Public Disclosure Authorized
Public Disclosure Authorized
Public Disclosure Authorized
Public Disclosure Authorized

THE WORLD BANK
IBRD • IDA | WORLD BANK GROUP

A Model Case for Simplification in Europe — Internal Ratings, Better Calibrated

Laurent Millischer
August 2026

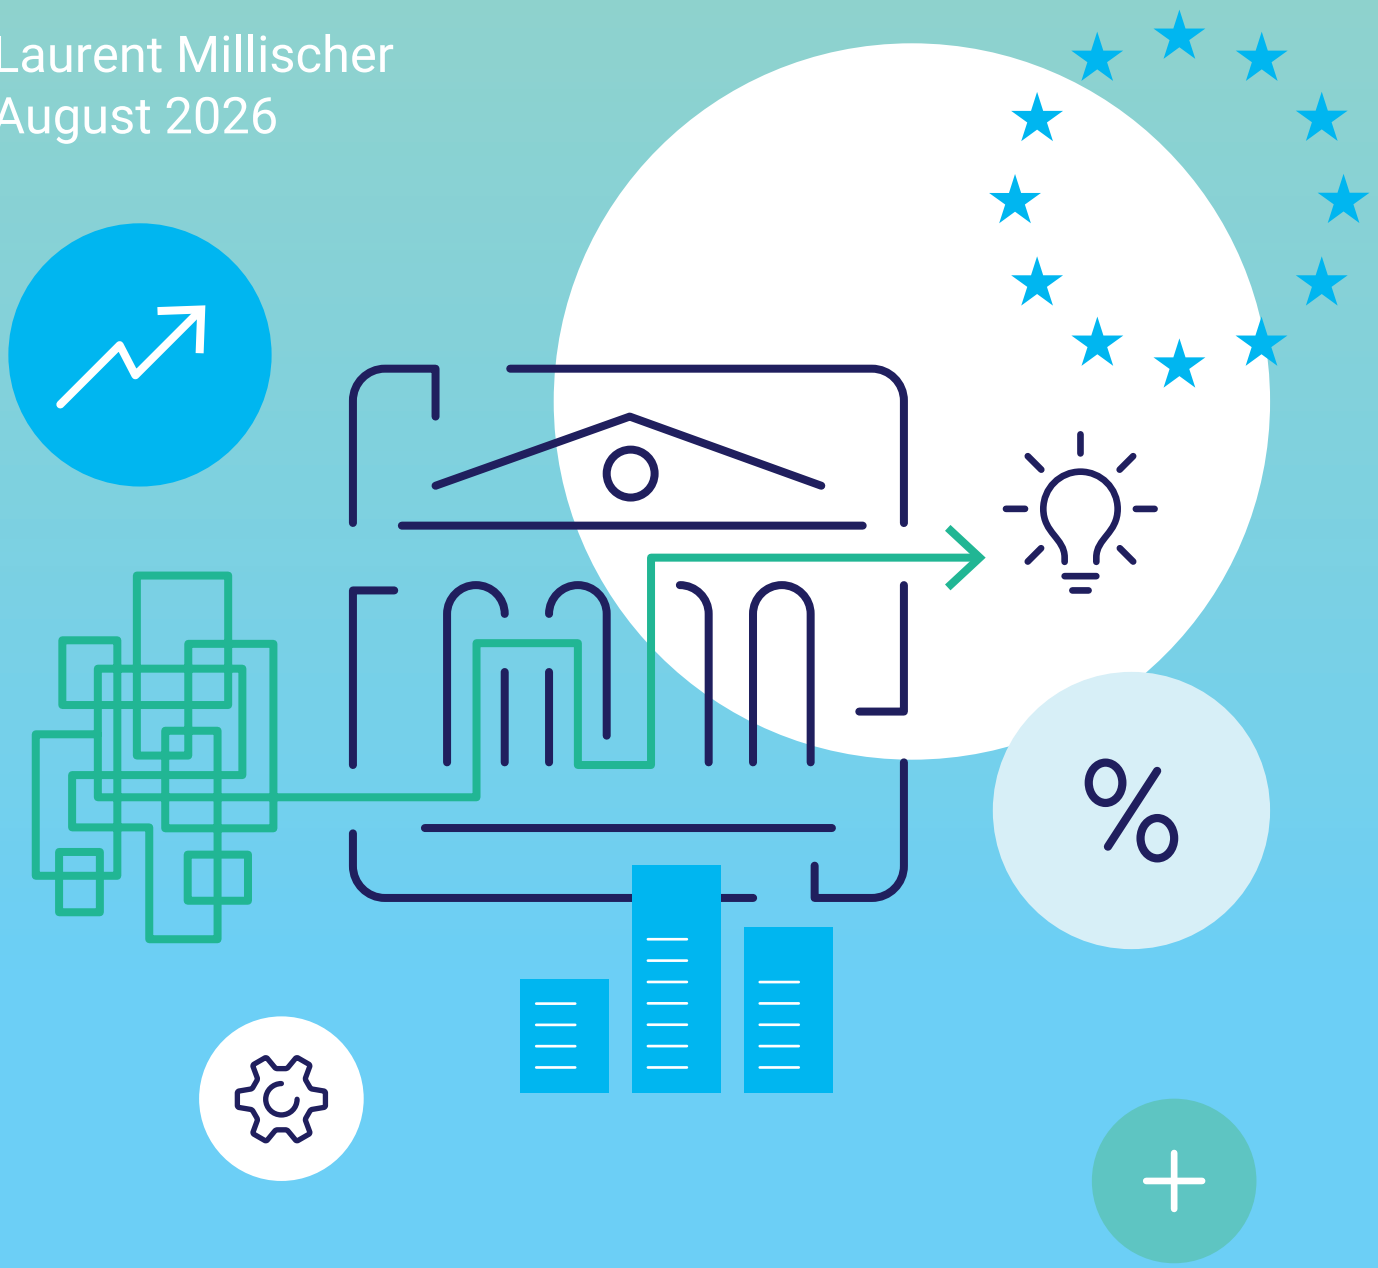

© 2026 International Bank for Reconstruction and Development / International Development Association or The World Bank

1818 H Street NW Washington DC 20433 Telephone: 202-473-1000 Internet: www.worldbank.org

This work is a product of The World Bank staff. The findings, interpretations, and conclusions expressed in this work do not necessarily reflect the views of The World Bank, its Board of Executive Directors, or the governments they represent. The World Bank does not guarantee the accuracy of the data included in this work. The boundaries, colors, denominations, and other information shown on any map in this work do not imply any judgment on the part of The World Bank concerning the legal status of any territory or the endorsement or acceptance of such boundaries.

Acknowledgments
I am grateful to Pamela Lintner, Vahe Vardanyan, Florian Resch, Max Fandl, Roberto Mosca, Nepomuk Dunz, Ezio Caruso, Ganbaatar Jambal, Marco Gross, Stefan Schmitz, Stefan Kavan, and the participants of the Austrian National Bank financial stability seminar for helpful comments. All remaining errors are my own. Contact: lmillischer@worldbank.org.

Rights and permissions
The material in this work is subject to copyright. Because The World Bank encourages dissemination of its knowledge, this work may be reproduced, in whole or in part, for noncommercial purposes as long as full attribution to this work is given.

Any queries on rights and licenses, including subsidiary rights, should be addressed to World Bank Publications, The World Bank Group, 1818 H Street NW, Washington, DC 20433, USA; fax: 202-522-2625; e-mail: pubrights@worldbank.org

Abstract

The Internal Ratings-Based (IRB) modeling framework, introduced under Basel II, was designed to align capital requirements with measured credit risk by allowing banks to estimate their own probabilities of default, loss given default, and exposure at default. Two decades of experience have revealed a structural flaw in this design: banks that estimate the parameters governing their own capital charges face a permanent incentive to underestimate risk, and the supervisory apparatus built to contain this incentive has grown into one of the most complex and costly layers of European banking regulation. This paper proposes an alternative in which the regulator calibrates the credit risk models based on banks' loan-level data. The proposal preserves the Basel risk-weight function, the 99.9 percent confidence level, and bank-specific through-the-cycle calibration, while removing the structural incentive to underestimate risk that has driven two decades of regulatory escalation. Against five alternative reform paths, ranging from the incremental simplification pursued by European institutions to a binding leverage ratio, the proposal is the only one that combines risk sensitivity, cross-bank comparability, and cost efficiency within the boundaries of the Basel framework. The resulting simplification is substantial, with minimal impact on the level and risk-sensitivity of capital requirements.

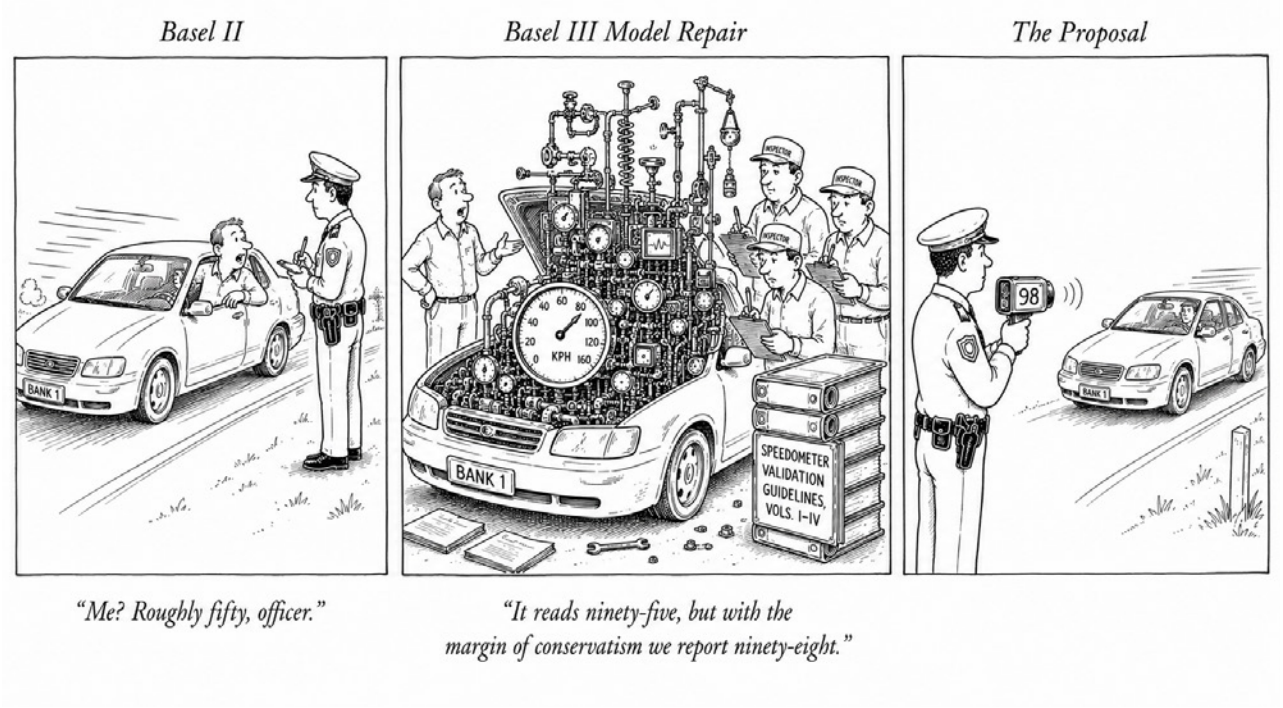

Source: ChatGPT, OpenAI, August, 2026, https://chat.openai.com/

Contents

Abstract 3
Contents 5
Executive Summary 6
Introduction 9
The Case for Phasing Out Bank-Developed Credit Risk Models 11
Benefits of IRB: Optimal Capital Allocation 11
Costs of IRB: An Inefficient Cat-and-mouse Game 12
Requirements for an Alternative 18
A Simple Approach for Risk-Sensitive Capital 21
Loss Given Default (LGD) 23
Credit Conversion Factors (CCF) 26
Probability of Default (PD) 27
Governance and Entry Conditions 35
Assessment 37
Assessment of the Proposal 37
Assessment of Alternative Approaches 41
Conclusion 45
References 48
Appendix 58

Executive Summary

i. **Regulatory simplification has become a central objective in European financial policy, yet one of the most complex and resource-intensive components of the prudential framework has hardly been examined: the Internal Ratings-Based (IRB) approach to credit risk.** The Draghi report, the European Commission's targeted consultation on banking competitiveness, and the ECB's High-Level Task Force on Simplification all identify regulatory burden as an impediment to growth and call for reduced complexity [Draghi, 2024, European Commission, 2026, European Central Bank, 2025d]. Given that credit risk weighted assets amount to EUR 8 trillion, approximately 80 percent of the total in EU banks [European Banking Authority, 2025c], the framework deserves scrutiny. This paper puts forward a proposal that aims to deliver comparable risk sensitivity to IRB models at a fraction of the current cost.

ii. **The IRB framework entrusts banks with estimating the very risk parameters that determine their capital charges, creating a permanent incentive to underestimate risk. Lower PD and LGD estimates translate mechanically into lower capital requirements and higher return on equity.** The academic literature confirms that banks have exploited this incentive in the past where permitted to do so, and that the degree of underestimation correlates with banks' capital positions [Behn et al., 2022, Mariathasan and Merrouche, 2014, Abbassi and Schmidt, 2018]. A back-of-envelope calculation illustrates the magnitude: a uniform reduction of just 10 basis points in estimated PDs saves approximately EUR 40 million per year for a mid-sized bank and nearly EUR 200 million for a large one.

iii. **The regulatory and supervisory apparatus built to contain this incentive absorbs vast resources yet still fails to deliver reliable and consistent risk measurement across banks.** The principal regulatory and supervisory texts governing the IRB framework in the EU exceed 1,000 pages, dedicated model inspection divisions operate at the ECB and at each national competent authority, and the aggregate annual cost to the banking industry and supervisors combined is estimated at EUR 0.7 to 1.2 billion. Despite this effort, less than half of cross-bank variation in estimated PDs can be explained by observed default rates, the remainder being driven by modeling choices [European Banking Authority, 2025a]. The resulting inconsistency distorts competition between banks using the IRB and the standardized approach.

iv. **Any alternative that achieves simplification of the IRB framework should satisfy five objectives:** (1) preserving risk sensitivity, i.e. ensuring capital requirements are higher for high-risk than for low-risk exposures, (2) producing comparable outcomes across banks for similar underlying risks, (3) achieving cost efficiency, (4) maintaining compatibility with the Basel framework, and (5) ensuring that capital requirements remain broadly stable through the economic cycle, avoiding procyclical swings that amplify the economic cycle.

v. **This paper proposes that the regulator calibrate credit risk models, using banks' loan-level data, while preserving the Basel risk-weight function, the 99.9 percent confidence level, and bank-specific through-the-cycle calibration.** Risk weights remain a function of borrower-level PD and exposure-level LGD, computed through the same equation as in the current Basel standard. What changes is not the formula but the way its inputs are determined: rather than requiring each bank to develop and maintain proprietary models, which require vast resources for regulation and supervision, the proposal assigns model calibration to the European Banking Authority (EBA). Banks contribute their loan-level data to the calibration. Banks should continue to develop models for their own risk management, now unconstrained by regulatory requirements and unbiased by the incentive to underestimate risk.

vi. **Bank-specific LGD models are replaced by a structural discounted-recovery formula; bank-specific credit conversion factors are replaced by supervisory flat rates.** For collateralized exposures, collateral values are reduced according to a supervisory haircut schedule and discounted by the bank's average time-to-foreclosure; for the uncollateralized portion, recovery is the bank's long-run average on fully unsecured exposures. A minimum LGD floor of 10 percent applies. Credit conversion factors (CCF) are calibrated by the EBA on pooled data by product type, consistent with the EBA's own assessment that CCF models exhibit low discriminatory power [European Banking Authority, 2026c].

vii. **For high-default portfolios, the EBA calibrates a logistic PD model every five years on pooled loan-level data, using five to ten standardized explanatory variables per model perimeter.** For portfolios such as residential mortgages, small and medium-sized enterprise (SME) lending, and consumer credit, for which many defaults are observed and statistical modeling is robust, each bank calibrates a single constant annually so that its portfolio-average PD matches its long-run average default rate. This preserves borrower-level risk differentiation while ensuring through-the-cycle stability. No bank-specific model development, validation, or supervisory model approval is required.

viii. **For low-default portfolios, where defaults are too rare for statistical calibration, PDs rely on external credit ratings where available; for unrated borrowers, the EBA calibrates a shadow rating model mapping firm fundamentals to default rates.** Low-default portfolios include large corporates, financial institutions, and sovereigns. Specialized lending exposures are treated under mandatory supervisory slotting. A minimum risk weight applies to all exposures, addressing calibration uncertainty at very low PDs where the risk-weight function is steepest and differences between adjacent rating grades are not statistically distinguishable.

ix. **The proposal meets all five objectives. Risk sensitivity is preserved: PDs are estimated from borrower-level data, LGDs are grounded in exposure-level collateral and recovery data, and the bank-specific constant ensures capital requirements reflect banks' default experience.** Cross-bank comparability improves as all banks apply the same model. Cost efficiency improves as model development, validation, and supervisory review are centralized. The proposal is compatible with the Basel agreement: the risk-weight function and the 99.9 percent confidence level are preserved in full and central risk parameters act as regulatory overrides. Through-the-cycle stability is embedded through the constant's calibration to the long-run average default rate.

x. **Alternative reform paths fail at least one of the five objectives. Incremental reform of the existing framework improves cost efficiency only at the margin.** Extending Foundation IRB (F-IRB) to all portfolios leaves bank-estimated PDs, the dominant driver of cost and risk-weight variation, in place. A full return to the standardized approach sacrifices risk sensitivity, while a granular standardized approach built on risk-driver tables comes closest to the proposal but offers coarser differentiation, more cyclicality, and weaker Basel compatibility. A binding leverage ratio abandons risk weighting altogether. The proposal is the only path that combines risk sensitivity, comparability, and cost efficiency within the Basel framework.

xi. **A natural roll-out would begin with low-default portfolios, where the statistical case against current models is strongest, complemented by a pilot on one high-default segment.** Banks on the standardized approach stand to gain access to more risk-sensitive capital requirements without building an IRB function. Current IRB banks face transition costs that are bounded in time and dwarfed by recurrent savings. The large body of IRB-related regulation, currently exceeding 1,000 pages, becomes largely superfluous.

Introduction

1. **Regulatory simplification has become a central objective in European financial policy.** The Draghi report on EU competitiveness identifies regulatory burden as a structural impediment to growth and calls for a systematic reduction in compliance costs across sectors [Draghi, 2024]. The European Commission has launched a targeted consultation on banking sector competitiveness that devotes an entire chapter to the complexity and effectiveness of the regulatory framework [European Commission, 2026], and the ECB's High-Level Task Force on Simplification has formulated recommendations to reduce undue complexity in the prudential, supervisory, and reporting framework [European Central Bank, 2025d]. Within this broader agenda, the EBA has proposed simplifications to specific areas of the credit risk framework, including margins of conservatism, downturn LGD estimation, and credit conversion factors [European Banking Authority, 2026c]. In prudential regulation, however, simplification often entails a trade-off: reducing regulatory requirements risks weakening the safeguards they were designed to provide. One area that has not yet been examined in this debate is among the most complex and resource-intensive components of the prudential framework: the Internal Ratings-Based (IRB) approach to credit risk. Given that credit risk weighted assets amount to EUR 8 trillion, approximately 80 percent of the total in EU banks [European Banking Authority, 2025c], the framework deserves scrutiny. This paper puts forward a proposal that aims to avoid the usual trade-off: comparable risk sensitivity with a fraction of the current cost.

