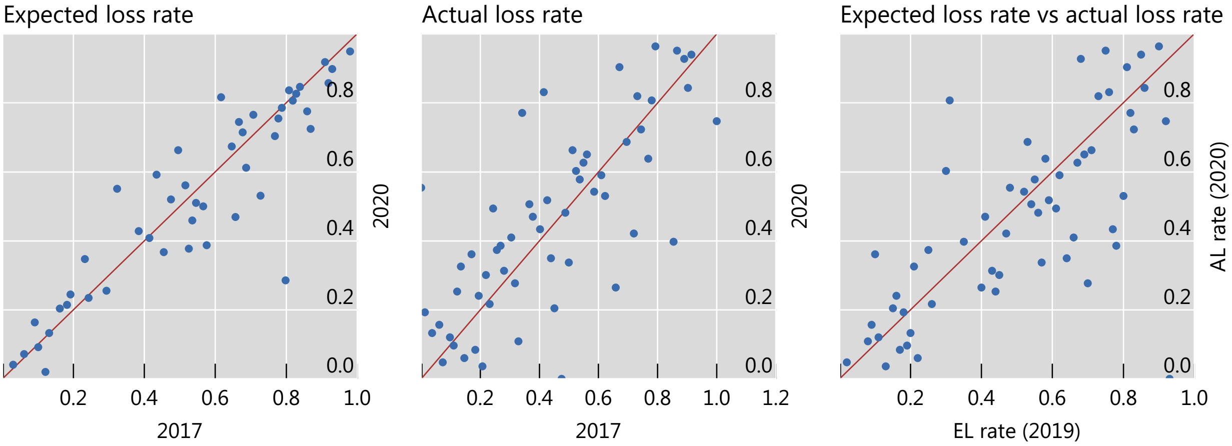

Source: BCBS and Fitch.

4.2 EL vs AL: time dimension

To study the extent to which banks' risk estimates anticipate the evolution of actual losses, we conduct two exercises: first, we calculate the correlation of year-on-year changes in bank-level ELs and the corresponding changes in ALs. Second, we assess the overall alignment between AL and the loss-absorbing resources implied by the internal risk-based (IRB) regulatory approach. For the latter exercise, we start by deriving – again at the bank level – the lowest *time-invariant* RCR that – together with the reported ELs and RWAs – would imply loss-absorbing resources at least as high as the actual losses in *any* year of our sample. We cumulate over time the AL and

[27] This range excludes losses in 2021, which lead to an insignificant correlation of 19%, in line with the exceptional nature of the Covid-19 pandemic.

[28] Again, this excludes losses in 2021, which give rise to an insignificant correlation of -9%.

divide them by the "just enough" cumulative loss-absorbing resources implied by this RCR. This ratio is at most one by construction. A value closer to one indicates that smaller loss-absorbing resources are needed to compensate for errors in banks' risk estimates.[29] Since each exercise embeds an implicit rebasing of the risk estimates, we abstract from any discrepancy between their average level and that of ALs.[30]

To benchmark the results of the two exercises, we refer to parallel results based on alternative risk estimates. These are also derived in (quasi) real-time, ie on the basis of data available before the corresponding ALs occur. One of the alternatives is based on the forecasting model in Juselius and Tarashev (2020) – henceforth, JT – which employs indicators of financial overheating to forecast the aggregate losses on loans extended to US banks. For our exercise, we use the actual realisation of the latter losses as ALs and the one-year forecasts of this model as ELs. For the second alternative, we use the one-year expected default frequencies (EDFs) from Moody's KMV as EL and the corresponding corporate default rates (from Moody's Investors Service) as AL. To translate the JT ELs and the EDFs into RWAs, we assume an LGD of 45% and a maturity of one year in the regulatory formula for RWA.

We find that changes in banks' EL estimates tend to be unrelated to corresponding AL changes in the following year. We see this in the first row of Table 1: the median correlation across banks is 0%; the correlations are positive and statistically significant for only six of the 60 banks in the sample for which at least five years of relevant data are available. This is in line with the statistically insignificant correlation based on EDFs but contrasts the high and statistically significant correlation implied by the JT model.

We interpret these results with caution. For one, they would reflect year-on-year blips that banks would see through with through-the-cycle forecasts that seek to capture slow-moving processes, such as a secular increase or decline in credit risk. Moreover, banks build loss absorbing resources to cover both expected and unexpected losses, not least because some forecast errors are inevitable. These considerations motivate the second exercise.

