

**ABI response to the consultation on EBA Discussion
Paper on Simplification and assessment of the credit
risk framework (EBA/DP/2026/01)**

May 2026

Answers to the questions for consultation

Simplification in the Standardised Approach

Q1. For the purpose of reporting under CRR Article 430a, which definition of loss should be used?

In principle, both definitions of losses (i.e., “estimated losses” and “observed losses”) are deemed admissible for the purposes of CRR Article 430a, as they lead to similar conclusions regarding the stability of national immovable property markets.

For the purpose of harmonisation, as regards Member States consistency with the loss rates provided by each jurisdiction should be preferably ensured via application of the approach based on estimated losses, as it is familiar to IRB institutions participating in the EBA Benchmarking exercise, allowing for continuity in existing reporting processes, and can be calculated using input data available to all institutions, whether under the IRB or Standardised Approach (SA).

For equivalent third countries, both definitions, estimated losses and observed losses, should be deemed valid subject to the rates published by the relevant competent authorities. Observed losses refer to losses actually recognised during the reporting period across all exposures outstanding in that period (defaulted and non-defaulted), irrespective of when the default event occurred. To face with the time-lag issue, larger period could be observed. These losses are anchored in accounting recognition events rather than model-based estimations, which may vary materially across institutions and jurisdictions.

In practice, equivalent third countries typically report loss rates based on observed losses. In addition to differences in the definition of losses, other differences in the calculation methodology may arise between loss rates published by different competent authorities (e.g., overall losses vs losses up to the secured value of the property). Although these methodologies may differ slightly from EU-defined loss rate measures, they often produce valid conclusions.

Loss rates calculated based on observed losses and published by competent authorities should be recognised or serve as inputs for establishing a mapping to “corresponding loss rates”, provided that: these metrics are accompanied by transparent metadata of portfolio perimeter, loss recognition rules, denominator definition, and statistical data is available and representative.

Anyway, any simplification or redefinition of Article 430a loss metric should consider not only its use for the purpose of the preferential IPRE treatment, but also other CRR provisions that reference the same metric. This is important because:

- the CRR already contains a definition of “loss” for specific purposes;
- Article 430a reporting is referenced in multiple parts of the real estate framework;
- inconsistent adjustments could create interpretational uncertainty or unintended inconsistencies.

Therefore, simplification efforts should ensure that changes to Article 430a remain coherent with the broader structure of the CRR and do not inadvertently affect related rules with unintended impacts.

Q2. Should the loss data (CRR Article 430a) be used for the assessment of RWs of real estate exposures under CRR Article 126(4) and CRR Article 465(11)?

While the intention to introduce stronger empirical grounding in the calibration of real estate risk weights is understandable, relying mechanically on observed loss data may introduce significant procyclical effects. Losses in real estate markets often exhibit strong temporal clustering, and a framework directly linked to short term loss observations could reinforce credit tightening during downturns. An approach that makes use of supervisory loss datasets in a manner that smooths volatility and avoids mechanically amplifying market cycles would be preferable.

Q3. Which elements of the real estate framework should be further simplified?

A key point where simplification is deemed necessary is the new “property value” requirement in Article 229(1)(b)(ii) CRR, which obliges institutions to adjust the value of buildings “to take into account the potential for the current market value to be significantly above the value that would be sustainable over the life of the loan.” While acknowledging that this concept stems from the Basel framework and was intended to ensure prudence in collateral valuation, it has proven extremely difficult to operationalize in a consistent and harmonized manner across Member States, given the absence of a clear methodology for determining long term “sustainable” values and the lack of alignment with national valuation standards.

The practical challenges associated with its application and the actual limitations in ensuring consistent outcomes across institutions have also been acknowledged outside the EU (and jurisdictions like the UK decided not to implement the “property value” adjustment).

In light of the above, it would be very important that the EBA supported a review of Article 229(1)(b)(ii) CRR aimed at a more harmonized and operationally workable approach, namely reverting to the long-standing, well-established and transparent concepts of “market value” and “mortgage lending value”.

Q4. Which other clarifications do you consider necessary to apply the new ECAI framework?

With specific regard to the EBA proposal, the pragmatic approach to facilitate the transition toward the new ECAI framework is appreciated and shared.

Anyway, the number of rating agencies producing an explicit rating without government support is limited and it will take time before a sufficient coverage of rating following the new rules is available. It is therefore essential that in the meantime the Competent Authorities make full use of the option available under Article 495e up to December 2029.

More generally speaking, another issue should be highlighted that deserves attention in the context of a simplification of the credit risk framework with regard to the ECAI ratings. The obligation for the bank to conduct “due diligence” on external ratings seems unduly complex considering that prudential regulations only allow the use of ratings assigned by

recognized ECAs operating within a specific regulatory framework. The benefit of a further assessment by institutions, therefore, appears to be very limited. Under the premise that we do not consider it necessary, if the legislator deems necessary an additional assessment of reliability of external ratings, such analysis should be performed by Authorities/Supervisors, as this would be much more efficient than requiring the same task to each single bank.