2. **The IRB approach entrusts banks with estimating the very risk parameters that determine their capital charges.** Introduced under Basel II in 2004, the framework allows banks to estimate the probability of default (PD) and, under the Advanced (A-IRB) variant, the loss given default (LGD) and exposure at default (EAD).[1] More granular estimates of borrower risk translate into lower capital charges for safer exposures and higher charges for riskier ones, improving the efficiency of capital allocation. Because lower PDs and LGDs translate mechanically into lower capital requirements and higher return on equity, the framework creates a structural incentive to underestimate risk. The academic literature documents how banks have used this discretion historically. It has given rise to a vast regulatory and supervisory apparatus to constrain, review, and challenge banks' modeling choices, and to correspondingly large compliance costs for banks.

1 The Basel framework also permits internal models for market risk and counterparty credit risk, but these cover a much smaller share of total capital requirements and are not addressed in this paper.

3. **Two decades of supervisory effort have not resolved the tension between bank discretion and reliable risk measurement.** At the international level, the Basel Committee's 2017 introduction of input floors on PD and LGD estimates and an aggregate output floor reflects a collective judgment that bank-produced risk estimates cannot be trusted without hard constraints [Basel Committee on Banking Supervision, 2017]; US and Australian regulators have progressively constrained the scope and discretion of internal risk estimation along the same lines. In Europe, the response has taken the form of layering constraint upon constraint on the modeling process itself. The ECB's Targeted Review of Internal Models (TRIM), conducted between 2017 and 2021, found widespread deficiencies in banks' modeling practices [European Central Bank, 2021]. The EBA's own benchmarking exercises find that less than half of cross-bank variation in estimated PDs can be explained by observed default rates; the remainder reflects modeling choices, calibration horizons, and prudential adjustments [European Banking Authority, 2025a]. The EBA itself acknowledges that "the gained risk sensitivity is limited" while "the burden on modeling and the related supervisory review process is high" [European Banking Authority, 2026c]. The ECB now actively encourages banks to decommission models for portfolios that are difficult to model [European Central Bank, 2021], a significant shift from the CRR's original requirement to progressively roll out internal models across all material exposure classes.

4. **This paper argues that a limited set of regulator-calibrated credit risk models can deliver comparable risk sensitivity at a fraction of the current cost.** The proposal is designed to simplify the regulatory framework, not to increase the level of capital requirements; it preserves the Basel risk-weight function and the 99.9 percent confidence level in full. The first section makes the case for dropping bank-developed IRB models by quantifying the costs and documenting the incentive problem. The second section derives five objectives that any alternative should meet. The third section proposes such an alternative: a framework in which the regulator calibrates PD models centrally using pooled bank data, banks contribute their loan-level data and calibrate a single bank-specific parameter, and LGD is determined through a structural formula with supervisory haircuts. The fourth section assesses the proposal and five alternative reform paths against these objectives, finding that the proposal alone meets all five. The proposal is deliberately specific: vague calls for simplification are easy to agree with and difficult to act on. By laying out a concrete alternative, this paper aims to shift the debate from whether the IRB framework should be simplified to how.

The Case for Phasing Out Bank-Developed Credit Risk Models

Benefits of IRB: Optimal Capital Allocation

5. **The IRB framework was built on a compelling premise: that capital requirements calibrated to measured credit risk would allocate bank capital more efficiently than any flat regulatory schedule.** If a bank can measure the riskiness of each exposure with reasonable precision, it can hold capital in proportion to actual risk rather than applying uniform regulatory buckets. This matters enormously. Credit risk accounts for approximately 80 percent of total risk-weighted assets (RWA) in EU banks, representing some EUR 8 trillion, making even modest improvements in risk measurement consequential for capital allocation across the European banking system [European Banking Authority, 2025c]. In other words, risk-sensitive capital requirements are economically more efficient: for a given level of bank capital, more loans can be extended to the real economy without increasing the overall level of financial stability risk.

6. **This efficiency gain operates through two distinct channels: the reduced need for conservatism buffers and the concave shape of the risk-weight function.** The first, and more fundamental, channel operates through conservatism. When a regulator must assign a single risk weight to a broad category of exposures, as is the case in the Basel standardized approach, prudence demands that the weight be set conservatively enough to also cover riskier borrowers in the bucket. More granular risk measurement relaxes this constraint: if the regulator (or the bank) can distinguish risky from safe loans, each category can be calibrated more precisely, and the blanket conservatism that inflates capital requirements for safer borrowers falls away. In regulatory parlance, better measurement reduces the need for margins of conservatism, and hence lowers capital requirements overall. The second channel is mathematical. The regulatory IRB risk-weight function is concave in PD, therefore the average of the risk weights applied to individual borrower PDs is strictly lower than the risk weight evaluated at the pool's average PD. Differentiating borrowers therefore reduces aggregate capital requirements even in the absence of any conservatism adjustment.[2]

2 See Appendix A.1 for a formal statement and illustration.

7. **The resulting differences in capital requirements between IRB and standardized approach banks are substantial, even for identical exposure types.** Under the standardized approach, an unrated corporate exposure attracts a flat risk weight of 100 percent regardless of its creditworthiness [European Parliament and Council of the European Union, 2024]. An IRB bank, by contrast, can differentiate within this broad category: a high-quality corporate with a PD of 10 basis points and an LGD of 45 percent attracts a risk weight of roughly 30 percent, less than a third of the standardized charge. A lower-quality corporate with a PD of 2 percent and the same LGD would face a risk weight of approximately 115 percent. The standardized approach treats both identically. It is this granularity, the ability to distinguish good risks from bad within each asset class, that constitutes the genuine economic benefit of the IRB framework. Empirically, the combined effect is large: average risk-weight densities for credit portfolios treated under the IRB approach were roughly half those of comparable portfolios under the standardized approach before the Basel III output floor was introduced [Pérez Montes et al., 2018].

Costs of IRB: An Inefficient Cat-and-mouse Game

8. **The IRB framework has given rise to an extraordinarily complex regulatory apparatus that absorbs vast supervisory and industry resources, yet still fails to deliver consistent risk measurement across banks.** What was designed as a flexible, risk-sensitive system has become a bureaucratic labyrinth of model development, validation, supervisory review, and regulatory constraints. The EBA itself acknowledges that "the burden on modeling and the related supervisory review process is high" [European Banking Authority, 2026c]. This section argues that the root cause is a structural misalignment of incentives that no amount of regulatory patching can resolve.

9. **The fundamental problem is that banks have a strong financial incentive to underestimate credit risk parameters under the IRB framework.** Lower PD and LGD estimates translate mechanically into lower risk weights, lower capital requirements, and therefore higher return on equity. A simple back-of-envelope calculation[3] illustrates the magnitude of the incentive. Consider a bank with EUR 100 billion of credit risk RWA (a mid-sized European institution). A uniform reduction of just 10

3 The full calculation, including sensitivity to PD level and bank size, is presented in Appendix A.2.

basis points in estimated PDs, from 1.0 to 0.9 percent, reduces risk weights by approximately 3.5 percentage points, freeing some EUR 570 million of regulatory capital and saving roughly EUR 40 million per year in funding costs. For a large bank with EUR 500 billion of credit risk RWA, the annual saving from the same 10 basis point manipulation reaches nearly EUR 195 million, equivalent to approximately 2 percent of net income. Crucially, this incentive is strongest at low PDs (where a 10 basis point reduction cuts risk weights by up to 25 percent), precisely the segment where banks compete most aggressively for investment-grade business.

10. **The academic literature confirms that banks have historically exploited this incentive when permitted to do so.** Behn et al. [2022] show that IRB risk weights are associated with higher subsequent default rates than their level would imply, and that they predict defaults less accurately than standardized risk weights for the same exposures, indicating that banks' internal models introduce optimistic bias rather than superior risk measurement. Mariathasan and Merrouche [2014] find that the reduction in risk weights following IRB approval is disproportionately large among weakly capitalized institutions, precisely those with the strongest incentive to economize on capital, suggesting strategic underestimation rather than improved risk measurement. Abbassi and Schmidt [2018], using German supervisory data, demonstrate that banks with lower capital ratios report systematically lower PDs for the same borrowers compared to better-capitalized peers. The pattern is robust and well-documented: where discretion exists, banks use it to reduce capital charges. Calza et al. [2021] find that IRB ratings for corporate loans used as Eurosystem collateral are systematically less conservative than central banks' own assessments of the same borrowers, suggesting that banks also adjust risk estimates to maximize collateral eligibility. Recent regulatory reforms have closed many of these channels of optimization and raised IRB capital requirements.

11. **Regulators and supervisors have responded to this incentive problem by layering constraint upon constraint on the modeling process.** The response has taken two distinct forms. At the international level, the Basel Committee has progressively curtailed the scope of internal models through structural constraints: input floors on PD and LGD estimates and an aggregate output floor capping the capital benefit of internal models at 27.5 percent relative to the standardized approach [Basel Committee on Banking Supervision, 2017], as well as a leverage ratio that bypasses risk weighting altogether [Basel Committee on Banking

Supervision, 2014a]. These are not refinements of the modeling framework; they are expressions of fundamental distrust in bank-produced risk estimates. At the European level, the regulatory response has taken the form of an expanding body of secondary legislation and supervisory guidance aimed at constraining the modeling process itself. The principal regulatory and supervisory texts governing the IRB framework in the EU, from primary legislation through EBA technical standards and guidelines to the ECB's Guide to Internal Models, amount to over 1,000 pages.[4] Within this framework, instruments such as margins of conservatism, ex-ante notification requirements, and prescriptive estimation methodologies attempt to constrain the degrees of freedom available to bank modelers at each step of the process.

12. **The supervisory infrastructure needed to police this framework is substantial.** At the ECB, the Directorate General On-site and Internal Model Inspections is devoted to on-site supervision. A dedicated Internal Model Investigations Division comprises specialized credit risk sections whose sole function is to assess and challenge bank-developed IRB models [European Central Bank, 2026a]. In a normal year, the ECB launches 70–100 credit risk internal model investigations, each typically lasting several months, covering some 40 to 50 significant institutions [European Central Bank, 2025b, 2024, 2023, 2018]. The ECB's Targeted Review of Internal Models (TRIM), conducted between 2017 and 2021, involved 161 on-site investigations of credit risk models across 65 significant institutions, partly integrated into regular supervisory operations and partly carried out as dedicated TRIM missions [European Central Bank, 2021]. At the national level, each major competent authority maintains comparable structures: the Bundesbank operates a dedicated division for on-site inspections and internal risk models [Deutsche Bundesbank, 2023], the Banco de España maintains an Inspections and Internal Models department within its supervisory directorate [Banco de España, 2024], the French Prudential Supervision and Resolution Authority (ACPR) houses a dedicated Risk Modeling Control Unit within its on-site inspection delegation [Autorité de contrôle prudentiel et de résolution, 2025], and the Banca d'Italia coordinates internal model validation through its Directorate General for Financial Supervision and Regulation together with the Inspectorate [Banca d'Italia, 2025]. These dedicated inspection teams, moreover, are only part of the picture: horizontal methodology functions, internal model specialists in frontline supervisory teams, and regulatory staff all contribute to internal model oversight. Across the

4 A detailed inventory is provided in Appendix A.3. The total excludes Q&As, national supervisory guidance, and consultation papers.

EU as a whole, the total supervisory workforce dedicated to IRB model oversight likely exceeds five hundred. Even at this scale, the IMF's 2025 FSAP found internal model oversight structurally under-resourced and the ECB compelled to operate a triage system, with resource constraints flagged in two consecutive FSAP cycles [International Monetary Fund, 2025].

13. **On the bank side, the costs likely approach EUR 1 billion per year.** Each IRB institution must maintain dedicated teams for model development, independent validation, and internal audit of credit risk models, three separate functions mandated by regulation [European Banking Authority, 2016b]. In large European banks operating IRB models across multiple countries and asset classes (retail, corporate, SME, sovereign, bank, project finance), the dedicated IRB workforce, spanning model development, validation, monitoring, and governance, typically numbers between 60 and 80 full-time equivalents (FTEs), with some of the largest cross-border groups exceeding 100.[5] This internal effort is supplemented by a substantial external consulting industry: every major advisory firm (Oliver Wyman, McKinsey, KPMG, Deloitte, EY, PwC, Accenture) maintains dedicated IRB and credit risk modeling practices in Europe, and specialist firms such as d-fine, Finalyse, and Prometeia have built their business models around IRB compliance.[6] The aggregate cost to the European banking industry, in staff, consulting fees, and management attention, likely falls in the range of EUR 0.7 to 1.2 billion annually.[7] These costs, troubleshooted, may none the less be rational for individual banks: as Appendix A.2 shows, even a modest reduction in reported PDs can save tens of millions of euros per year in capital costs for a single institution. The question is whether the same degree of risk sensitivity could be achieved through a less costly approach.

14. **Beyond the direct financial costs, the IRB framework imposes a significant opportunity cost on the allocation of quantitative talent.** Banks, supervisors, and consulting firms alike employ large teams of quantitative analysts whose primary task is to develop, validate, or review models whose outputs are then constrained by regulatory floors, margins of conservatism, and supervisory overrides. This matters in a broader context. The EU's productivity gap with the United States

5 Estimated by scaling the per-asset-class benchmarks reported in PricewaterhouseCoopers [2022] to the full scope of a large multi-country IRB bank; see Appendix A.5 for the derivation by bank tier.

6 Appendix A.4 provides a list of advisory firms with dedicated IRB and credit risk modeling practices in Europe, with references to the corresponding service pages.

7 A derivation is provided in Appendix A.5.

has widened substantially over the past two decades, driven in large part by Europe's failure to translate technological innovation into economic growth [Schnabel, 2024]. As of 2024, no EU company with a market capitalization above EUR 100 billion had been created from scratch in the last fifty years, and the EU trails both the US and China in private investment in advanced technologies [Draghi, 2024]. In this environment, tying up many thousands of quantitative specialists in a regulatory modeling exercise of diminishing marginal returns contributes to a broader pattern of misallocation of scarce human capital that Europe can ill afford.

15. **Despite this vast apparatus, the IRB framework fails to deliver consistent risk measurement across banks.** The EBA's annual benchmarking exercises document persistent and wide cross-bank dispersion in PD estimates, of which less than half can be explained by observed default rates; the rest reflects modeling choices, calibration horizons, and prudential adjustments [European Banking Authority, 2023a, 2024, 2025a]. For high-default portfolios where sufficient data exist, the regulatory repair effort has narrowed the scope for systematic bias: on aggregate, estimated PDs now exceed observed five-year default rates across all major asset classes [European Banking Authority, 2025a]. But for low-default portfolios, including large corporates, institutions, sovereigns, and specialized lending, defaults are too rare to validate the models' discriminatory power with statistical confidence, and cross-bank variability is largest precisely in this segment. This variability is not merely noise. Studies that use credit register data to compare PD estimates assigned by different banks to the same borrower confirm large and systematic cross-bank differences [Berg and Koziol, 2017, Abbassi and Schmidt, 2018, Plosser and Santos, 2018], and show that the differences correlate with banks' capital positions and exposure sizes: institutions with lower capital ratios assign more favorable ratings to the borrowers where they hold the largest exposures. In particular, for low-default portfolios where supervisory data cannot detect bias, the incentive to underestimate persists. The system, in short, absorbs enormous resources to produce an outcome that still falls short of its fundamental promise: comparable, risk-sensitive capital requirements.

16. **A claimed side benefit of bank-developed models, their use in banks' own risk management, is undermined by the models' regulatory design.** The framework requires that internal estimates play an essential role in credit approval, limit setting, pricing, and internal governance, the so-called use test [European Parliament and Council of the European Union, 2013b, Article 144(1)(b)]. Yet the estimates the framework produces are ill-suited for these purposes by construction. Regulatory PDs are calibrated through the cycle, the right property for stable capital requirements but the wrong one for risk management: a through-the-cycle PD does not estimate the probability that a borrower defaults over the next twelve months under current conditions, which is what underwriting, monitoring, and pricing require. Margins of conservatism and prescribed modeling choices [European Banking Authority, 2017] compound this, making estimates more conservative but not more discriminating. It is telling that banks maintain parallel point-in-time model suites for IFRS 9 provisioning, which demands forward-looking estimates that the regulatory models cannot supply. Whether the regulatory parameters genuinely drive credit decisions is, moreover, intrinsically difficult to verify: documentation and committee minutes can be reviewed, but overrides and parallel decision processes are not observable from the outside. A model calibrated for capital stability does not become a management tool by regulatory command.

17. **The inconsistency of IRB risk estimates is not only a prudential concern but also distorts competition.** A bank that achieves lower PDs for the same borrower, whether through genuine risk insight or through more aggressive calibration choices, holds less capital and can price loans more competitively. All else equal, a bank that calibrates its models more aggressively can offer tighter pricing and capture business, creating a competitive incentive to minimize reported risk. This dynamic is strongest at the low-PD end of the spectrum, precisely the investment-grade segment where banks compete most intensely on pricing. The distortion extends beyond IRB banks. Because the standardized approach assigns a flat 100 percent risk weight to unrated corporate exposures regardless of creditworthiness, standardized approach banks cannot compete for safer borrowers on comparable terms: an IRB bank facing a risk weight of 30 percent for a high-quality corporate can offer pricing that a standardized approach bank holding more than three times the capital simply cannot match. The result is a two-tier system in which access to lower capital requirements depends not on the underlying risk of the lending but on whether the bank has the resources to build and maintain an IRB apparatus.

Requirements for an Alternative

18. **Before proposing an alternative, this section establishes the objectives any simplification of the IRB framework should meet.** A growing institutional consensus holds that the framework needs simplification, though existing reform proposals remain within its current boundaries. Any alternative should satisfy five objectives: preserving risk sensitivity, producing comparable outcomes across banks, achieving cost efficiency, maintaining compatibility with the Basel framework, and ensuring through-the-cycle stability of capital requirements.

19. **There is a growing institutional consensus that the framework needs simplification.** The case was made over a decade ago: in the speech "The dog and the frisbee" Haldane and Madouros [2012] argued that complex, model-based capital regulation was not merely costly but counterproductive. The ECB's own position has shifted over the past decade: whereas the CRR requires IRB banks to progressively roll out internal models across all material exposure classes [European Parliament and Council of the European Union, 2013b, Article 148], after TRIM the ECB now actively encourages institutions to simplify their model landscapes and supports the decommissioning of models for portfolios that are difficult to model [European Central Bank, 2021]. The EBA's 2026 Discussion Paper on simplification of the credit risk framework takes a further step, proposing optional fallback approaches for several IRB estimation requirements, including margins of conservatism, downturn LGD, in-default LGD, and credit conversion factors [European Banking Authority, 2026c, Section 3.2]. In March 2026, the EBA and the ECB each took steps to streamline the approval process for IRB model changes [European Banking Authority, 2026b, European Central Bank, 2026b], a welcome reduction in administrative friction that nonetheless remains within the existing paradigm. In its response to the Commission's competitiveness consultation, the Eurosystem went further, explicitly encouraging banks to confine internal models to strategic portfolios and to switch to the standardized approach elsewhere [Eurosystem, 2026]. The European Banking Federation has called the IRB approval process "very burdensome and lengthy for both banks and supervisory teams" [European Banking Federation, 2025]. All of these initiatives, however, remain within the

boundaries of the current framework. This paper argues that a more fundamental rethinking is both possible and necessary. The following paragraphs set out the objectives against which any such alternative should be assessed.

20. **Risk sensitivity: capital charges must reflect the underlying risk of each exposure.** This is the genuine economic contribution of the IRB framework, and the reason it was introduced. When capital charges reflect the actual creditworthiness of each borrower, banks can extend more credit for a given level of loss-absorbing capital without increasing financial stability risk. Any alternative that sacrifices this property, would impose some economic cost. An alternative meets this objective when the resulting risk weights reflect genuine differences in default risk that can be statistically demonstrated against observed outcomes.