We find that, because of a misalignment between banks' risk estimates and the overall time profile of AL, loss absorbing resources need to be quite conservative in order to be sufficient to avoid a default (Table 1, second row). For more than three-quarters of the banks, these resources cumulate over the 14 years in our sample to more than twice as much as actual losses – ie the ratio we derive is smaller than 50%. Again, while this is comparable to the implications of EDFs, it is markedly worse than those of the JT forecasts.[31]

29 Of course, a real-world authority would not have the perfect foresight needed to determine the minimum time-invariant RCR that generates just enough loss absorbing resources to cover ALs. We assume perfect foresight to obtain a simple metric capturing the alignment of the *time profiles* of AL and the loss-absorbing resources implied by risk estimates.

30 This exercise assumes that any inconsistencies between the supervisory and accounting data do not influence the *relative* time profiles of EL and AL rates. We examine such inconsistencies in Section 5.

31 It is in principle possible that the poor performance of the EDFs is an artefact of comparing them to actual losses from a different data source. A strong indication that the two sources are consistent comes from an additional exercise in which we pretend that forecasts at the beginning of year *T* were constructed at the beginning of year *T*–1. This improves the performance, with the metrics in Table 1 (last column) changing to 80% and 0.63, respectively. We conclude that, rather than comparability issues, the poor performance of EDFs stems from a failure to capture turning points in loss rates.

Actual losses and risk estimates: alignment of time profiles

Table 1

	EL to AL						Benchmark (JT)		Benchmark (EDF)	
	Min	25th	Median	75th	Max	Nr	short	long	short	long
Correlation(ΔAL, ΔEL)	-91%	-34%	0%	27%	92%	60 (6)	82%***	81%***	38%	31.5%
Cum(AL)/Cum(LAR)	0.01	0.24	0.34	0.47	0.86	60	0.61	0.53	0.45	0.45

Note: The first six columns correspond to the sample of 60 banks in the supervisory data for which at least five years of relevant data are available. "Benchmark (JT)" uses the methodology in Juselius and Tarashev (2020) to forecast the aggregate charge-off rates of US banks' loan portfolios. "Benchmark (EDF)" corresponds to the juxtaposition of Moody's KMV EDFs and corporate credit loss rates. AL = actual (credit) loss rate, EL = expected loss rate, Δ indicates a one-year change. The correlation coefficient is Pearson's. Cum = cumulative sum over the relevant sample period (2008 to 2021 for the first five columns and the "short" benchmark columns; 2000 to 2021 for the "long" column for benchmark (JT) and 2002 to 2021 for the "long" column for benchmark (EDF)). Loss-absorbing resources (LAR) are equal to EL plus $k \cdot$ RWA, where k is the minimum time-invariant relative capital requirement (RCR) ensuring that LAR is at least as high as the corresponding AL in each year. Of the bank-level correlations, six are statistically significant at the 5% level. *** indicates significance at the 1% level.

We illustrate specific reasons for the conservatism of "just-enough" loss-absorbing resources in Graph 4. In the left panel of that graph, we focus on the bank with the median ratio in the second row of Table 1. The blue line, which plots these resources, is always above and touches once the red line, corresponding to actual losses. The area between the two lines is large because the bank misses the spike in 2009 and then forecasts a downward trend that is shallower than in reality. In comparison, loss-absorbing resources need to be less conservative under the JT forecasts, mainly because these capture more accurately the turning points in ALs and the relative magnitudes of high and low levels of AL.

"Just enough" capital requirements[1]