21. **Cross-bank comparability: two institutions holding similar underlying portfolios should face similar capital charges.** While true PDs and LGDs are unobservable, some cross-bank variation in estimated risk parameters is inherent to modeling under uncertainty and does not, by itself, indicate model failure. For high-default portfolios, this variability is manageable: observed defaults are frequent enough to backtest each bank's estimates against realized outcomes, providing a disciplining mechanism that bounds how far estimates can drift from one another. For low-default portfolios, no such discipline exists. Defaults are too rare to validate models statistically, cross-bank differences must be taken at face value, and the structural incentive to underestimate documented above cannot be controlled ex post. It is in this segment that comparability matters most, and where the current framework falls short: cross-bank variability is largest precisely where it is least verifiable [European Banking Authority, 2025a]. An alternative meets this objective when cross-bank variation can be traced to observable differences in risk rather than unverifiable modeling choices.

22. **Cost efficiency: the resources devoted to the framework should be proportionate to the risk sensitivity they deliver.** The current IRB apparatus absorbs an estimated EUR 0.7 to 1.2 billion per year in bank-side compliance costs and supervisory resources, and ties up thousands of quantitative specialists in a regulatory modeling exercise of diminishing marginal returns. An alternative meets this objective when it delivers comparable risk sensitivity with substantially lower aggregate resource consumption, freeing supervisory capacity and bank budgets for more productive uses.

23. **Compatibility with the Basel framework.** The EU's prudential standards derive their credibility in part from their consistency with internationally agreed norms. Any departure from the Basel standard invites scrutiny under the Basel Committee's Regulatory Consistency Assessment Program and risks undermining market confidence in the comparability of EU banks' capital ratios. An alternative meets this objective when it can be implemented within the existing Basel architecture, preserving the core risk-weight function and the minimum standards for risk parameter estimation, without requiring renegotiation of the international framework.

24. **Through-the-cycle stability: capital requirements should remain broadly stable across the economic cycle.** Procyclical capital requirements amplify the business cycle by forcing banks to raise capital or cut lending precisely when the economy is weakening. The Basel framework addresses this by requiring IRB risk parameter estimates to reflect long-run averages and downturn conditions, and by assigning largely flat risk weights under the standardized approach. An alternative meets this objective when capital requirements remain broadly stable as economic conditions change.

A Simple Approach for Risk-Sensitive Capital

25. **The approach proposed in this paper preserves the core architecture of the IRB framework: risk weights remain a function of borrower-level probability of default and exposure-level loss given default.** The entire risk-weight (RW) calculation continues to rest on a single equation, unchanged from the current Basel standard:

$$RW = 12.5 \times f(PD) \times LGD \qquad (1)$$

where f (PD) is the Vasicek function that maps each borrower's through-the-cycle probability of default into a stressed default probability at the 99.9th percentile confidence level, incorporating an asset correlation parameter and a maturity adjustment (see Appendix A.1 for the full specification). What changes in this proposal is not the formula but the way its inputs are determined. Rather than requiring each bank to develop, validate, and maintain proprietary statistical models for PD, LGD, and EAD, the proposal assigns these tasks to the regulator. This follows the same institutional logic as the standardized approach, where risk-weights are set centrally and banks apply them. In the European context, model calibration would be carried out by the EBA and adopted by the Commission through binding regulatory technical standards. The three components are described below: LGD, where a structural discounted-recovery approach replaces proprietary regression models; EAD, where standardized supervisory credit conversion factors replace bank-specific estimates; and PD, where regulator-calibrated models with a bank-specific constant preserve borrower-level risk differentiation. Before describing these components, two structural decisions that apply across all three must be established: the classification of portfolios into high-default and low-default segments, and the definition of model perimeters.

26. **The first cross-cutting structural decision is the classification of portfolios into high-default and low-default segments.**[8] High-default portfolios are those where defaults are frequent enough to support meaningful statistical calibration at the individual bank level, or at minimum across a pool of banks. Low-default portfolios are those where defaults are too rare for this purpose, whether because default rates are low or because the portfolio consists of few, large exposures, so that risk parameters must be inferred from external sources. This distinction governs the estimation approach for both PD and LGD: for high-default portfolios, parameters can be calibrated from observed default and loss data; for low-default portfolios, they must rely on supervisory flat rates, external benchmarks, or centrally pooled data.[9] Under the proposal, the classification would be formalized and defined centrally by the EBA. In broad terms, high-default portfolios would include residential mortgages, qualifying revolving retail exposures, other retail and consumer lending, SME, and mid-market corporates; low default portfolios would cover large corporates,[10] financial institutions, sovereigns and public-sector entities, and specialized lending.

27. **The second structural decision is the definition of model perimeters within each segment.** A model perimeter is the level at which a common estimation approach is defined. The minimum segmentation follows CRR exposure classes (Article 147), but finer splits may be needed where sub-segments exhibit materially different default or recovery behavior, for example by geography, product type, or borrower characteristics such as fixed versus floating rate loans. What will typically drive the segmentation is the explanatory power of the PD model: where a single model cannot adequately capture the default dynamics of a sub-segment, a separate perimeter is warranted, and the LGD parameters are computed accordingly. CRR Article 143 requires risk differentiation within each exposure class. The proposal operationalizes this by making the perimeter definition a central regulatory decision rather than a bank by-bank choice.

---

8 Neither the CRR nor the Basel framework formally defines "high-default" and "low-default" portfolios. The ECB's Targeted Review of Internal Models used the terms as an organizational grouping without formal definition [European Central Bank, 2021, Section 3.2.3, fn. 36]. The EBA supervisory benchmarking exercise classifies portfolios by exposure class (corporates, institutions, and specialized lending as "low-default"; retail and SME as "high-default"), which serves as a rough proxy but does not correspond to a regulatory threshold.

9 The CRR already recognizes this distinction operationally: Article 185(c) requires institutions to carry out "external benchmarking" for portfolios where internal default data are sparse, implicitly acknowledging that some exposure classes cannot be validated on the basis of internal experience alone.

10 The boundary between mid-market and large corporates would likely be defined by a metric of firm size such as consolidated revenue. See Figure 1 for an illustration of how the PD approaches this paper puts forward map along the dimensions of firm size and riskiness.

Loss Given Default (LGD)

28. **This paper proposes to replace bank-specific LGD models with a structural approach built on a small number of observable parameters.** For each exposure, a recovery value is computed separately for the collateralized and uncollateralized portions, and the LGD is derived from the aggregate, subject to a floor. The collateral recovery is obtained from a simple discounted-recovery formula: the recorded collateral value, reduced by a supervisory haircut set EU-wide by collateral type, is discounted to the date of default using the bank's average time-to-foreclosure. The uncollateralized recovery rate is the bank's long-run average observed recovery on fully unsecured exposures in the same portfolio. Guarantees are recognized through the existing substitution approach. The proposal retains meaningful risk differentiation through collateral type, coverage ratio, foreclosure timelines, discount rates, and recovery experience, while eliminating the proprietary modeling, validation, and supervisory oversight that drive cost and cross-bank inconsistency. The detailed equations are provided in Appendix A.6.

29. **The fundamental drivers of LGD are few and observable, which allows for a significantly simpler modeling approach.** Empirical research consistently finds that realized loss rates on defaulted exposures are dominated by a small number of loan-level characteristics [Schuermann, 2004, Khieu et al., 2012, Grunert and Weber, 2009]. For the collateralized portion, recovery depends on the type and value of collateral pledged and on the creditor's ability to seize and liquidate it, which is overwhelmingly determined by the efficiency of the judicial system [Davydenko and Franks, 2008, Djankov et al., 2008]. For the uncollateralized portion, recovery is driven primarily by the seniority of the claim and the efficiency of the bank's workout process, with economic conditions at the time of default playing a secondary role. Yet under the current framework, banks retain wide discretion over LGD estimation: segmentation variables and boundaries, treatment of incomplete workouts, definition of cure, and downturn calibration. The EBA's benchmarking exercises have found that significant cross-bank differences in LGD estimates reflect modeling choices rather than differences in actual risk [European Banking Authority, 2025a]. Each choice must be documented, validated, and reviewed. Because the underlying drivers are few and slow-moving, meaningful risk differentiation can still be achieved with a small number of prescribed parameters drawn mostly from the bank's own data, eliminating the modeling discretion and much of the supervisory apparatus built to control it.

30. **For the collateralized portion of each exposure, the recovery rate is computed using a structural discounted-recovery model with two key inputs: a supervisory haircut and a bank-specific time-to-foreclosure.** The collateral value recorded for the specific exposure is reduced by a supervisory haircut and then discounted to the date of default using the bank's average time-to-foreclosure for that collateral class, yielding a recovery rate on the secured portion.[11] The haircuts are set as EU-wide supervisory flat rates by collateral type, reflecting the gap between the recorded value and what is actually realized at forced liquidation (fire-sale discounts, transaction costs, appraisal optimism). They are calibrated by the EBA using pooled loss data and reviewed every five years. The time-to-foreclosure is computed as the bank's own average time from default to collateral realization over a rolling window of seven to ten years, differentiated by collateral type and, for banks operating in several countries, by the jurisdiction in which the collateral is located. It is observable, auditable, and moves slowly over time, reflecting the efficiency of the judicial environment in which the collateral is realized. All collateral categories currently eligible under the IRB collateral framework remain eligible, each with its own haircut schedule and bank-specific time-to-foreclosure.[12] This eliminates the A-IRB-specific modeling layer (bank-estimated haircuts, collateral-borrower correlation adjustments, downturn LGD calibration) [European Banking Authority, 2020]. In practice, the gap between recorded collateral values and realized recovery prices may vary across markets and where differences are material, the haircut schedule can be differentiated by groups of jurisdictions rather than set at a single EU-wide level.

31. **For the uncollateralized portion, the recovery rate is the bank's long-run average observed recovery on fully unsecured exposures or from regulatory rates where default data is unavailable.** Using fully unsecured loans avoids a double counting problem: borrowers who have pledged collateral tend to exhibit higher non-collateral recovery rates (because they have more at stake), so mixing collateralized and uncollateralized exposures would overstate the unsecured recovery parameter and, combined with the separately computed collateral recovery, over-

11 Existing CRR valuation rules are retained. Collateral must be revalued at the minimum frequencies prescribed by Article 208 (residential real estate at least every three years, commercial real estate annually, financial collateral daily), and the conservatism caps on upward revaluation under Article 229 continue to apply (property values may not be adjusted upward beyond the historical average over the preceding six years for residential or eight years for commercial real estate).

12 The IRB-eligible collateral universe comprises financial collateral, residential and commercial real estate, receivables linked to commercial transactions, other physical collateral, leases, and life insurance policies (CRR Articles 197 to 200, with the additional IRB-specific categories set out in Article 199). The eligibility and legal enforceability conditions for collateral recognition (CRR Articles 192 to 217) continue to apply. The detailed list of eligible collateral types, the conditions attached to each, and the calibration of the corresponding supervisory haircut schedules are outside the scope of this paper.

state total recovery [Acharya et al., 2007]. Restricting the sample to fully unsecured loans produces a conservative estimate, which is appropriate in a prudential framework. Because the bank retains an incentive to overstate recorded recoveries, the underlying workout data are subject to standard supervisory review. The parameter is computed over a minimum of 10 years at the portfolio level at which the bank operates, using only closed recovery processes: each portfolio has a defined maximum recovery period, after which the case is treated as final, and defaults less than two years old are excluded.[13] This avoids the need to forecast future recoveries on open cases, which is one of the most judgmental and model-dependent steps in current A-IRB LGD estimation. A bank's workout strategy affects both actual losses and the reported recovery rate; a more efficient workout process benefits the bank through lower LGD, which is the intended prudential incentive. The remaining supervisory task is correspondingly narrower than under current A-IRB: verifying the consistency of the bank's recovery data and the application of the prescribed maximum recovery periods, rather than reviewing an entire LGD model. Where the bank's loss history is shorter than 20 years, the competent authority applies a conservatism adjustment to the long-run average recovery rate, mirroring the treatment of the long-run average default rate. For low-default portfolios, where individual banks will not have enough resolved defaults to compute a robust long-run average, the EBA sets supervisory flat recovery rates, calibrated using pooled loss data across banks and published rating agency recovery studies. Where a bank's own unsecured default history for a given portfolio is insufficient or unavailable, a conservative default value applies, set below the pooled average to reflect the additional uncertainty.

32. **The discount rate reflects the opportunity cost of tied-up capital. It enters the LGD calculation at two points: discounting the collateral value to the date of default and converting the cash flows recovered on unsecured exposures into present-value terms.** It is set as the higher of two rates: (1) the loan's effective interest rate as defined under IFRS 9 (incorporating fees, commissions, and charges into a single annualized rate) plus any contractual default interest surcharge, and (2) a supervisory floor rate. The loan-specific rate captures the full contractual claim the bank foregoes while waiting for collateral realization, consistent with the approach used in empirical recovery studies [Khieu et al., 2012]. The supervisory floor is a single EU-wide rate set as the euro short-term rate (€STR) plus a fixed

13 This is consistent with the treatment of incomplete workouts under EBA/GL/2017/16, paragraph 156, which requires institutions to define a maximum recovery period by exposure type. The ECB's Guide to Internal Models further allows institutions to establish a minimum observation period before including defaults in LGD calibration samples.

add-on, calibrated to approximate the average cost of bank funding. The add-on is set by the EBA and reviewed every three to five years, anchored to observed average bank funding costs. The "higher of" construction ensures a minimum discount for all exposures while capturing the genuinely higher opportunity cost on high-rate loans, and avoids rewarding banks for underpricing risk.

33. **A minimum LGD floor of 10 percent applies to all exposures regardless of collateral coverage.** This ensures that capital requirements never fall to zero even for exposures that appear fully secured, reflecting residual uncertainty in collateral realization, legal risks, and operational frictions in the workout process. The 10 percent level is anchored to the Basel III A-IRB input floor for exposures secured by residential or commercial real estate [Basel Committee on Banking Supervision, 2017], the lowest floor for any non-financial collateral type. Since the structural formula already differentiates risk through collateral-specific haircuts, time-to-foreclosure, and discount rates, the floor serves purely as a backstop for risks the formula cannot capture.

Credit Conversion Factors (CCF)

34. **Bank-specific credit conversion factor models have low discriminatory power, and the regulatory framework already limits both the risk sensitivity and the capital benefit they can deliver.** Under the Basel III input floors, a bank's own CCF estimate may not fall below 50 percent of the corresponding standardized approach CCF. In practice, this means that even a bank with a demonstrably better model can reduce its capital requirement on a given off-balance-sheet item by at most half relative to the standardized treatment. The EBA's own Discussion Paper acknowledges that CCF models "exhibit low discriminatory power, raising validity questions" and that "the burden of CCF modeling may not be proportional to the benefits" [European Banking Authority, 2026c, paragraphs 54–55]. The Basel Committee's 2016 consultation on reducing variation in credit RWA identifies CCF dispersion as a material contributor to cross-bank RWA variation [Basel Committee on Banking Supervision, 2016]. A structural reason for this poor performance is that credit conversion factors tend to cluster at extreme values (zero or full drawdown), with little mass in between, making them unsuitable for standard regression-based modeling [Tong et al., 2016, Taplin et al., 2007]. The EBA's first comprehensive guidelines on CCF estimation were published only in July 2025, running to 114 pages [European Banking Authority, 2025b], nearly two

decades after Basel II introduced the IRB framework. The delay itself suggests that defining sound CCF estimation methodology is an inherently difficult task, reinforcing the case for supervisory flat rates.

35. **This paper proposes that the EBA calibrate supervisory CCFs by product type, using pooled bank data, replacing bank-specific estimation entirely.** Off-balance sheet commitments represent roughly a quarter of total credit exposures for EU significant institutions [European Central Bank, 2025a], but after applying credit conversion factors (most of which range between 20 and 50 percent) the resulting contribution to RWA is considerably smaller, limiting the capital impact of any refinement in CCF estimation. The EBA's Discussion Paper itself proposes allowing a fixed CCF for entire exposure types as an optional fallback [European Banking Authority, 2026c].[14] The approach proposed here makes this the standard treatment rather than an option. The CCFs would be differentiated by product type (committed credit facilities, trade related contingencies, note issuance facilities, and similar categories) and calibrated by the EBA using pooled data from banks' own draw down histories. For revolving retail exposures (credit cards, overdrafts), a more granular schedule of two to three CCF buckets based on current utilization rate is permitted. The supervisory CCFs would be reviewed every five years, but the expectation is that they would change infrequently.

Probability of Default (PD)

36. **This paper proposes that the regulator calibrate PD models centrally, using pooled data, and that banks apply these models with limited bank-specific adjustment. The approach distinguishes two regimes.** For high-default portfolios, where defaults are frequent enough for robust statistical estimation, the regulator calibrates the coefficients of a logistic model every five years using anonymized loan-level data collected from all participating banks; each bank then calibrates its own constant annually so that the average PD of its current portfolio matches its long-run average default rate. For low-default portfolios, where bank-level default data are too sparse for meaningful calibration, PDs for rated borrowers are derived directly from long-run default rates by external rating grade, while a centrally calibrated model maps firm fundamentals to default rates for unrated borrowers; in neither case is a bank-specific model required. In both cases, the

14 Proposal C14 in Discussion Box 7 of the EBA's Discussion Paper, subject to back-testing requirements.

standardized approach remains available as a fallback for institutions that do not meet the data requirements. The maturity adjustment is retained unchanged, as it is a regulatory formula rather than a bank-estimated parameter.

37. **Within each model perimeter, the centrally calibrated PD model is defined by a single set of explanatory variables and estimated coefficients, applied uniformly across all participating banks.** The PD-specific perimeter decisions, including the choice of variables, the estimation sample, and any geographic or product-type splits, are described below for each of the two regimes.

High-default Portfolios

38. **For each high-default portfolio, a logistic model is calibrated every five years, with all parameters set centrally except a bank-specific constant recalibrated annually.**[15] The PD model maps a vector of borrower and loan characteristics to a PD; the EBA estimates the coefficients on each explanatory variable centrally, using a pooled loan-level dataset to maximize statistical power and ensure consistency across institutions. The logistic specification is chosen because it is the standard functional form in credit scoring [Lessmann et al., 2015], producing well-calibrated probabilities bounded between zero and one, and because its coefficients have a direct interpretation as log-odds ratios. To account for differences in risk appetite and origination standards across banks, and to anchor capital requirements through the cycle, each institution calibrates its own constant annually so that the average predicted PD across its current portfolio equals its long-run average default rate. The concept of regulator operated credit risk models is not without precedent.[16]

39. **The EBA would calibrate each PD model centrally using pooled loan-level data collected from banks.** Every five years, each participating bank submits an anonymized loan-level dataset covering a minimum observation period (ideally 10 years, to capture a full economic cycle). Each observation is an annual snapshot

15 Appendix A.7 sets out the formal specification of the logistic model and the recalibration of the bank-specific constant (the intercept of the logistic regression, in statistical terminology).

16 Seven euro area national central banks operate In-house Credit Assessment Systems (ICASs), logistic credit risk models for non-financial corporations calibrated on financial statement and credit register data, subject to common Eurosystem guidelines and annual performance monitoring [Auria et al., 2021]. ICASs serve collateral eligibility rather than prudential capital and incorporate expert judgment, but they demonstrate that public authorities can build and maintain such models at scale using the data infrastructure (AnaCredit, national credit registers) that the proposed framework would also require.

of a loan: borrower and loan characteristics recorded at the beginning of the year, paired with a binary indicator of whether a default (as defined under CRR Article 178; see European Banking Authority, 2016a) occurred during the following twelve months. The EBA selects five to ten standardized explanatory variables for each model perimeter. For retail mortgages, typical variables would include the debt-service-to-income ratio, the loan-to-value ratio at origination, seasoning (time since origination), current delinquency status, loan type (fixed or variable rate), and borrower employment status. For SME and mid-market corporate exposures, relevant variables include leverage, interest coverage, profitability, firm size, sector, and delinquency history. The EBA pools the data across banks, estimates the logistic regression, and publishes the resulting coefficients. Banks are not required to build, validate, or maintain their own PD models; they apply the published coefficients to their own loan data.