In Percent

Graph 4

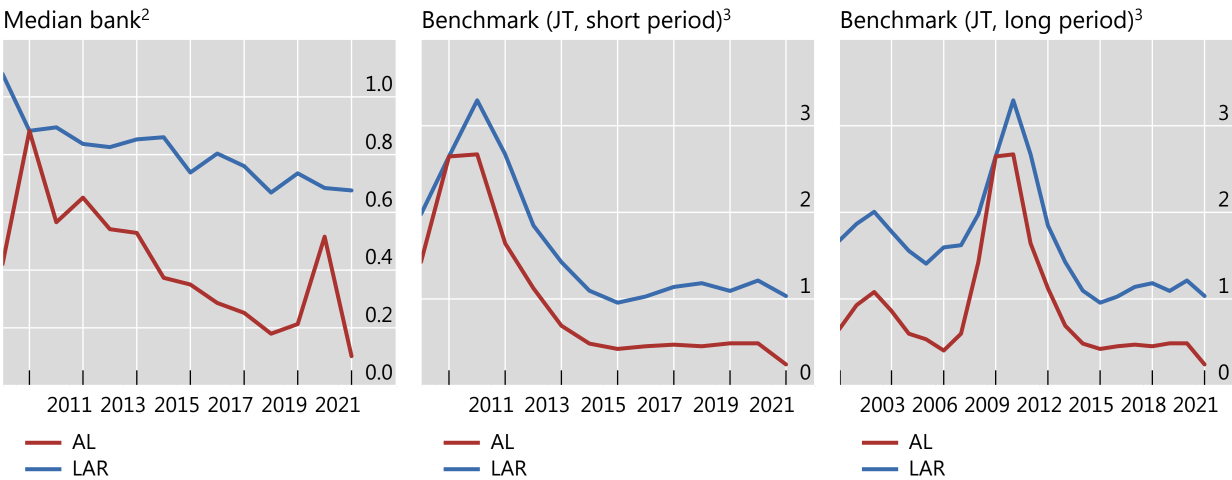

LAR = Loss-absorbing resources, AL = Actual Losses

[1] Based on the lowest time-invariant relative capital requirement (see expression (1)) that – together with the reported ELs and RWAs – would imply loss absorbing resources at least as high as the actual losses in each year. [2] Per Table 1 (send row). [3] Based on Juselius and Tarashev (2020), where actual losses are the aggregate charge-off rates on US banks' loan portfolios.

Source: BCBS and Fitch.

4.3 Conservatism in actual capital requirements

Given that the above analysis argues for bank capital that is materially above AL most of the time, we verify whether actual regulatory requirements exhibit conservatism. In particular, for each bank, we add its actual losses on credit risk exposures over time and divide the sum by the corresponding sum of loss-absorbing resources: regulatory buffers plus regulatory minima plus expected losses (which proxy for provisions) for SA and IRB credit exposures.[32] And we plot the distribution of the AL-to-LAR ratios across banks in Graph 5. This parallels the second row in Table 1 but now using actual capital requirements. One of the distributions reflects current capital requirements and the other one the requirements that would have been in place if a forthcoming tightening – stemming from the final Basel III framework and the so-called output floor – had been already implemented.[33]

Ratio of actual losses to loss-absorbing resources

Density of bank-level averages (2009–2022) Graph 5

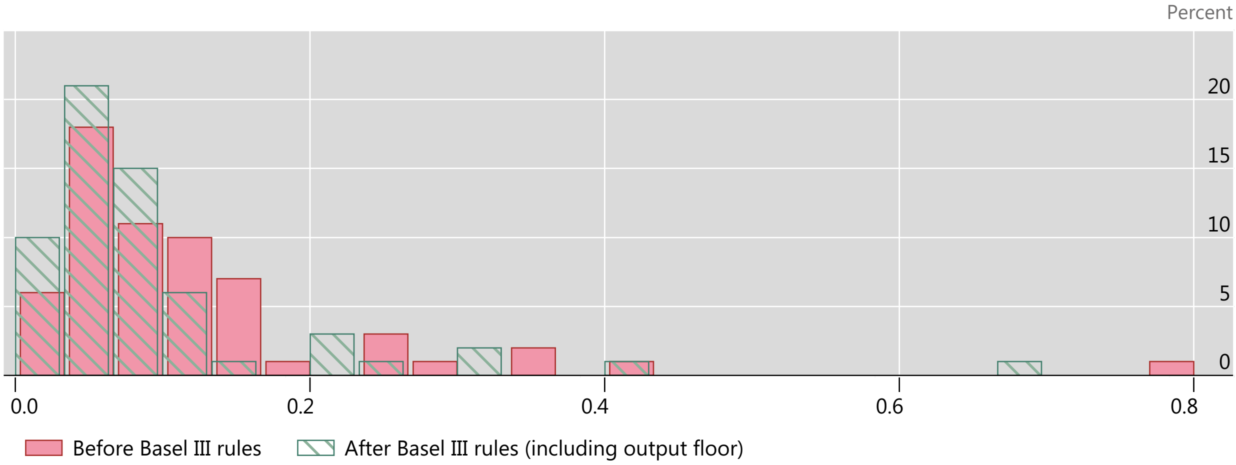

Note: Based on both IRB and SA credit risk exposures. Direct estimates of loss-absorbing resources (LAR) after the finalisation of Basel III are only available from 2018 to 2022. The minimum impact of finalisation reforms from 2018 to 2022, at the bank level, is applied retroactively to pre-2018 IRB exposures. LAR include minimum Pillar 1 requirements + capital conservation buffer + G-SIB buffer + countercyclical capital buffer.

Source: BCBS and Fitch.

We find that capital requirements are indeed conservative relative to cumulative losses. The AL-to-LAR ratios range between 0.6% and 80% for the banks in our sample, with a median of 8% – they are thus more conservative than the just-enough capital we derived in Section 4.2. This is unsurprising because the latter capital rests

32 RWA and EL are reported separately only under the IRB approach; the SA risk weights cover both expected and unexpected losses. The loss-absorbing resources used refer to non-defaulted exposures.