40. **The choice and definition of explanatory variables requires careful harmonization.** Some variables are defined uniformly across EU banks: delinquency status, loan-to-value ratio, and maturity are standardized by regulation or by accounting standards. Others present definitional challenges. Accounting ratios for smaller firms that do not report under IFRS may be computed differently across jurisdictions, reflecting national Generally Accepted Accounting Principles conventions. The proposal requires that these variables be reported in a standardized format defined by the EBA. Where reporting conventions differ across countries but the statistical relationship to default is similar, the data are pooled into a single model perimeter. Where the variable definitions or the associated default behavior diverge materially, the perimeter is split by jurisdiction. The five-year calibration cycle provides a natural moment to reassess these groupings and update the variable definitions as reporting practices converge. In practice, some explanatory variables may be noisy or unavailable for a subset of obligors, particularly smaller firms. The pooled estimation approach is robust to moderate data gaps, and the five-year recalibration cycle allows the variable set to be refined as data quality improves.[17]

17 For the small-firm segment, systematic data gaps are to be expected, since smaller firms rarely file IFRS-grade financials and often publish data with delays. Alternative data sources (such as transactional cash-flow data from bank accounts) can substitute for missing financial-statement variables in this segment [Hong Kong Monetary Authority and Hong Kong Applied Science and Technology Research Institute, 2020, Gambacorta et al., 2019]. EU data infrastructure is moving in the same direction: PSD2 already provides consent-based access to payment account data, and the Financial Data Access Regulation (FiDA) extends the same principle to broader commercial data, mirroring Hong Kong's Commercial Data Interchange [Hong Kong Monetary Authority, 2024]. The centrally calibrated model can incorporate such variables once the EBA harmonizes their definition, likely in a calibration cycle subsequent to the first.

41. **Each bank shifts the model's PD predictions up or down so that the portfolio average aligns with the bank's long-run average default rate (LRADR).** The LRADR is computed over a minimum of ten years. However, where the observation window is shorter than 20 years or contains no significant credit stress episode, the competent supervisory authority can adjust the LRADR. The bank recalculates the constant annually by solving for the value that equates the portfolio-average predicted PD to the updated LRADR. This is a straightforward numerical exercise: given the centrally published coefficients and the bank's own loan-level data, the constant is the unique scalar that aligns average predicted and observed default rates. No model development, validation infrastructure, or supervisory model approval is required for this step. The constant is auditable: it can be verified by any competent authority with access to the bank's default history and current portfolio composition.

42. **The annual recalibration of the constant ensures through-the-cycle stability and captures banks' risk appetite.** Because the constant is recalibrated annually to the long-run average default rate, it absorbs short-term fluctuations in borrower quality (rising debt-service ratios, increasing delinquency rates, or changes in employment status) rather than passing them through to capital requirements. At the same time, a bank that lends to lower-risk borrowers within a given model perimeter will accumulate a lower LRADR and therefore face lower capital requirements, while a bank that lends to riskier segments will face correspondingly higher capital charges. The mechanism is self-correcting and requires no supervisory judgment about model conservatism.

Low-default Portfolios

43. **Banks using F-IRB for low-default portfolios receive substantially lower capital requirements than the standardized approach, yet the models underpinning this relief cannot be statistically validated.**[18] As established above, low-default portfolios are defined precisely by the absence of sufficient default data to calibrate or backtest a model from internal experience. Whatever approach a bank adopts for risk differentiation, whether a shadow rating model that replicates external ratings, a scoring function trained on limited pooled default history, or expert judgment, the result cannot be verified against observed outcomes [Dedeaga and Crumpler, 2024, Pluto and Tasche, 2006]. PD calibration in these portfolios ultimately borrows from external sources, typically rating agencies' published long-run default rates. The risk differentiation these models provide over and above external ratings is therefore unverifiable. Yet the resulting capital relief is large: for a single-A rated corporate, the standardized approach assigns a risk weight of 50 percent (Credit Quality Step 2 (CQS 2) under CRR Article 122), while F-IRB, using a long-run default rate of approximately 6 basis points (the weighted-average one-year default rate for A− rated issuers over 1981– 2024; S&P Global Ratings, 2025, Table 9) and the supervisory LGD of 45 percent, produces a risk weight of approximately 22 percent, less than half the standardized charge. Under this paper's proposal, a minimum risk weight of 15 percent applies (see below).

44. **For rated borrowers, the risk differentiation that bank-specific IRB models provide does not demonstrably improve upon that of external credit ratings.** Credit rating agencies dedicate teams of analysts to each issuer, conduct detailed qualitative assessments, convene rating committees, and monitor ratings on an ongoing basis with upgrades and downgrades as conditions change. Their methodologies are subject to mandatory validation and backtesting under the EU Credit Rating Agencies Regulation [European Parliament and Council of the European Union, 2009, Article 8(3)], and their historical performance is publicly disclosed through ESMA's central repository (CEREP).[19] A bank's shadow rating model, by contrast, is a statistical approximation of these expert assessments, built on a narrower

18 Under the foundation IRB approach, banks estimate only the probability of default, while loss given default and exposure at default are set at supervisory values; under the advanced IRB approach, banks estimate all three parameters [Basel Committee on Banking Supervision, 2005b].

19 Article 11(2) of the CRA Regulation requires ESMA to establish a central repository making information on the historical performance of credit rating agencies available to the public. The repository is accessible at https://registers.esma.europa.eu/cerep-publication/.

set of financial variables, without the qualitative overlay, and validated against the very ratings it is trying to replicate. There is no evidence that shadow rating models produce more accurate default predictions than the agency ratings they emulate.

45. **For rated exposures in low-default portfolios, this paper proposes that PDs be derived directly from external credit ratings without any bank-specific model.** The external credit assessment institutions (ECAIs) used for this purpose must meet the eligibility criteria of CRR Articles 135–136, and their rating scales are already mapped to a common set of credit quality steps [European Commission, 2016]. Default rate data from multiple agencies are pooled at the credit quality step level [Tasche, 2013], providing a long-run average PD for each rating grade. Where a borrower is rated by more than one agency, the existing CRR rules for multiple assessments (Article 138) apply. This PD, combined with the supervisory LGD from the F-IRB approach, is entered directly into the risk-weight function, subject to a minimum risk weight of 15 percent. No bank-specific model is needed. All banks holding an unsecured exposure to the same single-A rated corporate would face the same PD, the same supervisory LGD, and therefore the same risk weight, eliminating the cross-bank variability that the current shadow rating approach produces. Where exposures are collateralized, the LGD would differ according to the type and value of collateral under the F-IRB framework, preserving risk sensitivity at the exposure level.

46. **For unrated borrowers in the low-default segment, this paper proposes that the regulator calibrate a centrally maintained shadow rating model.** A non-negligible fraction of corporates in the low-default segment will not carry an external credit rating.[20] For these exposures, the EBA would estimate a shadow rating model on the basis of publicly available firm fundamentals (leverage, interest coverage, profitability, firm size), using the observed relationship between these variables and default rates among rated firms to extrapolate PDs for unrated firms with comparable characteristics. This is functionally equivalent to the shadow rating models that IRB banks already operate, with the difference that a single centrally calibrated model replaces dozens of bank-specific ones. Figure 1 illustrates how this treatment fits within the overall framework.

20 For reference, approximately 75 percent of EU exposures to corporates with consolidated annual turnover above EUR 50 million are to unrated borrowers [European Banking Authority, 2019b, Table 6]. This class is broader than the low-default segment as defined here, but the figure gives a sense of scale.

Figure 1: Summary of proposed PD treatment by firm size.

The shaded area represents the universe of corporate borrowers. The upper horizontal dashed line indicates the threshold above which external credit ratings are typically available (ECAI mapping). The lower horizontal dashed line separates the centrally calibrated shadow rating models from the centrally calibrated statistical PD models. The vertical dashed line marks the boundary below which defaults are too rare for statistical calibration.

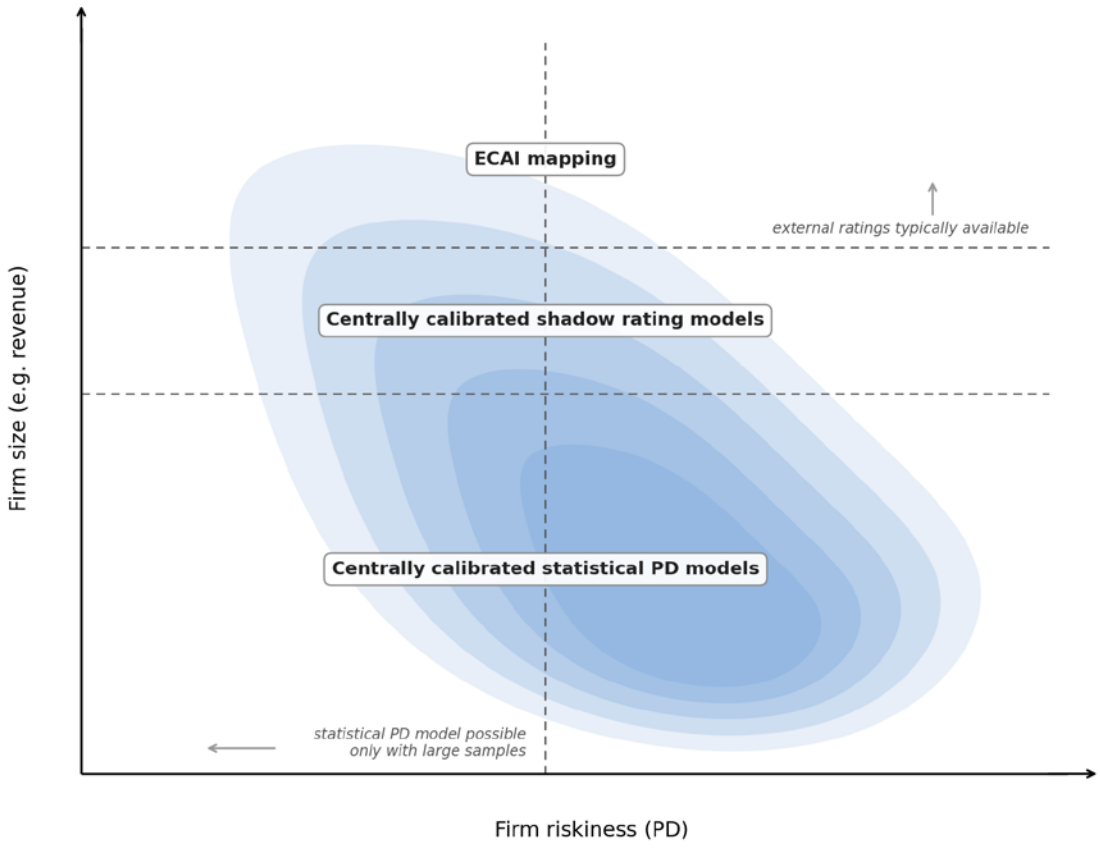

Source: World Bank.

47. **For specialized lending, the proposal mandates the use of supervisory slotting.** Specialized lending exposures, comprising project finance, income-producing real estate, commodities finance, and object finance, are characterized by a high degree of heterogeneity and a scarcity of defaults that make internal PD models difficult to calibrate and effectively non-backtestable. CRR3 already removes the A-IRB option for these exposures, restricting banks to F-IRB (own PD estimates, supervisory LGD). This paper proposes going one step further: removing the F-IRB option as well and requiring all specialized lending exposures to be treated under the supervisory slotting criteria of CRR Article 153(5). Under slotting, each exposure is assigned to one of five supervisory categories based on predefined criteria, with risk weights ranging from 50 to 250 percent depending on the category and remaining maturity. The EBA's May 2026 consultation paper on revised slotting criteria confirms that the framework is being actively developed as a standalone risk-differentiation tool [European Banking Authority, 2026a].

Risk-weight Floor

48. A minimum risk weight of 15 percent applies to all exposures, regardless of the PD produced by the model or the external rating. The rationale is statistical: at very low PDs, the confidence intervals around observed default rates are wide relative to the point estimates, and differences between adjacent rating grades are not distinguishable with any realistic sample.[21] The risk-weight function is steepest precisely in this region (Appendix A.1), amplifying statistically meaningless PD differences into large apparent capital differences. The 15 percent floor truncates this spurious precision at roughly the risk weight the Basel formula produces for a PD of 3 to 5 basis points with an LGD of 45 percent, the boundary below which default rate estimates lose statistical content. In practice, the floor binds only for the highest-rated investment-grade exposures; for all other portfolios, the formula risk weight exceeds the floor. Because the floor addresses calibration uncertainty at its source, the aggregate output floor of 72.5 percent of standardized approach RWA, designed to constrain bank-produced risk weights that supervisors could not fully trust, becomes redundant.

21 For example, with a true PD of 5 basis points and 1,000 exposures observed over 20 years, the 95 percent confidence interval spans roughly 2 to 9 basis points, covering two full rating grades. See the derivation in Appendix A.8. Once realistic asset correlations are incorporated, the uncertainty is considerably larger: Pluto and Tasche [2006] show that with 800 obligors and no observed defaults, the upper confidence bound on the PD of the best rating grade reaches 5.29 percent at the 99.9 percent level when default correlation consistent with the Basel corporate risk-weight function ($\rho = 12\%$) is assumed.

Governance and Entry Conditions

49. **Validation of the centrally calibrated models is the responsibility of the EBA, not of individual banks.** An independent unit in the regulator validates each model before endorsement, assessing the definition of model perimeters, discriminatory power, calibration accuracy, and stability on held-out samples of the pooled data. Once published, model parameters are fixed until the next recalibration cycle; if discriminatory power degrades over time, this is detected and corrected at the next recalibration. The need for traditional bank-level backtesting is also greatly reduced: under the current framework, backtesting monitors an agency problem (banks have an incentive to underestimate risk parameters), but under this proposal, bank-specific inputs are limited to observables (the long-run average default rate for PD, the long-run average loss rate for LGD) that cannot be optimized. Annual updates to these inputs automatically adjust predicted values to each bank's realized experience. For low-default portfolios, defaults are too rare for statistical backtesting, which is why external ratings serve as the anchor. Banks retain responsibility for data quality and the accuracy of the bank specific inputs they compute.

50. **The conditions for a bank to use the centrally calibrated models are substantially simpler than the current IRB approval process, but not straightforward.** Today, IRB approval requires years of prior experience with internal rating systems, comprehensive model documentation, an independent credit risk control unit, and a demonstrated "use test" [European Parliament and Council of the European Union, 2013b, European Banking Authority, 2016b, Articles 143–145]. Much of this becomes redundant when the bank no longer develops its own model. What remains essential is (i) loan-level data handling: the bank must be able to classify each exposure, assign the correct risk parameters, and report in the format required by the EBA; (ii) for high-default portfolios, a verified default history of at least five years and the ability to compute a reliable long-run average default rate (for PD) and long-run average loss rate (for LGD); and (iii) a supervisory assessment confirming that the bank's data infrastructure and risk governance meet the required standard.

51. **New entrants begin under the standardized approach and transition once they meet the entry conditions.** A newly licensed bank, or an existing bank entering a new exposure class, would apply the standardized approach until it can satisfy the requirements above. For high-default portfolios, the binding constraint is accumulating the minimum default history needed to compute a reliable LRADR, subject to the conservatism adjustment described above where the observation window is shorter than 20 years. The same applies to the long-run average loss rate on unsecured exposures, which requires a comparable history of resolved defaults. For low-default portfolios, no bank-specific calibration parameter is needed, but the bank must still demonstrate loan-level data handling before it may apply the centrally calibrated PD schedule or ECAI mapping in place of the standardized approach risk weights. Banks that do not wish to invest in this infrastructure simply remain under the standardized approach.

52. **Banks remain free to model their own risks, and should.** The proposal removes internal models from the calculation of capital requirements; it does not remove them from banks. Sound underwriting, pricing, limit setting, and early-warning systems will continue to rely on internal models, and banks should continue to build and backtest them. What changes is their purpose: freed from the incentive to minimize capital requirements and from the regulatory form this imposed, these models no longer need to be calibrated through the cycle, padded with margins of conservatism, or frozen pending supervisory approval. Banks can build point-in-time, agile models, recalibrated as often as useful, using machine learning and any data they find informative. How much to invest in them becomes a business decision disciplined by the bank's own loss experience rather than a compliance exercise, and none of the investment is spent parsing modeling regulation or engineering estimates that are defensibly low rather than accurate.

Assessment

53. This section assesses the proposal put forward in this paper and five alternative approaches against the five objectives set out above: risk sensitivity, cross-bank comparability, cost efficiency, Basel compatibility, and through the cycle stability. The first subsection examines how the proposal performs on each criterion. The second evaluates five alternative paths to IRB simplification: incremental reform of the existing framework, the application of F-IRB to all portfolios including retail, a full return to the standardized approach, a granular standardized approach structured along risk drivers, and a leverage ratio as the binding capital constraint. A comparative summary concludes the section.

Assessment of the Proposal

54. The proposal preserves PD risk sensitivity. The general approach to PD risk differentiation is unchanged: borrower and loan characteristics are combined in a logistic regression to predict defaults, and the resulting PDs feed into the same risk-weight function at the same confidence level. The credit scoring literature consistently finds that simple logistic models with a small number of well-chosen variables capture the large majority of achievable discriminatory power [Lessmann et al., 2015]; additional bank-specific model complexity yields diminishing marginal returns. A genuine trade-off is that the proposal uses fewer, standardized explanatory variables, which may sacrifice some portfolio-specific signal. This loss is partially compensated by the fact that coefficients are estimated on a pooled dataset that is substantially larger than any individual bank's sample, improving statistical precision and out-of-sample stability. The bank-specific constant ensures that each institution's average PD reflects its own portfolio composition and risk appetite, so the proposal does not force a one size fits all level of predicted risk.

55. **Whether bank-specific PD models differentiate risk better or worse than centrally calibrated ones is an empirical question, and it can be tested.** The comparison requires applying both models to the same portfolio and measuring their out of sample discriminatory power, using the AUC or the conditional likelihood statistic,[22] taking into account the statistical uncertainty around each metric. The AUC measures ranking accuracy: an AUC of 0.70 means that for any pair consisting of a defaulted and a non-defaulted borrower, the model assigns the wrong ranking 30 percent of the time. The conditional likelihood is a more demanding test: it asks whether the observed pattern of which rating grades experienced defaults is distinguishable from what would occur if the grade-level PDs had been assigned randomly. In practice, two well-constructed models applied to the same portfolio would likely produce similar discriminatory power metrics. The question is whether the difference, if any, is statistically significant, and the answer depends entirely on the number of defaults available for the comparison. In many portfolios, particularly low-default portfolios, the available data may be insufficient to distinguish modest differences in model performance with statistical confidence [Pluto and Tasche, 2006]; see Appendix A.8.

56. **The proposal preserves LGD risk sensitivity.** Risk sensitivity is preserved through the structural formula's observable inputs: collateral type and coverage, time-to-foreclosure, and bank-specific unsecured recovery rates. These parameters differentiate LGD at the exposure level without requiring a regression model, and because they are drawn from the bank's own data, they reflect genuine differences in collateral environments and workout efficiency across institutions. The proposal also dispenses with the requirement for banks to estimate a separate downturn LGD [Basel Committee on Banking Supervision, 2005a]. Because the LGD treatment follows the F-IRB logic of supervisor-set parameters rather than bank-estimated models, the downturn conservatism is embedded directly in the calibration of the supervisory haircuts, which reflect forced-liquidation values under adverse market conditions.