33 A key objective of the revisions to the Basel III framework finalised in December 2017 is to reduce excessive variability of risk-weighted assets (RWA) by constraining the use of internally modelled approaches and complementing inter alia the risk-weighted capital ratio with the output floor (see footnote 11). These reforms took effect from 1 January 2023 and will be phased in over five years. The supervisory data include banks' estimates of the impact of these reforms starting from 2018. For the years before 2018, we calculate the minimum impact of these reforms across the 2018–2022 period and we apply it retroactively to pre-2018 IRB exposures.

on perfect foresight, which is of course unavailable for real-life capital requirements. Furthermore, imposing the output floor brings in greater conservatism, with a median impact of about -2 percentage points.

5. EL vs AL: panel econometric analysis

We combine the cross-section and time dimensions in a panel regression analysis. In a first step, we derive aggregate measures of EL rates' accuracy as forecasts of AL rates. In a second step, we seek interpretations of the step-one forecast errors.

How accurate are banks' reported ELs as predictors of ALs?

We address this question on the basis of the following regression model:

$$AL\,rate_{i,t} = \alpha + \beta \cdot EL\,rate_{i,t-1} + e_{i,t} \tag{2}$$

where i is a bank and t is a year. The EL rate reported at t–1 is a real-time forecast of the AL rate that materialises over t. We do not include any other explanatory variables in this regression – including fixed effects – because our objective for now is *not* to explain AL rates but to evaluate banks' forecasts of AL rates.

The better EL rates are as predictors of AL rates, the closer would be the intercept to zero and the slope coefficient to one, and the higher would be the goodness-of-fit measure (adjusted R^2). Indeed, both a pooled OLS and a random effects specification deliver intercept estimates that are insignificantly different from zero and slope coefficient estimates that are significantly different from zero but not from one (Table 2). The adjusted R^2 indicates that banks' reported EL rates capture about one-fifth of the variation in AL rates. The explanatory power has mostly to do with variation across banks – the between R^2 (in the random-effects specification) is roughly 69%. This is consistent with our earlier findings that EL rates are quite successful in rank-ordering banks according to the respective AL rates but not when it comes to forecasting the evolution of these rates.

Expected credit losses as predictor of actual losses

End-2009 to end-2022 — Table 2

	(1)	(2)
Dependent variable:	AL rate$_t$	AL rate$_t$
Regression model:	Pooled OLS	Random effects
EL rate $_{t-1}$	1.003***	0.859***
Constant	4.81E-04	9.09E-04
Observations	643	643
No. of banks	65	65
Adjusted R^2	0.198	0.198
Between R^2		0.690
Within R^2		0.044

Note: Standard errors are clustered at the bank level. *** p<0.01, ** p<0.05, * p<0.1

What explains the forecast errors?

We saw above that EL rates leave unexplained a large share of AL rates' variation. We can provide interpretation of the underlying errors to the extent that regression residuals relate to variables observable at the time when EL rates are set. As candidate explanatory variables in a step-two regression, we choose bank-level characteristics, macro-economic variables and an indicator of potential inconsistencies between the datasets underpinning EL and AL rates (see next).

We define the step-one errors in the standard way:

$$e_{i,t} = AL\ rate_{i,t} - \widehat{AL\ rate}_{i,t}$$

where $\widehat{AL\ rate}_{i,t}$ is the fitted value in the OLS specification (equation (2)).[34]

As an indicator of dataset inconsistencies, we use the ratio of accumulated defaults (RAD) in the regulatory data.[35]

$$RAD_{i,t-1} = \frac{defaulted\ exposues_{i,t-1}}{total\ exposues_{i,t-1}}$$

where the numerator captures all exposures that have been identified as "in default" up to year t–1 but have not been written off and thus have not been dropped from the data in that year. $RAD_{i,t}$ reflects defaults over year t that should underpin $AL\ rate_{i,t}$. Its lag, $RAD_{i,t-1}$, could thus explain step-one errors, given serial correlation in AL rates – 54% across banks – that through-the-cycle EL rates may abstract from. Alternatively, a higher $RAD_{i,t-1}$ could stem from an accumulation of not fully resolved past defaults that can generate a new flow of losses (ie a revision of LGD) in year t. This flow of losses would be reflected in $AL\ rate_{i,t}$ but is not something that $EL\ rate_{i,t-1}$ – which refers only to non-defaulted exposures – is supposed to forecast.