57. **Cross-bank comparability improves substantially under the proposal over the status quo.** Under the current framework, the EBA's benchmarking exercises document persistent and wide cross-bank dispersion in PD estimates for similar exposures [European Banking Authority, 2025a]. Under the proposal, all banks in a given segment apply the same model with the same parameters. For low-default port-

22 The Area Under the Receiver Operating Characteristic Curve (AUC) and the conditional likelihood statistic are defined formally in Appendix A.8.

folios, where the proposal maps external ratings to centrally determined PDs, the resulting risk weights are identical across banks. For high-default portfolios, two banks applying the model to borrowers with identical characteristics will produce identical relative PD differentiation, but PD levels may differ because each bank calibrates its own constant to its long-run default rate. This residual variation reflects genuine differences in default experience rather than modeling choices.

58. **Cost efficiency improves substantially.** Model development, validation, and supervisory review are centralized at the EBA and performed once for the entire banking system. Individual banks no longer need to maintain the three separate functions (development, validation, audit) currently mandated for IRB [European Banking Authority, 2016b], nor the documentation, ex-ante notifications, material model change approvals, and multi-year rollout plans that the current framework demands [European Commission, 2014]. The supervisory apparatus built to police bank-specific models, including dedicated inspection divisions, TRIM-style reviews, and annual benchmarking exercises in their current form, can be scaled back substantially. Centralization does create new costs at the regulator: the EBA would need a permanent modeling, data, and validation capacity. These costs are, however, incurred once for the entire system rather than in parallel at each IRB institution and its supervisors, and part of the required infrastructure already exists (the EBA's annual benchmarking exercises already process and analyze banks' model outputs, and its biennial EU-wide stress tests handle detailed bank data collected through the competent authorities). While a central function of several hundred staff may cost tens of millions of euros per year, this is one to two orders of magnitude below the apparatus it replaces – Appendix A.5 estimates the current aggregate cost of IRB compliance at EUR 0.7 to 1.2 billion per year.

59. **The proposal is compatible with the Basel framework.** The core risk-weight function (CRE 31) is preserved in full: the Vasicek/Asymptotic Single Risk Factor (ASRF) model, the 99.9 percent confidence level, and the maturity adjustment are all unchanged. F-IRB already relies on supervisory estimates for LGD and EAD (CRE 30.36), establishing the principle that risk-sensitive capital requirements can coexist with supervisor-provided risk components. The supervisory slotting criteria for specialized lending (CRE 33) go further, providing a direct precedent for supervisor-determined risk differentiation: banks map exposures to supervisory categories, each associated with specific risk weights, without estimating

any risk component themselves. The proposal generalizes this approach to all asset classes with greater granularity (continuous PD via regression rather than five coarse buckets) and with bank-specific calibration through the constant. CRE 30.1 explicitly allows "supervisory values" for risk components, and CRE 36 defines a "rating system" functionally, without prescribing who designs the model. The proposal can therefore be framed as a European implementation choice within the existing Basel architecture, using the terminology "standardized IRB models", comparable to the United States, where the Collins Amendment floor has de facto neutralized the capital benefits of internal models by requiring that standardized requirements always bind when they exceed advanced approaches estimates [Risk.net, 2022].[23]

60. **Through-the-cycle stability is embedded in the design.** For high-default portfolios, the bank-specific constant is calibrated to the LRADR, which by definition smooths across the economic cycle; even a severe single-year shock shifts the average by at most one-tenth of its magnitude. The model parameters, recalibrated every five years on pooled data, are also inherently stable. The main potential source of PD cyclicality is that borrower-level explanatory variables (leverage, debt-service ratios) move with the cycle, but the fixed constant dampens this: the portfolio-average predicted PD remains anchored to the LRADR regardless of where individual borrowers sit on the model's PD scale. For low-default portfolios, PDs are derived from long-run average default rates by external rating grade; rating migrations introduce some cyclicality, but external ratings are designed to reflect permanent changes in creditworthiness, limiting migration-driven PD movements [European Securities and Markets Authority, 2024]. On the LGD side, collateral values are the main cyclical channel, but the structural formula limits the impact: the unsecured portion is a long-run average, and CRR revaluation rules prevent stale appraisals from inflating coverage during upswings.

61. **A centralized model concentrates model risk, but the current system's model diversity is distorted by capital incentives, not the pursuit of accuracy.** Under the proposal, all banks in a given segment apply the same model, so a systematic

23 The Basel Committee's 2014 RCAP graded the EU's IRB implementation as "materially non-compliant" [Basel Committee on Banking Supervision, 2014b], and the IMF's 2025 FSAP repeated this finding for CP16 (Capital Adequacy) [International Monetary Fund, 2025]. The bulk of the materially-non-compliant grading reflects statutory EU policy choices (the Danish Compromise, the SME sup porting factor, the limited scope of CVA, the permanent partial use of the standardized approach for sovereign exposures in IRB banks) that lie outside the scope of this paper. One strand of the FSAP finding, however, is directly addressed by the proposal: the FSAP highlights significant ECB resource constraints in internal model validation, the introduction of a triage system, and the predominantly reactive nature of internal model reviews. By replacing bank-developed models with centrally calibrated ones, the proposal removes the bulk of the supervisory work that creates these resource constraints.

error in the centrally calibrated parameters would affect the entire banking system simultaneously. This is a legitimate concern. However, model risk diversification delivers its theoretical benefits only when each model independently aims for the most accurate estimate. When every model builder has a structural incentive to produce the lowest possible risk estimate, as is the case under the current framework, the resulting dispersion does not reflect genuine methodological diversity but rather variation in calibration aggressiveness. The EBA's benchmarking exercises confirm this: less than half of observed cross-bank variation in estimated PDs can be explained by realized default rates [European Banking Authority, 2025a]. Banks use similar estimation techniques, draw on similar explanatory variables, and operate under the same tight regulatory constraints on model design. A centrally calibrated model makes the concentration of model risk explicit and therefore auditable and correctable: the model's specification and coefficients are published and subject to public scrutiny, its out-of-sample performance on pooled data is observable, and the 15 percent risk-weight floor provides an explicit buffer against calibration error at the segment of the PD distribution where model uncertainty is greatest.

Assessment of Alternative Approaches

62. **This subsection assesses five alternative paths to simplifying the IRB framework against the same objectives.** The alternatives considered are: (1) incremental reform of the existing framework as proposed by European institutions, (2) the application of F-IRB to all portfolios including retail, (3) a full return to the standardized approach, (4) a granular standardized approach with risk-driver-based tables, and (5) a leverage ratio as the binding capital constraint. Table 1 summarizes the assessment.

63. **Incremental reform, as put forward by European institutions, streamlines the existing IRB framework but improves its cost efficiency only at the margin.** The proposals fall into three categories: simplified estimation requirements, with optional fallbacks for margins of conservatism, downturn LGD, in-default LGD, and credit conversion factors [European Banking Authority, 2026c]; streamlined supervisory processes, with fewer model changes requiring prior approval, no automatic on-site investigations, and a less prescriptive model review cycle [European Banking Authority, 2026b, European Central Bank, 2026b, 2025d]; and

a reduced model scope, with internal models confined to banks' most material exposure classes [Eurosystem, 2026]. The core apparatus, however, remains: banks develop, validate, and maintain their own models, the structural incentive to underestimate risk parameters persists, and where approval friction is removed, ex-ante supervisory scrutiny gives way to banks' own validation, with attendant risks to model quality. Every other dimension is inherited from the current framework: risk sensitivity and through-the-cycle stability remain high, Basel compatibility remains full, and cross-bank comparability remains low.

64. **A universal application of F-IRB to all portfolios, including retail, would simplify much of the framework but leave its central weakness, bank-estimated PDs, untouched.** Under CRR3, F-IRB is already mandatory for large corporates and institutions and unavailable for retail; adopting it for the remaining mid-market and SME books would barely change the status quo. The alternative considered here therefore extends supervisory LGD and CCF values to all portfolios, including retail. This eliminates LGD and EAD modeling costs and the associated cross-bank variability; and because Basel does not foresee a foundation approach for retail, Basel compatibility rests on the same supervisory-values logic as this paper's proposal. PD estimation, however, remains with banks, preserving the incentive to underestimate the parameter that drives the largest share of risk-weight variation, together with the rating systems, validation functions, inspections, and approval processes built around it. Universal FIRB thus follows the same institutional logic as this paper's proposal, removing LGD and EAD from bank discretion, but in a cruder form: its flat supervisory values sacrifice the risk differentiation that the proposal's structural LGD retains through collateral type, coverage, and workout experience, and it stops short of the step that delivers most of the benefit, the centralization of PD. Cost efficiency and cross-bank comparability improve for LGD but not for PD, where the bulk of the modeling and supervisory apparatus lies, while through-the-cycle stability remains high, as long-run PD calibration is unchanged and fixed supervisory values do not move with the cycle.

65. **A full return to the standardized approach eliminates internal models entirely, achieving maximum simplicity at the cost of risk sensitivity.** Some coarse differentiation survives through exposure classes, external ratings, and loan-to-value ratios, which distinguishes the standardized approach from the leverage ratio. In Europe, however, this differentiation has little bite for the corporate book: approximately 75 percent of EU exposures to larger corporates are to unrated borrowers [European Banking Authority, 2019b, Table 6], which attract a flat risk weight of 100 percent regardless of creditworthiness. Within each bucket, capital charges penalize safer borrowers and undercharge riskier ones. Cross-bank comparability and cost efficiency are maximal. Through-the-cycle stability is high since flat risk weights do not move with economic conditions. Basel compatibility is full. The loss of risk sensitivity, however, is substantial, particularly for well-diversified portfolios of high-quality exposures.

66. **A granular standardized approach structures risk-weight tables along observable risk drivers, recovering meaningful risk sensitivity without internal models, but at the price of coarser differentiation and procyclicality.** Instead of a flat risk weight for all exposures in one exposure class, the regulator defines a grid along the two strongest risk drivers, for example leverage and delinquency status for corporates, or debt-service-to-income ratio and loan-to-value ratio for residential mortgages. The natural calibration for the cell values is the average of the Basel risk-weight function for borrowers in each bucket, with the aggregate capital level of the current standardized approach available as an anchor (Appendix A.9 provides illustrative tables for two asset classes). Cross-bank comparability is full and cost efficiency is high. The principal limitations are that a grid built on two risk drivers with discrete buckets provides coarser risk differentiation than a continuous model drawing on a fuller set of variables, and that the risk drivers defining the grid move with the economic cycle, introducing procyclicality. Basel compatibility is only partial: although the cell values derive from the Basel risk-weight function, the grid in effect defines a new, more risk-sensitive standardized approach, with risk weights that depart from those prescribed by the Basel framework.

67. **Making the leverage ratio the binding capital constraint, as advocated by Admati and Hellwig [2024], eliminates risk weighting altogether, achieving perfect comparability at the cost of abandoning risk sensitivity entirely.** Banks hold a fixed percentage of equity against total unweighted assets, removing both the incentive to underestimate risk and the need for risk models. Cross-bank comparability is perfect, cost efficiency is maximal, and through-the-cycle stability is very high. The trade-off is the complete loss of risk sensitivity: a government bond and a speculative-grade corporate loan attract the same capital charge, creating an incentive to shift toward riskier, higher-yielding assets as was the case under Basel 1. Basel compatibility is absent: the leverage ratio exists in the Basel framework only as a backstop to risk-weighted requirements, and replacing those requirements outright would abandon the risk-based minimum at the core of the framework.

Table 1: Comparative assessment of the status quo and six approaches to credit risk capital requirements. Each cell summarizes how the approach performs against the objective. "This proposal" refers to the regulator-calibrated framework described in the preceding sections.

	Risk sensitivity	Cross-bank comparability	Cost efficiency	Basel compatibility	Through-the-cycle stability
Status quo	High	Low	Very low	Full	High
This proposal	High	High	High	Compatible	High
Incremental reform	High	Low	Low	Full	High
Universal F-IRB	Moderate	Moderate	Low	Compatible	High
Return to SA	Low	Full	Very high	Full	High
Granular SA	Moderate	Full	Very high	Partial	Moderate
Leverage ratio	None	Full	Very high	None	Very high

Source: World Bank.

Conclusion

68. **Risk-sensitive capital requirements do not require bank-developed models.** This paper has argued that the core benefit of the IRB framework can be preserved while centralizing model calibration. The proposed approach retains the Basel risk-weight function, the 99.9 percent confidence level, and calibration to each bank's own long-run average default rate, while eliminating the structural incentive for banks to underestimate risk parameters. When the IRB framework was introduced in the early 2000s, banks may have needed a regulatory incentive to invest in what was then novel data and modeling infrastructure. Two decades later, logistic credit scoring models are computationally trivial. The scarce input is not computational capacity or modeling expertise; it is an unbiased calibration. The proposal redirects the regulatory framework accordingly: banks contribute data, and the regulator ensures that the models built on those data serve prudential objectives rather than capital optimization.

69. **Among the assessed alternatives, the proposal is the only one that meets all five objectives.** The comparative assessment (Table 1) makes the trade-offs explicit: incremental reform leaves cost and comparability largely where they are; universal F-IRB follows the same institutional logic as the proposal but in cruder form, leaving PD, the dominant driver of both cost and risk-weight variation, with banks; a return to the standardized approach buys simplicity at the price of risk sensitivity; a granular standardized approach comes closest in spirit, reading centrally calibrated risk weights from tables, but offers coarser differentiation, more cyclicality, and weaker Basel compatibility; and a binding leverage ratio abandons risk weighting altogether. Centralized calibration with bank-specific constants is the only configuration that combines risk sensitivity, comparability, and cost efficiency within the boundaries of the Basel framework.

70. **The regulatory simplification without deregulation that follows is substantial.** At the European level, the body of regulation governing IRB models, currently exceeding 1,000 pages, becomes largely redundant. The supervisory structures built to police bank-developed models can be scaled back to a narrower mandate focused on data quality and the integrity of banks' reported default and loss histories. This is not deregulation; it is simplification, the removal of an apparatus built to manage an agency problem that the proposal addresses at its source. The Commission's targeted consultation on banking competitiveness [European Commission, 2026] frames the central question directly: whether the prudential framework strikes the right balance between risk sensitivity and complexity. This paper argues that it does not and proposes a specific alternative. The Eurosystem's response to the consultation [Eurosystem, 2026] endorses its premise: it calls for simplification and harmonization. The proposal could also prompt a reconsideration of the Basel safeguards designed to compensate for the shortcomings of bank-developed models. The aggregate output floor was introduced because regulators did not trust bank-produced risk weights; if risk weights derive from regulator-calibrated models, the original rationale for capping their capital benefit is substantially weakened. The proposal's own 15 percent risk weight floor addresses calibration uncertainty directly, making a blanket floor tied to the standardized approach redundant. Over time, the same logic may also bear on the leverage ratio, which was conceived in part as a backstop against risk-weight manipulation.

71. **This proposal creates more winners than losers.** Banks currently on the standardized approach stand to gain: any institution that can construct a robust default database and meet the data infrastructure requirements could access centrally calibrated models, obtaining more risk-sensitive, hence lower, capital requirements without building a dedicated IRB function. This broadens the efficiency gains of risk-sensitive capital to a larger share of the European banking system. On the other side, banks that have invested heavily in IRB infrastructure face transition costs as they wind down modeling, validation, and audit functions that the framework no longer requires. These costs are bounded in time, and they are dwarfed by the recurrent savings: the banking system as a whole obtains comparable risk sensitivity at a fraction of the current compliance cost, freeing financial resources and quantitative talent for more productive uses.

The proposal would also reduce competitive distortions created by the current IRB framework. By giving all participating banks access to the same centrally calibrated models, it eliminates the capital advantage that currently accrues to institutions with more aggressively calibrated IRB models and removes the structural barrier that has kept standardized approach banks from competing on equal terms for lower-risk lending.

72. **A natural roll-out would begin with low-default portfolios, complemented by a pilot on one high-default segment.** For low-default portfolios, the statistical case against current IRB models is strongest (discriminatory power cannot be verified with available default data), the capital impact of PD differentiation is largest (the risk-weight function is steepest at low PDs), and the proposal is simplest to implement (PDs are derived from external ratings without bank-specific slope calibration). A quantitative impact study comparing the capital outcomes of the proposed approach against current IRB results for large corporate, sovereign, and institutions would provide the evidence base for transition. In parallel, a pilot exercise on one high-default segment (such as retail mortgages) would test the central calibration process end to end: data collection and pooling, model perimeter definition, explanatory variable selection, model estimation, and bank-specific intercept recalibration. The pilot would thereby settle empirically whether a centrally calibrated model with a small set of standardized variables matches the discriminatory power of bank-specific models, tested out of sample. The results of both exercises would determine the pace and sequencing of a broader transition.

73. **The IRB framework offers a model case for simplification without deregulation.** Calls for regulatory simplification often stall on the concern that reducing requirements means weakening safeguards. The proposal in this paper avoids that trade-off: it preserves the Basel risk-weight function, the 99.9 percent confidence level, and borrower-level risk differentiation while largely eliminating the agency problem and the regulatory apparatus built to manage it. The bank discretion that remains, over the classification of defaults, the computation of the explanatory variables, and the LGD inputs, is anchored in auditable data and disciplined by standard supervisory practice. The EBA itself acknowledges that the burden of IRB modeling may not be proportional to its benefits. This paper offers a specific alternative – the next step is to test it.

References

Puriya Abbassi and Michael Schmidt. A comprehensive view on risk reporting: Evidence from supervisory data. *Journal of Financial Intermediation*, 36:74–85, 2018. doi:10.1016/j.jfi.2018.03.001.

Viral V. Acharya, Sreedhar T. Bharath, and Anand Srinivasan. Does industry-wide distress affect defaulted firms? Evidence from creditor recoveries. *Journal of Financial Economics*, 85(3):787–821, 2007.doi:10.1016/j.jfineco.2006.05.011.

Anat R. Admati and Martin Hellwig. *The Bankers' New Clothes: What's Wrong with Banking and What to Do about It.* Princeton University Press, expanded edition, 2024.

Laura Auria, Markus Bingmer, Carlos Mateo Caicedo Graciano, Clémence Charavel, Sergio Gavilá, Alessandra Iannamorelli, Aviram Levy, Alfredo Maldonado, Florian Resch, Anna Maria Rossi, and Stephan Sauer. Overview of central banks' in-house credit assessment systems in the euro area. Occasional Paper Series 284, European Central Bank, October 2021. URL https://www.ecb.europa.eu/pub/pdf/scpops/ecb.op284~bbce5257bf.en.pdf.

Autorité de contrôle prudentiel et de résolution. Banking supervision directorates, 2025. URL https://acpr.banque-france.fr/en/acpr/organisation/banking-supervision-directorates. Risk Modelling Control Unit within the On-site Inspection Delegation.

Banca d'Italia. Dipartimento vigilanza bancaria e finanziaria, 2025. URL https://www.bancaditalia.it/chi-siamo/organizzazione/ac/vigilanza-bancaria-finanziaria/index.html. Accessed March 2026.

Banco de España. The banco de españa supervisory model, 2024. URL https://www.bde.es/f/webbde/COM/Supervision/funciones/Ficheros/en/Banco_de_Espana_supervisory_model_clean.pdf. Accessed March 2026.

Basel Committee on Banking Supervision. Guidance on paragraph 468 of the framework document: The irb use of loss given default under a foundation and advanced approach. Technical report, Bank for International Settlements, July 2005a. URL https://www.bis.org/publ/bcbs115.pdf.

Basel Committee on Banking Supervision. An explanatory note on the Basel II IRB risk weight functions. Technical report, Bank for International Settlements, July 2005b. URL https://www.bis.org/bcbs/irbriskweight.pdf.

Basel Committee on Banking Supervision. Studies on the validation of internal rating systems. Technical Report Working Paper No. 14, Bank for International Settlements, May 2005c. URL https://www.bis.org/publ/bcbs_wp14.htm. Revised version.