In the second step, we estimate versions of the following panel equation that incorporate different subsets of regressors:

$$e_{i,t} = \alpha_i + \beta_1 \cdot AL\ rate_{i,t-1} + \beta_2 \cdot RAD_{i,t-1} + controls + \delta_t + u_i + \varepsilon_{i,t} \tag{3}$$

where δ_t and u_i are the time and bank fixed effects respectively, and $\varepsilon_{i,t}$ is the overall error term. Limited availability of AL rates reduces the number of observations in columns 1 and 3 from 643 (ie bank-level observations available for the entire time span) to 638, while limited coverage of other vendor data further reduces the number of observations in columns 4 to 7 to 463.

The first set of results (columns 1 to 3) do indicate that EL rates miss useful information stemming from serial correlation in AL rates. For one, lagged AL rates have explanatory power for step-one errors, as indicated by their statistically significant (positive) coefficient in column 1. Even though they generate a low goodness-of-fit along the time dimension (within R^2 just below 3%), they explain a large fraction of the dispersion of errors across banks (between R^2 at 64%). This

[34] We do not report step-two results based on the random-effect specification. Under such a specification, bank-specific bias in EL rates would result in a correlation between the regressor and the random effect, implying inconsistent estimates. That said, these results are very closely aligned to those reported in the text, based on pooled OLS.

[35] Again, we focus exclusively on exposures subject to the internal ratings-based (IRB) approach.

variable is also highly correlated with RAD (66%). But the latter brings in additional information, as adding it raises the within R^2 from 3% to 5% (column 3).

Step 2 regressions

Residuals from step 1 pooled OLS regression — Table 3

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
				step-one residual, e_t			
AL rate$_{t-1}$	0.366***		0.189*	0.146*	0.179**	-0.034	0.058
RAD$_{t-1}$		0.051***	0.036**	0.052***	0.053***	0.066***	0.072***
DepShare$_{t-1}$				-3.85e-05*	-2.68e-05	-1.34e-04*	-1.11e-04
RoA$_{t-1}$				2.47e-03**	2.67e-03***	2.26e-03**	3.13e-03**
PtB ratio$_{t-1}$				-2.33e-03***	-2.04e-03***	-2.04e-03	-1.71e-03
Debt service ratio$_{t-2}$				8.26e-05	7.27e-05	-5.40e-05	-5.70e-05
Credit-GDP gap$_{t-2}$				3.79e-05**	2.80e-05	5.36e-05**	3.99e-05**
GDP growth$_{t-2}$				1.46e-04*	-1.10e-04	2.13e-04**	-1.26e-05
Constant	-1.41e-03***	-1.26e-03***	-1.60e-03***	-3.00e-05	5.42e-03**	8.60e-03*	0.011*
Observations	638	643	638	463	463	463	463
No. of banks	65	65	65	49	49	49	49
Adjusted R^2	0.13	0.16	0.18	0.29	0.38	0.35	0.43
Within R^2	0.03	0.05	0.05	0.15	0.29	0.18	0.31
Between R^2	0.64	0.49	0.65	0.66	0.71	0.41	0.53
FE: time					✓		✓
FE: bank						✓	✓

Note: Standard errors are clustered at the bank level. *** p<0.01, ** p<0.05, * p<0.1

Do dataset discrepancies drive the forecast errors uncovered by the step-one regression? The dual interpretation of RAD – as discussed earlier – and its robust statistical significance are an indication that this may indeed be the case. Moreover, the coefficient of RAD is stable and maintains its significance across the entire set of richer specifications (columns 4 to 7). In these latter specifications, we introduce additional control variables that can offer forecast error interpretations related to: banks' business model, profitability, and market performance, or to macro-financial conditions that risk models may omit, or to aggregate unexpected shocks. Such economic interpretations would in turn suggest that the drivers of forecast errors go beyond data discrepancies.

The introduction of additional lagged explanatory variables brings in useful information for step-one errors (columns 4 to 7).[36] For one, these variables raise meaningfully the goodness of fit (adjusted R^2 rises from about 18% to 43%). This is mostly because they result in a better account of step-one errors' variability along the time dimension (within R^2 rising from 5% to 31%).[37]

[36] The decline in the number of banks included is due to: (i) some of them being not listed and thus lacking price-to-book ratios; (ii) gaps in the vendor data as regards bank-specific variables.

[37] The results are robust to modifications that control for bank size, country measures of concentration of the banking sector, and the proportion of exposures to specific asset classes vis-à-vis a banks' total credit exposures.

From the additional explanatory variables, three stand out with statistical and economic significance in the presence of bank and/or year fixed effects.[38] First, the positive coefficient of RoA is consistent with more profitable banks being more optimistic in their assessments of credit risk. A one-standard deviation change in this variable accounts for a 0.35 standard-deviation change in step-one errors. Second, a one-standard deviation increase in the price-to-book ratio accounts for a 0.17 standard-deviation decline in the step-one error, consistent with high-valuation banks being able to afford greater conservatism in their EL rates. Finally, a one standard-deviation change in the two-year lag of the credit-to-GDP gap amounts to roughly a 0.12 standard-deviation increase in the step-one error. This indicates that banks ignore valuable information in indicators of financial overheating.