Basel Committee on Banking Supervision. Basel III leverage ratio framework and disclosure requirements. Technical report, Bank for International Settlements, January 2014a. URL https://www.bis.org/publ/bcbs270.pdf.

Basel Committee on Banking Supervision. Regulatory consistency assessment programme (RCAP): Assessment of Basel III regulations – European Union. Technical Report d300, Bank for International Settlements, December 2014b. URL https://www.bis.org/bcbs/publ/d300.pdf.

Basel Committee on Banking Supervision. Reducing variation in credit risk-weighted assets: Constraints on the use of internal model approaches. Technical report, Bank for International Settlements, March 2016. URL https://www.bis.org/bcbs/publ/d362.pdf. Consultative Document.

Basel Committee on Banking Supervision. Basel III: Finalising post-crisis reforms. Technical report, Bank for International Settlements, December 2017. URL https://www.bis.org/bcbs/publ/d424.pdf.

Markus Behn, Rainer Haselmann, and Vikrant Vig. The limits of model-based regulation. *The Journal of Finance*, 77(3):1635–1684, 2022. doi:10.1111/jofi.13124.

Tobias Berg and Philipp Koziol. An analysis of the consistency of banks' internal ratings. *Journal of Banking & Finance*, 78:27–41, 2017. doi:10.1016/j.jbankfin.2017.01.013.

Andreas Blöchlinger and Markus Leippold. Economic benefit of powerful credit scoring. *Journal of Banking & Finance*, 30(3):851–873, 2006. doi:10.1016/j.jbankfin.2005.07.014.

Alessandro Calza, Julius-Benjamin Hey, Alessandro Parrini, and Stephan Sauer. Corporate loans, banks' internal risk estimates and central bank collateral: Evidence from the euro area. Working Paper Series 2579, European Central Bank, July 2021. URL https://www.ecb.europa.eu/pub/pdf/scpwps/ecb.wp2579~e7d577c5ad.en.pdf.

Sergei A. Davydenko and Julian R. Franks. Do bankruptcy codes matter? A study of defaults in France, Germany, and the U.K. *The Journal of Finance*, 63(2):565–608, 2008. doi:10.1111/j.1540-6261.2008.01325.x.

Cem Dedeaga and Michael Crumpler. How to improve IRB modelling for low-default portfolios. Oliver Wyman Insights, November 2024. URL https://www.oliverwyman.com/our-expertise/insights/2024/nov/how-to-improve-irb-modelling-for-low-default-portfolios.html. Accessed March 2026.

Deloitte. 2025 EMEA model risk management survey. Technical report, Deloitte, 2025. URL https://www.deloitte.com/dk/en/services/consulting/research/2025-emea-model-risk-management-survey.html. Accessed March 2026.

Elizabeth R. DeLong, David M. DeLong, and Daniel L. Clarke-Pearson. Comparing the areas under two or more correlated receiver operating characteristic curves: A nonparametric approach. Biometrics, 44(3):837–845, 1988. doi:10.2307/2531595.

Deutsche Bundesbank. Directorate general banking and financial supervision reorganised, 2023. URL https://www.bundesbank.de/en/tasks/topics/directorategeneral-banking-and-financial-supervision-reorganised-856146. Division B 4: On-site Inspections / Internal Risk Models.

Simeon Djankov, Oliver Hart, Caralee McLiesh, and Andrei Shleifer. Debt enforcement around the world. *Journal of Political Economy*, 116(6):1105–1149, 2008. doi:10.1086/595015.

Mario Draghi. The future of European competitiveness: A competitiveness strategy for Europe. Technical report, European Commission, September 2024. URL https://commission.europa.eu/topics/competitiveness/draghi-report_en. Presented 9 September 2024.

Bernd Engelmann and Robert Rauhmeier. Measuring the discriminative power of rating systems. In *The Basel II Risk Parameters: Estimation, Validation, and Stress Testing*, pages 83–103. Springer, 2006. doi:10.1007/3-540-33087-9_5.

European Banking Authority. Guidelines on the application of the definition of default under article 178 of regulation (EU) no 575/2013. Technical Report EBA/GL/2016/07, European Banking Authority, September 2016a. URL https://www.eba.europa.eu/regulation-and-policy/credit-risk/guidelines-on-the-application-of-the-definition-of-default. Applied from 1 January 2021.

European Banking Authority. Final draft regulatory technical standards on the specification of the assessment methodology for competent authorities regarding compliance of an institution with the requirements to use the IRB approach. Technical Report EBA/RTS/2016/03, European Banking Authority, 2016b. URL https://www.eba.europa.eu/activities/single-rulebook/regulatory-activities/credit-risk/regulatory-technical-standards-2.

European Banking Authority. Final draft regulatory technical standards on assigning risk weights to specialised lending exposures. Technical Report EBA/RTS/2016/02, European Banking Authority, 2016c. URL https://www.eba.europa.eu/regulation-and-policy/credit-risk/regulatory-technical-standards-on-specialised-lending-exposures.

European Banking Authority. Guidelines on PD estimation, LGD estimation and the treatment of defaulted exposures. Technical Report EBA/GL/2017/16, European Banking Authority, 2017. Final report published November 2017.

European Banking Authority. Final draft regulatory technical standards on the specification of the nature, severity and duration of an economic downturn. Technical Report EBA/RTS/2018/04, European Banking Authority, 2018. URL https://www.eba.europa.eu/regulatory-technical-standards-and-guidelines-estimation-and-identification-economic-downturn-irb.

European Banking Authority. Guidelines on the estimation of LGD appropriate for an economic downturn. Technical Report EBA/GL/2019/03, European Banking Authority, 2019a. URL https://www.eba.europa.eu/publications-and-media/press-releases/eba-publishes-final-guidelines-estimation-lgd-under-economic.

European Banking Authority. Policy advice on the Basel III reforms: Credit risk – standardised approach and IRB approach. Technical Report EBA-Op-2019-09a, European Banking Authority, August 2019b. URL https://www.eba.europa.eu/sites/default/files/Policy%20Advice%20on%20Basel%20III%20reforms%20-%20Credit%20Risk.pdf.

European Banking Authority. Guidelines on credit risk mitigation for institutions applying the IRB approach with own estimates of LGDs. Technical Report EBA/GL/2020/05, European Banking Authority, 2020. URL https://www.eba.europa.eu/guidelines-credit-risk-mitigation-institutions-applying-irb-approach-own-estimates-lgds.

European Banking Authority. Summary report on the 2022 credit risk benchmarking exercise. Report EBA/REP/2023/09, European Banking Authority, 2023a. URL https://www.eba.europa.eu/publications-and-media/publications.

European Banking Authority. Supervisory handbook for the validation of internal ratings based rating systems. Technical report, European Banking Authority, 2023b. URL https://www.eba.europa.eu/activities/single-rulebook/regulatory-activities/model-validation/supervisory-handbook-validation.

European Banking Authority. Eba report on the 2023 credit risk benchmarking exercise. Report EBA/REP/2024/05, European Banking Authority, April 2024. URL https://www.eba.europa.eu/publications-and-media/publications.

European Banking Authority. Eba report on the 2024 credit risk benchmarking exercise. Report EBA/REP/2025/13, European Banking Authority, March 2025a. URL https://www.eba.europa.eu/publications-and-media/publications.

European Banking Authority. Consultation paper on guidelines on the methodology to estimate and apply credit conversion factors. Technical Report EBA/CP/2025/10, European Banking Authority, 2025b. URL https://www.eba.europa.eu/activities/single-rulebook/regulatory-activities/model-validation/guidelines-methodology-estimate-and-apply-credit-conversion-factors-under-capital-requirements.

European Banking Authority. Risk assessment report. Technical report, European Banking Authority, December 2025c. URL https://www.eba.europa.eu/publications-and-media/publications/risk-assessment-report-december-2025.

European Banking Authority. Consultation paper on draft regulatory technical standards amending delegated regulation (eu) 2021/598 on the supervisory slotting criteria approach. Consultation Paper EBA/CP/2026/09, European Banking Authority, May 2026a. URL https://www.eba.europa.eu/publications-and-media/press-releases/eba-consults-amendments-rts-assignment-risk-weights-specialised-lending-exposures-under-supervisory.

European Banking Authority. Draft regulatory technical standards on amending delegated regulation (EU) no 529/2014 with regard to regulatory technical standards for assessing the materiality of extensions and changes of the internal ratings based approach. Technical report, European Banking Authority, March 2026b. URL https://www.eba.europa.eu/publications-and-media/press-releases/eba-streamlines-supervisory-approvals-irb-model-changes.

European Banking Authority. Discussion paper on the simplification and assessment of the credit risk framework. Discussion Paper EBA/DP/2026/01, European Banking Authority, February 2026c. URL https://www.eba.europa.eu/publications-and-media/press-releases/eba-launches-consultation-simplifying-credit-risk-framework. EBA/DP/2026/01.

European Banking Federation. Simply competitive: How simplification can strengthen the eu banking sector. Technical report, European Banking Federation, July 2025. URL https://www.ebf.eu/wp-content/uploads/2025/07/Simply-Competitive-report-EBF.pdf.

European Central Bank. Targeted review of internal models (TRIM): Media briefing conference call, February 2017. URL https://www.bankingsupervision.europa.eu/about/banking-supervision-explained/shared/pdf/trim_presentation_slides.en.pdf.

European Central Bank. Guide to on-site inspections and internal model investigations. Technical report, European Central Bank, Banking Supervision, September 2018. URL https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.osi_guide201809.en.pdf.

European Central Bank. TRIM: Targeted review of internal models — project report. Technical report, European Central Bank, Banking Supervision, 2021. URL https://www.bankingsupervision.europa.eu/ecb/pub/pdf/ssm.trim_project_report~aa49bb624c.en.pdf.

European Central Bank. Annual report on supervisory activities 2022. Technical report, European Central Bank, Banking Supervision, March 2023. URL https://www.bankingsupervision.europa.eu/press/other-publications/annual-report/pdf/ssm.ar2022~e4b57f3b89.en.pdf.

European Central Bank. Annual report on supervisory activities 2023. Technical report, European Central Bank, Banking Supervision, March 2024. URL https://www.bankingsupervision.europa.eu/press/other-publications/annual-report/pdf/ssm.ar2023~2def923d71.en.pdf.

European Central Bank. Supervisory banking statistics, fourth quarter 2024. Technical report, European Central Bank, Banking Supervision, 2025a. URL https://www.bankingsupervision.europa.eu/press/other-publications/supervisorybankingstats/pdf/ssm.supervisorybankingstatistics_fourth_quarter_2024_202503.en.pdf.

European Central Bank. Annual report on supervisory activities 2024. Technical report, European Central Bank, Banking Supervision, March 2025b. URL https://www.bankingsupervision.europa.eu/press/other-publications/annual-report/html/ssm.ar2024~700cba1314.en.html.

European Central Bank. ECB guide to internal models (EGIM), consolidated version. Technical report, European Central Bank, Banking Supervision, 2025c. URL https://www.bankingsupervision.europa.eu/press/pr/date/2025/html/ssm.pr250728~2b36305822.en.html.

European Central Bank. Simplification of the European prudential regulatory, supervisory and reporting framework. Technical report, European Central Bank, December 2025d. URL https://www.ecb.europa.eu/pub/pdf/other/ecb.simplification_supervisory_reporting_framework202512.en.pdf. Report of the High-Level Task Force on Simplification, endorsed by the ECB Governing Council.

European Central Bank. Total annual supervisory fees, 2025e. URL https://www.bankingsupervision.europa.eu/about/fees/total/html/index.en.html. Accessed 4 March 2026.

European Central Bank. ECB list of managers, January 2026a. URL https://www.ecb.europa.eu/ecb/pdf/orga/ecb.ECB_managers.en.pdf. As of 15 January 2026.

European Central Bank. ECB streamlines how it supervises banks' internal models, March 2026b. URL https://www.bankingsupervision.europa.eu/press/pr/date/2026/html/ssm.pr260330~7f6b503e0c.en.html. Press release, 30 March 2026.

European Commission. Commission delegated regulation (EU) no 529/2014 supplementing regulation (EU) no 575/2013 with regard to regulatory technical standards for assessing the materiality of extensions and changes of the internal ratings based approach and the advanced measurement approach, March 2014. As amended by Commission Delegated Regulation (EU) 2022/439.

European Commission. Commission implementing regulation (EU) 2016/1799 laying down implementing technical standards with regard to the mapping of credit assessments of external credit assessment institutions for credit risk, 2016. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX:32016R1799.

European Commission. Targeted consultation on the competitiveness of the EU banking sector. Technical report, European Commission, Directorate-General for Financial Stability, Financial Services and Capital Markets Union, March 2026. URL https://finance.ec.europa.eu/regulation-and-supervision/consultations-0/targeted-consultation-competitiveness-eu-banking-sector-2026_en. Consultation closing 19 April 2026.

European Parliament and Council of the European Union. Regulation (EC) no 1060/2009 on credit rating agencies, 2009. URL https://eur-lex.europa.eu/eli/reg/2009/1060/oj/eng. Article 8(3) requires rating methodologies to be "rigorous, systematic, continuous and subject to validation based on historical experience, including back-testing".

European Parliament and Council of the European Union. Directive 2013/36/EU on access to the activity of credit institutions and the prudential supervision of credit institutions (capital requirements directive IV), 2013a. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:32013L0036. Articles 77–78 and 101 concern IRB permission, benchmarking and ongoing review.

European Parliament and Council of the European Union. Regulation (EU) no 575/2013 on prudential requirements for credit institutions (capital requirements regulation), 2013b. URL https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=celex:32013R0575. Part Three, Title II, Chapter 3 (Articles 142–191) governs the IRB approach.

European Parliament and Council of the European Union. Regulation (EU) 2024/1623 amending regulation (EU) no 575/2013 (capital requirements regulation III), 2024. CRR3.

European Securities and Markets Authority. Report on CRA market share calculation and performance statistics. Technical report, European Securities and Markets Authority, 2024. URL https://www.esma.europa.eu/document/cra-market-share-calculation-2024.

Eurosystem. Eurosystem response to the EU Commission's targeted consultation on the competitiveness of the EU banking sector. Technical report, European Central Bank, April 2026. URL https://www.ecb.europa.eu/press/consultationresponse/pdf/ecb.conresp202604.en.pdf. Includes Annex 1: HLTF report on simplification of the European prudential regulatory, supervisory and reporting framework.

Leonardo Gambacorta, Yiping Huang, Han Qiu, and Jingyi Wang. How do machine learning and non-traditional data affect credit scoring? New evidence from a Chinese fintech firm. BIS Working Papers 834, Bank for International Settlements, December 2019. URL https://www.bis.org/publ/work834.pdf.

Jens Grunert and Martin Weber. Recovery rates of commercial lending: Empirical evidence for german companies. *Journal of Banking & Finance*, 33(3):505–513, 2009. doi:10.1016/j.jbankfin.2008.09.002.

Andrew G. Haldane and Vasileios Madouros. The dog and the frisbee. Paper given at the Federal Reserve Bank of Kansas City's 36th economic policy symposium, Jackson Hole, Wyoming, August 2012. URL https://www.bankofengland.co.uk/paper/2012/the-dog-and-the-frisbee.

James A. Hanley and Barbara J. McNeil. The meaning and use of the area under a receiver operating characteristic (ROC) curve. *Radiology*, 143(1):29–36, 1982. doi:10.1148/radiology.143.1.7063747.

Hong Kong Monetary Authority. The use of digital innovation for reducing SMEs' credit constraints: Hong Kong SAR's recent experience. BIS Papers 148, ch. h, Bank for International Settlements, September 2024. URL https://www.bis.org/publ/bppdf/bispap148_h.pdf.

Hong Kong Monetary Authority and Hong Kong Applied Science and Technology Research Institute. Alternative credit scoring of micro-, small and mediumsized enterprises. White paper, Hong Kong Monetary Authority, November 2020. URL https://www.astri.org/wp-content/uploads/2016/11/White-paper-on-credit-scoring.pdf.

International Monetary Fund. Euro area: Publication of financial sector assessment program documentation – detailed assessment of observance of the Basel Core Principles for effective banking supervision. IMF Country Report 25/215, International Monetary Fund, July 2025. URL https://www.imf.org/-/media/files/publications/cr/2025/english/1eurea2025012.ashx.

Hinh D. Khieu, Donald J. Mullineaux, and Ha-Chin Yi. The determinants of bank loan recovery rates. *Journal of Banking & Finance*, 36(4):923–933, 2012. doi:10.1016/j.jbankfin.2011.10.005.

Stefan Lessmann, Bart Baesens, Hsin-Vonn Seow, and Lyn C. Thomas. Benchmarking state-of-the-art classification algorithms for credit scoring: An update of research. *European Journal of Operational Research*, 247(1):124–136, 2015. doi:10.1016/j.ejor.2015.05.030.

Mike Mariathasan and Ouarda Merrouche. The manipulation of Basel risk-weights. *Journal of Financial Intermediation*, 23(3):300–321, 2014. doi:10.1016/j.jfi.2014.04.004.

Carlos Pérez Montes, Carlos Trucharte Artigas, María Elizabeth Cristófoli, and Nadia Lavín San Segundo. The impact of the IRB approach on the risk weights of European banks. *Journal of Financial Stability*, 39:147–166, 2018. doi:10.1016/j.jfs.2016.11.004.

Matthew C. Plosser and João A. C. Santos. Banks' incentives and inconsistent risk models. *The Review of Financial Studies*, 31(6):2080–2133, 2018. doi:10.1093/rfs/hhy028.

Katja Pluto and Dirk Tasche. Estimating probabilities of default for low default portfolios. In *The Basel II Risk Parameters: Estimation, Validation, and Stress Testing*, pages 79–103. Springer, 2006. doi:10.1007/3-540-33087-9_5.

PricewaterhouseCoopers. IRB benchmarking survey 2022. Technical report, PwC UK, 2022. URL https://www.pwc.co.uk/financial-services/assets/pdf/irb-benchmarking-survey-2022.pdf.

Rafael Repullo and Javier Suarez. The procyclical effects of bank capital regulation. *The Review of Financial Studies*, 26(2):452–490, 2013. doi:10.1093/rfs/hhs118.

Risk.net. The collins flaw: Backstop turned binding constraint. *Risk.net*, 2022. URL https://www.risk.net/our-take/7928701/the-collins-flaw-backstop-turned-binding-constraint.

Isabel Schnabel. From laggard to leader? Closing the euro area's technology gap. Inaugural lecture of the ECB and its Watchers XXV conference, February 2024. URL https://www.ecb.europa.eu/press/key/date/2024/html/ecb.sp240216~df6f8d9c31.en.html. 16 February 2024.

Til Schuermann. What do we know about loss given default? Working Paper 04-01, Wharton Financial Institutions Center, 2004. URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=525702.

S&P Global Ratings. Default, transition, and recovery: 2024 annual global corporate default and rating transition study. Technical report, S&P Global Ratings, March 2025. URL https://www.spglobal.com/ratings/en/regulatory/article/250327-default-transition-and-recovery-2024-annual-global-corporate-default-and-rating-transition-study-s13452126.

Ross Taplin, Hien Minh To, and John Hee. Modeling exposure at default, credit conversion factors and the Basel II accord. Journal of Credit Risk, 3(2):75–84, 2007.

Dirk Tasche. The art of probability-of-default curve calibration. Journal of Credit Risk, 9(4):63–103, 2013. doi:10.21314/JCR.2013.169.

Edward N. C. Tong, Christophe Mues, Iain Brown, and Lyn C. Thomas. Exposure at default models with and without the credit conversion factor. European Journal of Operational Research, 252(3):910–920, 2016. doi:10.1016/j.ejor.2016.01.054.

Appendix

A.1 Concavity of the Risk-weight Function

The Basel IRB risk-weight function for corporate exposures (CRE 31) maps each borrower's probability of default *PD* into a risk weight *RW(PD)*, given a fixed loss given default *LGD* and maturity *M*. The function is defined as:

$$RW(PD) = LGD \cdot \left[\Phi\left(\frac{\Phi^{-1}(PD)}{\sqrt{1-R}} + \sqrt{\frac{R}{1-R}}\,\Phi^{-1}(0.999) \right) - PD \right] \cdot MA(PD) \cdot 12.5 \quad (2)$$

where Φ denotes the standard normal cumulative distribution function and $R = R(PD)$ is the asset correlation (a decreasing function of *PD*). The maturity adjustment is:

$$MA(PD) = \frac{1 + (M - 2.5) \cdot b(PD)}{1 - 1.5 \cdot b(PD)}, \qquad b(PD) = \left(0.11852 - 0.05478 \cdot \ln PD\right)^2 \quad (3)$$

This function is strictly concave in *PD* on the entire domain (0, 1).[24] By Jensen's inequality, for any concave function f and any set of values $x_1, \ldots, x_n$ with weights w_i summing to one:

$$\sum_i w_i\, f(x_i) \;\le\; f\left(\sum_i w_i\, x_i \right) \quad (4)$$

Applied to the IRB risk-weight function, this means that the weighted average of the risk weights computed at each borrower's individual PD is strictly lower than the risk weight evaluated at the portfolio's average PD:

$$\overline{RW} \;=\; \frac{1}{n} \sum_{i=1}^{n} RW(PD_i) \;<\; RW\left(\bar{PD}\right) \quad (5)$$

In other words, differentiating borrowers by their individual default probabilities produces lower aggregate capital requirements than treating them all at a single pooled *PD*, even in the absence of any conservatism adjustment. Figure 2 illustrates this property for two borrowers with PD_1 = 0.5% and PD_2 = 5% (LGD = 45%, *M* = 2.5 years). The average of the two individual risk weights (109.7%) is 15.7 percentage points below the risk weight evaluated at the average PD (125.4%).

24 Strict concavity can be verified numerically; it holds for both the core Vasicek term and the full riskweight function including the PD-dependent asset correlation and maturity adjustment. For a derivation of the function see Basel Committee on Banking Supervision [2005b].

Figure 2: Concavity of the IRB risk-weight function and Jensen's inequality. The solid curve shows *RW(PD)*. The dashed chord connects two borrowers. At the average PD, the average of the individual risk weights (square, on the chord) lies below the risk weight at the average PD (triangle, on the curve).

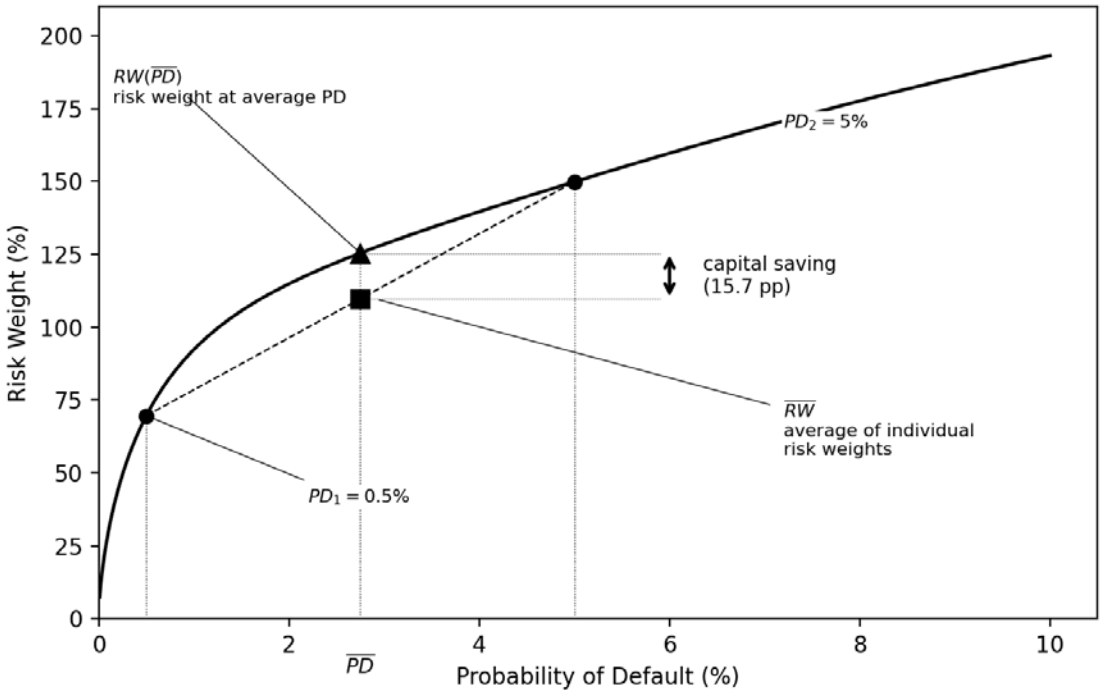

Source: World Bank.

A.2 Quantifying Financial Incentives to Underestimate PD

This appendix quantifies the financial incentive for banks to underestimate PD estimates under the IRB framework. The calculation uses the standard Basel risk-weight function for corporate exposures (CRE 31) with an LGD of 45 percent (the F-IRB supervisory value for unsecured exposures) and a maturity of 2.5 years.

The corporate formula is the most conservative baseline for this exercise because its maturity adjustment partially dampens sensitivity to PD changes. For residential mortgage and other retail exposures, where the Basel framework prescribes lower asset correlations and no maturity adjustment, the relative risk-weight impact of a given PD reduction is materially larger, as noted below.

Assumptions. We consider a hypothetical bank with EUR 100 billion of credit risk RWA, a total capital ratio of 15 percent, a cost of equity of 10 percent, and a weighted average cost of funding of 3 percent. The relevant measure of capital savings is the equity-debt spread (7 percentage points), since freed regulatory capital is replaced by cheaper debt funding on the liability side. Profitability parameters are drawn from EU-wide averages as reported in the EBA's December 2025 Risk Assessment Report [European Banking Authority, 2025c]: return on assets of 0.75 percent, return on equity of 10.7 percent, and a cost-to-income ratio of 52.3 percent. We assume an average risk-weight density of 40 percent, implying total assets of EUR 250 billion.

Impact of a 10 basis point PD reduction. We use a base PD of 1.0 percent. For reference, the EBA's 2024 Credit Risk Benchmarking Exercise reports EAD-weighted average PDs of approximately 0.5 percent for large corporates, 1.3 percent for corporates and 1.5 percent for SME corporates [European Banking Authority, 2025a]. Using the IRB risk-weight formula, a uniform reduction from PD = 1.0 percent to PD = 0.9 percent reduces the risk weight from 92.3 percent to 88.8 percent, a decline of 3.5 percentage points or 3.8 percent in relative terms. Applied to EUR 100 billion of credit risk RWA, this frees approximately EUR 570 million of regulatory capital, saving roughly EUR 40 million per year (the freed capital multiplied by the 7 percentage point equitydebt spread).

In context, this annual saving represents approximately 2 percent of the bank's net income (EUR 1.9 billion) and 1 percent of total revenue (EUR 3.9 billion). For a large bank with EUR 500 billion of credit risk RWA, the corresponding figures are EUR 2.8 billion of freed capital and approximately EUR 195 million per year.

Sensitivity to PD level. Because of the concavity of the IRB risk-weight function as shown on Figure 2, the incentive to underestimate PDs is markedly stronger at lower PD levels, as shown in Table 2.

Table 2: Risk-weight impact of a 10 basis point PD reduction (LGD = 45%)

Base PD	Risk-weight (base)	Risk-weight (−10bp)	Risk-weight reduction	Relative drop
0.25%	49.5%	37.4%	12.0 pp	24.4%
0.50%	69.6%	62.7%	6.9 pp	9.9%
1.00%	92.3%	88.8%	3.5 pp	3.8%
2.00%	114.9%	113.2%	1.7 pp	1.4%
5.00%	149.9%	148.9%	1.0 pp	0.7%

At PD = 0.25 percent (a typical investment-grade corporate), a 10 basis point reduction lowers the risk weight by 24.4 percent, nearly seven times the relative impact at PD = 1 percent. This is precisely the segment of the market where competitive pressure among banks is highest.

As noted above, these figures understate the incentive for retail portfolios. Under the residential mortgage formula (asset correlation R = 0.15, no maturity adjustment), the same 10 basis point reduction at PD = 0.25 percent lowers the risk weight by 31.7 percent. Under the qualifying revolving retail formula (R = 0.04), the relative drop reaches 34.2 percent. The financial incentive to underestimate PDs is therefore strongest precisely where IRB banks hold their largest exposures.

A.3 IRB-related Regulatory and Supervisory Texts

Table 3 provides an inventory of the principal regulatory and supervisory texts that govern the use of internal models for credit risk in the European Union. The table illustrates the sheer volume of prescriptive material that institutions and supervisors must navigate to operate and oversee the IRB framework.

Table 3: Principal IRB-related regulatory and supervisory texts in the EU

Document	Authority	Year	Pages
Primary legislation			
CRR Part Three, Title II, Chapter 3 (Articles 142–191): Internal Ratings Based Approach	EU	2013	50
CRD Articles 77–78, 101: Permission, benchmarking and ongoing review of internal approaches	EU	2013	3
EBA Regulatory Technical Standards			
RTS on assessment methodology for IRB approach (EBA/RTS/2016/03)	EBA	2016	167
RTS on specialized lending exposures (EBA/RTS/2016/02)	EBA	2016	74
RTS on materiality of model extensions and changes (CDR 529/2014)	EBA	2014	14
RTS on economic downturn (EBA/RTS/2018/04)	EBA	2018	37
EBA Guidelines			
GL on PD estimation, LGD estimation and treatment of defaulted exposures (EBA/GL/2017/16)	EBA	2017	57
GL on downturn LGD estimation (EBA/GL/2019/03)	EBA	2019	80
GL on credit risk mitigation for A-IRB institutions (EBA/GL/2020/05)	EBA	2020	103
CP on GL on CCF estimation (EBA/CP/2025/10)	EBA	2025	114
EBA other publications			
Supervisory handbook for the validation of IRB rating systems	EBA	2023	96
ECB Banking Supervision			
ECB Guide to Internal Models (EGIM), consolidated version	ECB	2025	370
Total			**1,165**

The total of over 1,000 pages covers only the principal texts listed above. It excludes the 53 IRB-related Q&As published by the EBA (68 pages in aggregate), national supervisory expectations and guidance issued by individual competent authorities, the annual EBA benchmarking reports, and the various consultation papers and impact assessments that preceded each final document. A conservative estimate of the total IRB-specific regulatory and supervisory corpus, including these additional materials, would exceed 1,500 pages.[25]

25 Page counts refer to the published PDF versions of each document. Where documents have been revised, the page count refers to the most recently published version as of early 2026.

A.4 Advisory Firms with Dedicated IRB and Credit Risk Modeling Practices

Table 4 lists the principal advisory firms that maintain dedicated teams for IRB credit risk modeling, model validation, and regulatory compliance in Europe. The list, while not exhaustive, illustrates the breadth of the external consulting ecosystem that has developed around the IRB framework.

Table 4: Advisory firms with dedicated IRB and credit risk modeling practices in Europe

Firm	IRB-related services
Major consulting and audit firms	
Oliver Wyman	IRB model development, calibration, and analytics
McKinsey	Credit risk modeling, rating systems, risk function transformation
KPMG	IRB model development and validation, regulatory compliance
Deloitte	Credit risk framework design, model risk management advisory
EY	Quantitative advisory for IRB, model validation and governance
PwC	IRB credit risk modeling, model governance, EGIM compliance
Accenture	IRB model validation, RWA computation frameworks
Specialist and quantitative consulting firms	
d-fine	Quantitative modeling and validation for credit risk
Finalyse	Dedicated IRB model validation; PD, LGD, EAD toolkits
Prometeia	Credit risk modeling, model development and validation platform
Zanders	Independent risk model validation, credit risk assessment
BearingPoint	Basel, IRB, IFRS 9, and stress testing advisory
Capgemini	Credit risk process redesign, advanced analytics
Forvis Mazars	IRB and IFRS 9 model development, quantitative finance

A.5 Estimating the Aggregate Cost of IRB Compliance

This appendix estimates the total annual cost of maintaining the IRB framework across the European banking industry, using three independent approaches that converge on a central range of EUR 0.7–1.2 billion. The estimate covers both bank-side and supervisor-side costs. IT infrastructure costs are excluded: banks would need robust credit risk data systems regardless of whether they use internal models, so these costs cannot be attributed to the IRB framework specifically.

IRB population. The precise number of IRB institutions in the EU is not publicly reported as a single figure. However, the ECB's supervisory statistics indicate that the majority of the roughly 110 significant institutions under direct ECB supervision use the IRB approach for at least part of their credit portfolio [European Central Bank, 2025a]. A number of less significant institutions and non-SSM banks hold IRB permissions as well. Approximately 90 institutions use the IRB approach for a material share of their credit exposures: roughly 12 large A-IRB banks operating across multiple jurisdictions,[26] 25 mid-sized institutions using either A-IRB or F-IRB, and 50 smaller F-IRB banks.[27]

Bottom-up FTE estimation. The starting point is PwC's UK IRB Benchmarking Survey, which reports that 62 percent of IRB banks employ 5–10 FTEs for retail model development alone, with validation teams of one to four FTEs at the majority of institutions [PricewaterhouseCoopers, 2022]. The Deloitte EMEA Model Risk Management Survey provides supporting context: large banks maintain an average of 365 models across all model types, of which credit risk models form the largest single category [Deloitte, 2025]. However, the PwC figures cover only retail portfolios in a single jurisdiction. A large European A-IRB bank typically operates IRB models across multiple countries (each requiring separate PD model calibrations), multiple asset classes (retail mortgages, revolving credit, other retail, corporate, SME, sovereign, bank and financial institution exposures, and in many cases project finance and securitization portfolios), and multiple risk parameters (PD, LGD, and EAD, plus application and behavioral scoring models). Each model requires development, independent validation, and ongoing monitoring, with annual recalibration and periodic redevelopment cycles.

26 BNP Paribas, Crédit Agricole, Santander, Société Générale, Deutsche Bank, ING, UniCredit, Nordea, Danske Bank, KBC, Erste Group, and Groupe BPCE. All operate Advanced IRB models across three or more EU/EEA jurisdictions.

27 The estimate of 90 material IRB users is conservative. If less significant institutions with partial IRB permissions are included, the total may exceed 100.

Scaling the PwC baseline for this broader scope, and adding governance and data infrastructure functions, yields the following FTE estimates by bank tier:

Table 5: Estimated IRB-dedicated FTEs by bank tier

Bank tier	Count	FTE per bank	Total FTE	Annual cost (million)
Large A-IRB (multi-country)	12	60–80	720–960	EUR 115–155
Mid-sized A-IRB or large F-IRB	25	25–35	625–875	EUR 100–140
Smaller F-IRB	50	6–12	300–600	EUR 50–95
Total	**~90**		**1,645–2,435**	**EUR 265–390**

The all-in cost per FTE is assumed at EUR 160,000 for major European financial centers (London, Frankfurt, Paris, Amsterdam), reflecting an average base salary of EUR 85,000, bonus of 30 percent, employer social contributions, and allocated overhead.[28] This covers model development, independent validation, monitoring, recalibration, governance, and data infrastructure staff. It does not include external consulting, which is estimated separately below.

Top-down cross-check: consulting market. The breadth of the external consulting ecosystem documented in Appendix A.4 provides a second, independent check. Every major advisory firm maintains dedicated IRB practices in Europe, and specialist firms have built their business models around IRB compliance. While aggregate spending data are not publicly available, a conservative assumption of EUR 2–5 million per year per IRB bank (covering both ongoing advisory retainers and project-based work for model redevelopments, validation, and regulatory examinations) implies total external consulting costs of approximately:

≈ EUR 200–500 million per year

Cross-check: TRIM deficiency analysis. The ECB's Targeted Review of Internal Models reviewed 200 credit risk models across 65 significant institutions and identified more than 5,800 deficiencies [European Central Bank, 2021], an average of 29 per model. The resulting aggregate RWA increase of EUR 275 billion (12 percent of the reviewed portfolio) confirms that model maintenance is not a trivial exercise. The high deficiency rate validates the allocation of approximately 10–15 percent of total IRB FTEs to ongoing monitoring and remediation, consistent with the bottom-up estimate.

28 For smaller financial centers, the cost per FTE may be 10–15 percent lower, which would reduce the aggregate by EUR 25–60 million without changing the order of magnitude.

Supervisor-side costs. The ECB's annual supervisory fee, which covers its own expenditure on banking supervision (but not the costs borne by national competent authorities), was EUR 681 million in 2024 [European Central Bank, 2025b,e]. While no public breakdown by function exists, internal model oversight absorbs a significant share of total supervisory resources. A useful reference point is the TRIM project: in 2017, the ECB disclosed that approximately 15 percent of the total SSM budget would be invested in TRIM alone, making it the largest single project investment in ECB Banking Supervision history [European Central Bank, 2017]. Applying a range of 10–15 percent to the 2024 ECB supervisory fee implies EUR 70–100 million for the ECB's IRB-related supervisory costs in normal operations. Adding the IRB-related supervisory costs of the 21 national competent authorities in the euro area (which bear their own costs for participating in Joint Supervisory Teams and for supervising less significant institutions) and the competent authorities of non-SSM EU Member States, total supervisor-side costs across the EU are approximately:

≈ EUR 200–300 million per year

Total estimate. Combining the three components:

Table 6: Estimated annual cost of IRB compliance in the EU

Component	Estimated annual cost
Bank internal staff (1,650–2,450 FTEs)	EUR 265–390 million
External consulting	EUR 200–500 million
Supervisor-side costs	EUR 200–300 million
Total	**EUR 0.7–1.2 billion**

The three estimation approaches converge on a central range of approximately EUR 0.7–1.2 billion per year. Even under conservative assumptions (fewer IRB banks, lower FTE counts, lower cost per FTE), the total approaches EUR 700 million; under more expansive assumptions (including partial IRB users and regulatory-driven cost escalation from evolving EBA guidelines), it exceeds EUR 1 billion.

This estimate excludes several additional cost categories: the management attention absorbed by IRB governance at senior level, the regulatory affairs and policy specialists that banks employ to engage with evolving EBA and ECB standards, and the cost of the regulatory apparatus itself (EBA standard-setting, consultation processes, benchmarking exercises). The true all-in cost of the IRB framework to the European economy is therefore likely to be higher than the figures reported above.

A.6 Structural LGD Formula

This appendix provides the detailed equations and variable definitions for the structural LGD approach proposed in Section 2.

Notation. For a given exposure with outstanding claim E at the date of default:

- C = recorded collateral value at the most recent valuation date (updated per CRR Article 208), capped at E
- h = supervisory haircut for the relevant collateral type (EU-wide, set by the EBA)
- T = bank-specific average time-to-foreclosure for the relevant collateral type (in years), computed over a rolling window of 7–10 years
- r = discount rate, defined as $r = \max(r_{loan}, r_{floor})$, where:
  - r_{loan} = the loan's effective interest rate (as defined under IFRS 9, incorporating fees, commissions, and charges) plus any contractual default interest surcharge. For variable-rate loans, r_{loan} is the effective rate as fixed at the date of default.
  - r_{floor} = the supervisory floor rate, set as €STR + a fixed add-on, calibrated to approximate the average cost of bank funding
- R_u = the bank's long-run average observed recovery rate on fully unsecured exposures in the same portfolio, computed over a minimum of 10 years (see below for derivation)
- LGD_{min} = minimum LGD floor of 10 percent (applying to all exposures regardless of collateralization)

Exposure at default. The denominator E is the outstanding principal plus any accrued unpaid interest at the moment of default recognition.