In the appendix, we confirm that the above results are robust to excluding observations for which there are large measurable discrepancies between the underlying supervisory and accounting datasets.

6. Conclusion

Using confidential BCBS data, we obtain two key takeaways as regards banks' credit risk forecasts. First, banks´ stable credit risk estimates, consistent with post-GFC efforts to mitigate the sector's pro-cyclicality, do not account well for the evolution of actual losses. There is thus a strong case for a conservative mapping from risk weighted assets to capital requirements, as observed in practice. Second, banks provide an accurate picture of their credit portfolios' *relative* riskiness. This is consistent with post-GFC efforts to reduce practice-based variation in regulatory metrics, thus helping ensure that banks with riskier investments have a higher loss-absorbing capacity.

The explanations we provide for shortcomings in EL rates as forecasts of AL rates are of relevance for prudential authorities. The possibility that systematically optimistic EL rates increase the vulnerability of high-RoA banks seems worthy of investigation. In addition, authorities may need to assess whether their use of the counter-cyclical capital buffer (CCyB) – whose activation is at national discretion (BCBS (2011)) – has been compensating sufficiently for the failure of banks' EL rates to capture the evolution of AL rates. We find that this failure is at least partly due to banks' tendency to de-emphasise macro indicators of financial overheating, which are the indicators that the CCyB is supposed to draw on in order to ensure that banks build up loss-absorbing capital ahead of a spike in losses.

While our analysis is based on the best available cross-jurisdictional data for the largest internationally active banks, potential inconsistencies between the underlying sources hamstring the analysis. Notably, we cannot establish with reasonable certainty if some banks report biased forecasts, even though this is of key importance for financial stability. Collecting systematic supervisory data on the AL rates that banks' EL rates are supposed to forecast would thus be of tremendous value.

[38] Year fixed effects reveal that AL rates in 2020 (the Covid-19 outbreak) were significantly higher than in non-crisis years.

References

Abad J and J Suárez (2017): "Assessing the cyclical implications of IFRS 9 – a recursive model", *ESRB Occasional Paper Series* No 12.

Bank for International Settlements (BIS) (2023): *Annual Economic Report*.

Basel Committee on Banking Supervision (2005): *An explanatory note on the Basel II IRB risk weight functions*, July, www.bis.org/bcbs/irbriskweight.htm.

——— (2011): *Basel III: A global regulatory framework for more resilient banks and banking systems*, June, www.bis.org/publ/bcbs189.htm.

——— (2016): *Regulatory consistency assessment programme (RCAP) – Analysis of risk-weighted assets for credit risk in the banking book*, April, www.bis.org/bcbs/publ/d363.htm.

——— (2017): *Basel III: Finalising post-crisis reforms*, December, www.bis.org/bcbs/publ/d424.htm.

——— (2021): *Buffer usability and cyclicality in the Basel framework*, October, www.bis.org/bcbs/publ/d542.htm.

——— (2022): *Evaluation of the impact and efficacy of the Basel III reforms*, December, www.bis.org/bcbs/publ/d544.htm.

——— (2023a): *Basel framework*, 2023 version, www.bis.org/basel_framework/.

——— (2023b): *Basel III monitoring report February 2023*, www.bis.org/bcbs/publ/d546.htm.

Bernanke, B S (2009): "Lessons of the Financial Crisis for Banking Supervision", 7 May 2009 speech at the Federal Reserve Bank of Chicago Conference on Bank Structure and Competition, Chicago, Illinois.

Board of Governors of the Federal Reserve System (2023): "Agencies request comment on proposed rules to strengthen capital requirements for large banks", 27 July 2023.

Chae S, R Sarama, C Vojtech and J Wang (2018): "The impact of the current expected credit loss standard (CECL) on the timing and comparability of reserves", *Finance and Economics Discussion Series* 2018–020.

Claessens, S, G Dell'Ariccia, D Igan and L Laeven (2010): "Lessons and Policy Implications from the Global Financial Crisis", *IMF Working Paper 10/44*, February.

Covas F and W Nelson (2018): "Current expected credit loss: Lessons from 2007-2009", *BPI Staff Working Paper* 2018-1.

Goncharenko R and A Rauf (2020): "Loan loss provisioning requirements in a dynamic model of banking", mimeo.

Hardy, D C and C Schmieder (2013): "Rules of Thumb for Bank Solvency Stress Testing", IMF Working Paper no 2013/232, International Monetary Fund.

Herz, B and J Keller (2023): "How Do Regulators Set the Countercyclical Capital Buffer?", *International Journal of Central Banking*, vol 19, no 3, pp 99–137.

Juselius M and N Tarashev (2020): "Forecasting expected and unexpected losses", *BIS Working Papers*, no 913.

Lewrick, U, Schmieder, C, Sobrun, J and E Takáts (2021): "Releasing bank buffers to cushion the crisis - a quantitative assessment", BIS Bulletin no 11, May.

Krüger S, D Rösch and H Scheule (2018): "The impact of loan loss provisioning on bank capital requirements", *Journal of Financial Stability*, vol 36, pp 114–129.

Loudis B and B Ranish (2019): "CECL and the credit cycle", *Finance and Economics Discussion Series* 2019-061, Board of Governors of the Federal Reserve System.

Lu Y and V Nikolaev (2021): "Expected loan loss provisioning: an empirical model", *Chicago Booth Research Paper* No. 19-11.

Ong, Li Lian, C Schmieder and M Wei (2023): "Insights into credit loss rates: a global database", *BIS Working Papers,* no. 1101, May.

Appendix: Additional regression results

Filtered observations

We rerun the regression analysis after applying filters to remove observations for which the absolute difference of the denominators of *EL rate*$_{t-1}$ and *AL rate*$_{t-1}$ is 50% or more apart. Since the denominator of the AL rate comprises all credit risk exposures in the accounting data, we calculate this difference after augmenting the denominator of the EL rate with SA exposures. This filter removes one bank fully from the sample, as well as some observations for other banks. The total number of observations declines from 643 to 555.

The results for the step-one regression are in Table 2b. These are in line with those obtained using the unfiltered sample in Table 2, even if the EL coefficient declines marginally in value and significance.

Step-one regressions

Expected credit losses as predictor of actual losses

End-2009 to end-2022 – filtered sample — Table 2b

	(1)	(2)
Dependent variable:	AL rate$_t$	AL rate$_t$
Regression model:	Pooled OLS	Random effects
EL rate$_{t-1}$	0.877**	0.671**
Constant	8.47e-04	1.37e-03
Observations	555	555
No. of banks	64	64
Adjusted R^2	0.21	0.21
Between R^2		0.63
Within R^2		0.07

Note: Standard errors are clustered at the bank level. *** p<0.01, ** p<0.05, * p<0.1

Step-two regressions

The results for the step-two regressions are shown in Table 3b below. They are again in line with those obtained with the unfiltered sample (Table 3): robust significance of RAD, RoA and the PtB ratio (the latter only in columns 4 and 5 as for the unfiltered sample); as well as the credit-to-GDP; the other explanatory variables remain largely insignificant.

Step 2 regressions

Residuals from step 1 pooled OLS regression (filtered sample) — Table 3b

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	step-one residual, e_t						
AL rate$_{t-1}$	0.398***		0.171	0.178	0.215*	-0.059	-0.05
RAD$_{t-1}$		0.051***	0.038**	0.057***	0.056***	0.076***	0.016***
DepShare$_{t-1}$				-5.51e-05**	-4.76e-05*	-1.03e-04	-1.02e-04
RoA$_{t-1}$				3.36e-03***	3.39e-03***	2.70e-03**	3.05e-03***
PtB ratio$_{t-1}$				-3.01e-03***	-2.66e-03***	-2.15e-03	-1.52e-03
Debt service ratio$_{t-2}$				1.14e-04*	9.33e-05	-3.24e-04	-2.93e-04
Credit-GDP gap$_{t-2}$				3.48e-05**	3.03e-05*	8.35e-05***	6.97e-05***
GDP growth$_{t-2}$				8.33e-05	2.21e-04**	2.21e-04**	-3.80e-05
Constant	-1.59e-03***	-1.38e-03***	-1.67e-03***	1.59e-04	4.03e-03*	1.09e-02**	1.28e-02**
Observations	550	555	550	402	402	402	402
No. of banks	64	64	64	49	49	49	49
Adjusted R^2	0.20	0.25	0.27	0.39	0.47	0.46	0.52
Within R^2	0.02	0.10	0.07	0.20	0.32	0.24	0.36
Between R^2	0.74	0.39	0.56	0.71	0.75	0.38	0.46
FE: time					✓		✓
FE: bank						✓	✓

Note: Standard errors are clustered at the bank level. *** p<0.01, ** p<0.05, * p<0.1

Filtered observations and constant sample

We next rerun the regressions in Table 2b and Table 3b while dropping additional banks for which we do not have data on all the regressors to keep the sample of banks constant across all regression specifications. This reduces the sample to 49 banks, for a total of 402 observations. As in the main text, we work with an unbalanced sample.