Recovery value on the collateralized portion. The present value of the collateral recovery is:

$$V_{secured} = \frac{C \cdot (1 - h)}{(1 + r)^T} \tag{6}$$

This is the recorded collateral value, reduced by the supervisory haircut and discounted by the time-to-foreclosure.

Recovery value on the uncollateralized portion. The recovery value on the unsecured share is the uncollateralized amount multiplied by the bank's long-run average recovery rate on fully unsecured loans in the same portfolio:

$$V_{\text{unsecured}} = (E - C) \cdot R_u \tag{7}$$

where R_u is derived from the bank's own workout history. For each resolved default j in the sample (using only closed recovery processes over a minimum of 10 years), the recovery rate is the present value of all cash flows received during the workout, discounted at the same rate r defined above:

$$R_j = \frac{1}{E_j} \sum_{k=1}^{K_j} \frac{CF_{j,k}}{(1 + r_j)^{t_{j,k}}} \tag{8}$$

where $r_j = \max(r_{\text{loan},j}, r_{\text{floor}})$ is the discount rate for loan j, $CF_{j,k}$ is the k-th cash flow recovered on defaulted exposure j, $t_{j,k}$ is the time in years from default to that cash flow, and E_j is the exposure at default. The portfolio-level recovery rate R_u is the average across all resolved defaults.

Overall LGD. The exposure-level LGD is one minus the total recovery value as a share of the exposure, subject to a floor:

$$LGD = \max\left(1 - \frac{V_{\text{secured}} + V_{\text{unsecured}}}{E},\ LGD_{\text{min}}\right) \tag{9}$$

Numerical example. Consider a EUR 1 million corporate loan with EUR 600,000 of commercial real estate collateral (C = 600,000). The bank operates in a jurisdiction where the average time-to-foreclosure for CRE is 3.5 years. The supervisory haircut for CRE is 15 percent. The loan's effective interest rate including penalty interest is 8 percent, and the supervisory floor rate is 4.8 percent (€STR of 2.6% + 220 bp). The bank's long-run average recovery rate on fully unsecured corporate loans is 38 percent. The minimum LGD floor is 10 percent.

Since $r_{loan} > r_{floor}$, the discount rate is $r = 8\%$.

$$V_{secured} = \frac{600{,}000 \times (1 - 0.15)}{(1 + 0.08)^{3.5}} = \frac{510{,}000}{1.309} = 389{,}600 \tag{10}$$

$$V_{unsecured} = 400{,}000 \times 0.38 = 152{,}000 \tag{11}$$

$$LGD = 1 - \frac{389{,}600 + 152{,}000}{1{,}000{,}000} = 1 - 0.542 = 45.8\% \tag{12}$$

This compares with the F-IRB supervisory LGD of 45 percent for an unsecured exposure, or, for a comparable secured exposure with CRE collateral, 35 percent (after applicable haircuts under F-IRB). The structural approach produces an intermediate value that reflects the specific collateral coverage, judicial environment, and unsecured recovery experience of the bank.

A.7 The Logistic PD Model and the Bank-specific Constant

This appendix sets out the formal specification of the centrally calibrated logistic PD model proposed for high-default portfolios, and shows how each bank recalibrates its constant (the intercept of the logistic regression) to its own long-run average default rate.

Model specification. Let $x_i = (x_{i1}, x_{i2}, \ldots, x_{iK})$ denote the vector of K explanatory variables for borrower i (for example, debt-service-to-income ratio, loan-to-value ratio, seasoning, delinquency status). The logistic model predicts the probability of default as:

$$PD_i = \frac{1}{1 + \exp\left(-\left(\beta_0^{(b)} + \beta_1 x_{i1} + \beta_2 x_{i2} + \cdots + \beta_K x_{iK}\right)\right)} \tag{13}$$

where $\beta_1, \ldots, \beta_K$ are the slope coefficients estimated centrally by the regulator on pooled data, and $\beta_0^{(b)}$ is the bank-specific constant for bank b. The slope coefficients determine how each explanatory variable affects the relative ranking of borrowers: a higher debt-service-to-income ratio, for instance, shifts the predicted PD upward by the same amount for all banks. The bank-specific constant shifts the entire PD distribution up or down without altering the relative ordering of borrowers within the bank's portfolio.

Central calibration. The regulator estimates the slope coefficients $\hat{\beta}_1, \ldots, \hat{\beta}_K$ by maximum likelihood on a pooled dataset comprising anonymized loan-level observations from all participating banks. The pooling maximizes the number of observed defaults available for estimation and ensures that coefficients reflect system-wide relationships between borrower characteristics and default risk. The estimation is repeated every five years to incorporate new data and structural changes in lending markets.

Bank-specific recalibration. Each bank b recalibrates its constant $\beta_0^{(b)}$ annually by solving:

$$\frac{1}{N_b}\sum_{i=1}^{N_b}\frac{1}{1+\exp\left(-\left(\beta_0^{(b)}+\hat{\beta}_1 x_{i1}+\hat{\beta}_2 x_{i2}+\cdots+\hat{\beta}_K x_{iK}\right)\right)} = LRADR_b \qquad (14)$$

where N_b is the number of borrowers in bank b's current portfolio and LRADR_b is the bank's long-run average default rate, computed over a minimum of ten years. Because the left-hand side is strictly increasing in $\beta_0^{(b)}$ (the logistic function is monotone), a unique solution exists and can be found by standard numerical methods.

Interpretation. The constant absorbs the difference between the portfolio-average PD predicted by the explanatory variables alone and the bank's observed long-run average default rate. If the centrally estimated variables fully explain why one bank's borrowers default more frequently than another's (for example, because they carry higher leverage or higher debt-service ratios), the constants will be identical across banks. The constant differs only to the extent that a bank's default experience is not explained by the common explanatory variables, for instance because of differences in origination practices, underwriting culture, or portfolio characteristics that the standardized variable set does not capture. The mechanism is automatic and auditable: given the published coefficients and the bank's default history, any competent authority can verify the constant independently.

What the bank controls and what it does not. Under this specification, the bank has no influence over the model structure, the choice of explanatory variables, or the slope coefficients. Its only degree of freedom is the constant, which is mechanically determined by the bank's own observed default rate. This closes the principal channel of the agency problem at the heart of the current IRB framework: the bank's ability to reduce capital requirements by choosing model specifications, estimation samples, and calibration assumptions that produce lower PD estimates. Under the proposed approach, the only way to achieve lower capital requirements is to lend to borrowers who default less frequently over the long run.

A.8 Statistical Limitations of Discriminatory Power Assessment

This appendix demonstrates that the discriminatory power of IRB rating systems cannot be reliably assessed given the default counts observed in realistic banking portfolios. The implication is that the claimed superiority of IRB granularity over a well-designed standardized approach is not an empirically demonstrated fact but an article of faith: numerical precision masquerading as statistical evidence.

The discrimination question. A credit rating system assigns obligors to discrete grades, each associated with a PD. Discriminatory power refers to the ability of the system to rank obligors such that those who subsequently default were assigned worse grades ex ante. The standard measure is the Area Under the ROC Curve (AUC), equivalent to a concordance probability estimable via the Mann-Whitney U statistic [Hanley and McNeil, 1982]:

$$\hat{\theta} = \frac{1}{n_D \cdot n_{ND}} \sum_{i \in D} \sum_{j \in ND} \left[\mathbf{1}(g_i > g_j) + \tfrac{1}{2} \mathbf{1}(g_i = g_j) \right] \tag{15}$$

where D denotes the set of defaulters, ND the set of non-defaulters, g is the grade assignment, and higher grades correspond to higher PDs. Under no discriminatory power $\theta = 0.5$; a perfect model gives $\theta = 1$. The Gini coefficient relates to AUC via Gini = $2\theta - 1$.

Precision under sparse defaults. Consider the minimal discrimination claim: two adjacent grades with PDs p and p + δ, each containing $N/2$ obligors, with k total defaults. The number of defaults in the higher-PD grade follows a hypergeometric distribution under the null hypothesis of no discrimination. The standard error of the AUC estimator simplifies to $SE(\hat{\theta}) \approx 1/(2\sqrt{k})$, giving a 95 percent confidence interval of approximately ±0.98/√k [DeLong et al., 1988]. Table 7 reports the resulting confidence intervals for realistic default counts.

Table 7: Precision of AUC estimates by number of observed defaults

Expected defaults k	95% CI half-width on AUC	95% CI half-width on δ/p
10	±0.31	±316%
25	±0.20	±200%
100	±0.10	±100%
400	±0.05	±50%
2500	±0.02	±20%

With 25 defaults, an observed AUC of 0.70 is consistent with a true AUC anywhere in [0.50, 0.90]: the observation is uninformative about whether the model has genuine discriminatory power. The fine PD gradations claimed by IRB master scales (typically 20–30 percent relative steps between adjacent grades) require approximately 2,500 defaults to be statistically detectable. This threshold exceeds the total number of defaults observed across most European banking portfolios over a full credit cycle.

Grade granularity does not help. The number of rating grades G is a reporting and bucketing choice, not a discriminatory power choice. The true discriminatory power resides in the underlying continuous score; discretizing it into G grades can only lose information. For a discrete system with G equally-sized grades and k defaults, $\mathrm{Var}(\hat{\theta}) = (G^2 - 1)/(12\, G^2\, k)$. Going from 2 grades to 7 recovers 99 percent of the precision achievable with a continuous score. Even replacing the master scale with the raw logit output for backtesting narrows the confidence interval by a factor of only 1.16. The fundamental unverifiability is not a discretization artefact; it is intrinsic to the scarcity of defaults [Basel Committee on Banking Supervision, 2005c].

Calibration versus discrimination. AUC is entirely invariant to the absolute PD levels assigned to each grade: a model with AUC = 0.85 can be wildly miscalibrated. This means that even a statistically significant AUC tells us nothing about whether the grade-level PDs are correct. Conversely, grade-level calibration backtesting (comparing predicted to realized default rates per grade) faces the same $1/\sqrt{k}$ precision constraint applied to even smaller within-grade samples. Banks calibrate to the long-run average default rate (a through-the-cycle concept) but annual backtests compare against point-in-time realized rates, guaranteeing a mismatch in every year that is unfalsifiable by construction.

The conditional likelihood alternative. A more powerful test conditions on the total number of defaults and asks whether the model assigns high PDs to precisely the obligors who defaulted:

$$P(D = \text{observed} \mid |D| = k) = \frac{\prod_{i \in D} p_i \prod_{j \notin D}(1 - p_j)}{\sum_{S:|S|=k} \prod_{i \in S} p_i \prod_{j \notin S}(1 - p_j)} \tag{16}$$

This statistic uses the PD levels (not just ranks) and is therefore more informative than AUC when the PD spread across the portfolio is wide. However, its precision also scales as $1/\sqrt{k}$: no choice of statistic escapes the constraint imposed by default scarcity. Moreover, the conditional likelihood is most advantageous relative to AUC precisely in portfolios with wide PD spreads, where discrimination is easiest to demonstrate anyway. For the hard cases (homogeneous investment-grade portfolios where the precision problem is most acute) it offers essentially no improvement [Engelmann and Rauhmeier, 2006].

In-sample versus out-of-sample. In-sample AUC is an upward-biased estimator of true discriminatory power and should never serve as validation evidence. Out-of-sample testing (holdout, time-based, or walk-forward cross-validation) removes this bias but does not resolve the precision problem: with k out-of-sample defaults, the confidence interval retains the same width. An out-of-sample AUC of 0.70 with 30 validation defaults carries a confidence interval of roughly [0.52, 0.88], entirely consistent with a model whose true discriminatory power does not exceed the regulatory minimum [Blöchlinger and Leippold, 2006].

The within-cell heterogeneity claim. IRB advocates sometimes argue that standardized approach risk-weight tables leave meaningful within-cell heterogeneity unpriced. This claim is the discrimination hypothesis restated at the level of within-SA-cell residual variation, and it requires the same statistical demonstration, now applied to sub-samples with even fewer defaults. Without statistical evidence, the within-cell heterogeneity argument is circular: the model's internal logic is presented as evidence for the model's accuracy.

Policy implication. The discriminatory power of IRB models cannot be statistically demonstrated for the portfolios (low-default, investment-grade) where IRB is most extensively applied and where the capital impact of PD differentiation is greatest (Appendix A.2). A well-designed standardized approach table encoding two or three credit-relevant risk drivers captures genuine information about default risk. Whether a bank's internal model captures this information more accurately than the centrally calibrated alternative is exactly the question that low-power backtesting cannot answer. The burden of proof rests with IRB proponents to demonstrate superiority using honest statistical standards; for many portfolios, this demonstration cannot be made.

Precision of observed default rates. The same scarcity problem applies to the estimation of default rates themselves. When a portfolio of n exposures is observed over T years, the total number of defaults k in the nT exposure-years follows a Poisson distribution with parameter $\lambda = nT \cdot p$, where p is the true annual PD. The observed default rate is $\hat{p} = k/(nT)$, and an approximate 95 percent confidence interval for p is:

$$\left[\frac{\chi^2_{2k,\,0.025}}{2\,nT} \,,\; \frac{\chi^2_{2(k+1),\,0.975}}{2\,nT} \right] \tag{17}$$

where $x^2_{v,\alpha}$ denotes the α-quantile of the chi-squared distribution with v degrees of freedom. For a portfolio of 1,000 exposures observed over 20 years (nT = 20,000) with a true PD of 5 basis points, the expected number of defaults is $\lambda = 10$. The resulting 95 percent confidence interval for the default rate spans approximately 2 to 9 basis points, covering two full external rating grades (for reference, long-run average default rates are roughly 1 basis point for AAA/AA, 3 basis points for AA/A, and 6 basis points for single-A; S&P Global Ratings, 2025).

Even doubling the observation window to 40 years, or the portfolio to 2,000 exposures, reduces the interval only by a factor of √2. At these PD levels, the data cannot distinguish a AAA borrower from a single-A borrower with conventional statistical confidence.

A.9 Illustrative Granular standardized approach

This appendix illustrates the granular standardized approach discussed in the assessment of alternative approaches. The construction starts from the same foundation as this paper's proposal: every five years, the regulator calibrates a credit risk model on pooled loan-level data submitted by banks. The result, however, is published in a different form. Instead of releasing model coefficients to which each bank adds its own constant, the regulator identifies the two risk drivers with the highest explanatory power in each portfolio, partitions each driver into a small number of buckets, and publishes for every cell of the resulting grid the average risk weight obtained from equation (1) across the borrowers in that bucket. Banks assign each exposure to a cell on the basis of two observable characteristics and read the risk weight directly from the table. Because the LGD embedded in each cell is a supervisory value, reading a risk weight from the table and reading a supervisory value for *f (PD)* are equivalent; the tables below show risk weights, which permits a direct comparison with the standardized approach.

Table 8 presents illustrative tables for two asset classes. The first replaces the flat 100 percent risk weight that the standardized approach assigns to unrated corporates. Its dimensions are leverage, measured as net debt over EBITDA, and delinquency status. The embedded LGD is the supervisory senior unsecured value of 40 percent used in the F-IRB approach, with an effective maturity of 2.5 years; exposures secured by eligible collateral would instead be treated under the credit risk mitigation framework, as under the standardized approach today. The second table covers residential mortgages, with dimensions debt-service-to-income ratio (DSTI) and loan-to-value ratio (LTV). Here the LTV serves a double function: it is a driver of default probability, and it determines the embedded LGD, which rises from 10 percent for the lowest LTV bucket to 30 percent for loans above 100 percent LTV. This mirrors the logic of the standardized approach, whose capital charge for mortgages rises with LTV for the same reason: directly through the LTV-based risk-weight table under the Basel whole-loan approach, or through the growing unsecured portion under the loan-splitting approach implemented in the CRR. Cell values are subject to the 15 percent risk-weight floor that this paper's proposal applies. Exposures more than 90 days past due are in default and drop out of the tables; they would retain the defaulted-exposure treatment of the standardized approach.

Table 8: Illustrative risk-weight tables for a granular standardized approach. Cell values are risk weights in percent, computed from equation (1) at illustrative bucket-average PDs and supervisory LGD values, rounded to the nearest five percentage points. Italicized cells are at the 15 percent risk-weight floor.

Unrated corporates (standardized approach: flat 100%)

	Days past due		
Net debt/EBITDA	None	1–30	31–90
≤ 1×	50	95	135
1–2.5×	65	110	160
2.5–4×	85	125	185
> 4×	110	150	205

Residential mortgages (standardized approach: 20%/75% loan split)

	Loan-to-value ratio (%)			
DSTI (%)	≤ 60	60–80	80–100	> 100
< 20	*15*	*15*	*15*	25
20–30	*15*	*15*	20	40
35–50	*15*	20	35	60
> 50	20	35	50	90

Notes: All values are illustrative and not calibrated on supervisory data. Corporates: assumed bucket-average PDs rise from 0.3 percent (lowest leverage, no arrears) to 18 percent (highest leverage, 31–90 days past due); LGD fixed at 40 percent; maturity 2.5 years. Mortgages: assumed bucket-average PDs rise from 0.2 percent to 4.3 percent with DSTI and LTV; embedded LGD of 10, 15, 20, and 30 percent across the four LTV buckets.

The values are meaningful when set against the current standardized approach. For unrated corporates, the flat 100 percent risk weight is replaced by a range from 50 percent for low-leverage performing borrowers to 205 percent for highly leveraged borrowers in arrears, restoring the risk differentiation that the flat treatment suppresses while remaining anchored to the same average level. For residential mortgages, the standardized approach already differentiates by LTV: under the loan-splitting approach of the CRR, the portion of the exposure up to 55 percent of the property value receives a 20 percent risk weight and the remainder the 75 percent retail risk weight, yielding an effective whole-loan risk weight that rises from 20 percent for low-LTV loans to roughly 45 percent at an LTV of 100 percent. The granular table extends this differentiation to affordability, so that a high-LTV loan to a stretched borrower attracts a materially higher risk weight than the same loan to a borrower with ample debt-service capacity. An actual calibration would be produced by the EBA from pooled supervisory data, with bucket boundaries and cell values set so that each cell contains a sufficiently homogeneous and sufficiently populated set of exposures.

A.10 List of Abbreviations

A-IRB	Advanced Internal Ratings-Based approach
APRA	Australian Prudential Regulation Authority
ASRF	Asymptotic Single Risk Factor (model)
AUC	Area Under the Receiver Operating Characteristic Curve
BCBS	Basel Committee on Banking Supervision
CCF	Credit Conversion Factor
CEREP	Central Repository (ESMA repository for credit rating agency performance data)
CQS	Credit Quality Step
CRA	Credit Rating Agency
CRD	Capital Requirements Directive
CRR	Capital Requirements Regulation
EAD	Exposure at Default
EBA	European Banking Authority
ECAI	External Credit Assessment Institution
ECB	European Central Bank
ESMA	European Securities and Markets Authority
€STR	Euro Short-Term Rate
EU	European Union
F-IRB	Foundation Internal Ratings-Based approach
FINMA	Swiss Financial Market Supervisory Authority
FSA	Financial Services Agency (Japan)
FTE	Full-Time Equivalent
GAAP	Generally Accepted Accounting Principles
HKMA	Hong Kong Monetary Authority
IFRS	International Financial Reporting Standards
IRB	Internal Ratings-Based (approach)
LGD	Loss Given Default
LRADR	Long-Run Average Default Rate
MAS	Monetary Authority of Singapore
OCC	Office of the Comptroller of the Currency (United States)
OSFI	Office of the Superintendent of Financial Institutions (Canada)
PD	Probability of Default
PRA	Prudential Regulation Authority (United Kingdom)
RCAP	Regulatory Consistency Assessment Program (Basel Committee)
RW	Risk Weight
RWA	Risk-Weighted Assets
SA	SA standardized approach
SME	SME Small and Medium-Sized Enterprise(s)
TRIM	Targeted Review of Internal Models (ECB)

THE WORLD BANK
IBRD • IDA | WORLD BANK GROUP