Step-one regressions

Expected credit losses as predictor of actual losses

End-2009 to end-2022 – filtered and constant sample — Table 2c

	(1)	(2)
Dependent variable:	AL $rate_t$	AL $rate_t$
Regression model:	Pooled OLS	Random effects
EL $rate_{t-1}$	0.901**	0.710**
Constant	1.05e-03	1.56e-03
Observations	402	402
No. of banks	49	49
Adjusted R^2	0.21	0.21
Between R^2		0.68
Within R^2		0.07

Note: Standard errors are clustered at the bank level. *** p<0.01, ** p<0.05, * p<0.1

Step-two regressions

The results for the step-two regressions are shown in Table 3c below. The results are again in line with those obtained with the unfiltered sample (Table 3) and with the filtered sample (Table 3b).

Step 2 regressions

Residuals from step 1 pooled OLS regression (filtered and consistent sample) — Table 3c

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	step-one residual, e_t						
AL rate$_{t-1}$	0.408***		0.127	0.171	0.210*	-0.063	-0.007
RAD$_{t-1}$		0.054***	0.044**	0.057***	0.056***	0.076***	0.082***
DepShare$_{t-1}$				-5.47e-05**	-4.78e-05*	-1.01e-04	-1.02e-04
RoA$_{t-1}$				3.32e-03***	3.35e-03***	2.67e-03**	3.05e-03***
PtB ratio$_{t-1}$				-3.00e-03***	-2.66e-03***	-2.13e-03	-1.48e-03
Debt service ratio$_{t-2}$				1.14e-04*	9.38e-05	-3.33e-04*	-3.03e-04
Credit-GDP gap$_{t-2}$				3.48e-05**	3.07e-05*	8.38e-05***	7.01e-05***
GDP growth$_{t-2}$				8.57e-05	-1.34e-04	2.23e-04**	-3.50e-05
Constant	-1.82e-03***	-1.53e-03***	-1.80e-03***	1.11e-04	3.74e-03*	1.06e-02**	1.26e-02*
Observations	402	402	402	402	402	402	402
No. of banks	49	49	49	49	49	49	49
Adjusted R^2	0.21	0.29	0.30	0.39	0.46	0.45	0.52
Within R^2	0.02	0.12	0.09	0.20	0.32	0.24	0.36
Between R^2	0.77	0.48	0.61	0.70	0.74	0.38	0.45
FE: time					✓		✓
FE: bank						✓	✓

Note: Standard errors are clustered at the bank level. *** p<0.01, ** p<0.05, * p<0.1

Previous volumes in this series

1124 September 2023	The effect of monetary policy on inflation heterogeneity along the income distribution	Miguel Ampudia, Michael Ehrmann and Georg Strasser
1123 September 2023	Global supply chain interdependence and shock amplification – evidence from Covid lockdowns	Sally Chen, Eric Tsang and Leanne (Si Ying) Zhang
1122 September 2023	gingado: a machine learning library focused on economics and finance	Douglas K G Araujo
1121 September 2023	Margins, debt capacity, and systemic risk	Sirio Aramonte, Andreas Schrimpf and Hyun Song Shin
1120 September 2023	Energy Shocks as Keynesian Supply shocks: Implications for Fiscal Policy	Enisse Kharroubi and Frank Smets
1119 September 2023	Keep calm and bank on: panic-driven bank runs and the role of public communication	Olivier Coibion, Yuriy Gorodnichenko, Francesco Grigoli and Damiano Sandri
1118 August 2023	The origins of monetary policy disagreement: the role of supply and demand shocks	Carlos Madeira, João Madeira and Paulo Santos Monteiro
1117 August 2023	An impossibility theorem on truth-telling in fully decentralized systems	Rodney Garratt and Cyril Monnet
1116 August 2023	Absolute blockchain strength? Evidence from the ABS market in China	Jing Liu, Ilhyock Shim and Yanfeng Zheng
1115 August 2023	Sharks in the dark: quantifying HFT dark pool latency arbitrage	Matteo Aquilina, Sean Foley, Peter O'Neill and Thomas Ruf
1114 August 2023	The term structure of inflation forecasts disagreement and monetary policy transmission	Alessandro Barbera, Fan Dora Xia and Xingyu Sonya Zhu
1113 August 2023	To Lend or Not to Lend: the Bank of Japan's ETF purchase program and securities lending	Mitsuru Katagiri, Junnosuke Shino and Koji Takahashi
1112 July 2023	Trust bridges and money flows	Tobias Adrian, Rodney Garratt, Dong He, and Tommaso Mancini-Griffoli

All volumes are available on our website www.bis.org.