Simplification in the IRB Approach

Q5. Should the consolidation of regulatory products for credit risk be a priority or should the regulatory stability be preferable instead? Have you identified any redundancies in IRB products?

While acknowledging the merits of a rationalisation of the regulatory landscape by means of a consolidation of regulatory products in the field of internal models, it has also to be noted the huge effort of banks, given the extensive and resource intensive CRR3 implementation already underway and the long supervisory approval cycles associated with IRB model changes. In this regard, we would recommend that any rationalisation exercise should avoid introducing new requirements or additional interpretative layers. The objective should be to improve readability and consistency of the framework, not to expand supervisory expectations or conservatism through consolidation, neither to introduce regulatory changes that would trigger changes in modelling and consequent approval processes at a time where banks are still engaged in model changes following to CRR3.

Q6. Do you consider that the integration of environmental and social risks into the credit risk framework could be further enhanced without undermining its simplicity? Which areas, if any, would you prioritise for further work or clarification?

Given current data limitations, the forward looking nature of ESG risks, and the need to avoid double counting between Pillar 1, Pillar 2 and stress testing, we believe that for the time being E&S risks should continue to be primarily addressed through Pillar 2 processes, while incorporation in Pillar 1 should only take place upon condition of statistical and historical robustness of data and models.

It is therefore recommended that EBA clarifies that the ESG drivers should be included in Pillar 1 models only as long as they would make it possible to achieve a degree of statistical robustness that is suitable to the very strict requirements for these models.

The integration of environmental and social risks into the IRB framework is underway but structurally incomplete.

The main challenges are:

1. Data scarcity: before any methodological integration, the availability and comparability of ESG data must be addressed. Currently, each institution selects its own ESG proxies based on availability and internal model compatibility, not on

harmonised criteria. This means that even for identical portfolios, different ESG inputs will produce non-comparable estimates. EBA, or alternatively other Competent Authorities (e.g. Banca d'Italia) in coordination with the ECB, should develop and maintain a standardised ESG data repository, in particular for environmental data, establishing a common minimum reference base for IRB model calibration.

2. Representativeness historical data: Credit origination processes have historically not incorporated ESG variables into their decisions — neither at rating nor pricing level. Consequently, available historical series reflect pre-ESG behaviours and policies: attempts to identify stable patterns between ESG variables and default or loss rates on long historical datasets are often methodologically weak, because the empirical relationship being estimated did not exist in the behaviour of market participants when the data were produced. Forcing direct integration of ESG variables into the core structure of IRB models (as PD or long-run and LGD calibration factors) currently risks introducing more bias than informational value, potentially worsening the quality of estimates rather than improving it. This will progressively diminish as ESG historical series lengthen and sustainability considerations become embedded in origination policies, but it remains an operational constraint that cannot be ignored today.

Q7. Which requirements should apply in relation to the measurement of the performance of continuous models (e.g. back-testing)? How could testing requirements be facilitated and enhanced for continuous models that are compliant with CRR, Part three, Title II, Chapter 3, Section 6 (Requirements for the IRB approach)?

[...]

Q8. Which requirements should apply in the application phase of continuous models (e.g. overrides)?

[...]

Q9. Which challenges have you encountered in implementing the new CRR definition of facility?

It is considered important that institutions retain a degree of flexibility in the application of facility definitions.

The primary challenge in implementing the CRR3 facility definition is not operational but conceptual: a facility is not a static entity but an object whose economic nature changes over time, and the three IRB parameters — PD, LGD and CCF — each capture ontologically distinct moments of the credit lifecycle, with their own facility logic that cannot be reduced to that of the others, for example:

- CCF measures the obligor's behaviour in the 12 months preceding default based on the contractual structure at the reference date. The CCF facility is the one existing at the reference date — an object anchored in the client's contractual present;

- LGD measures the recovery process after default, traversing product transformations, restructurings and reclassifications. The LGD facility must be reconstructed over time and its optimal perimeter depends on the actual recovery process, not on the original contractual structure. The parameters therefore operate on different representations of the same obligor — and no single facility definition can reconcile them without distorting at least one.

In addition, there are also clear tensions between the current GLs:

- EBA GL/2017/16 paragraph 136: which requires LGD and cash flows to be attributed to the original facility in restructuring cases;
- EBA CP GL CCF: the CCF rule on product transformation, which requires the same attribution for CCF.

The two rules formally converge but diverge in implementation, because management systems register the product as classified at the time, not as reconstructed retrospectively. In component models, the same obligor with a revolving product transformed into instalment contributes to the LGD pre-litigation of instalment products while to the revolving facility for EAD: there is no way to satisfy both conditions simultaneously without an explicit assumption that introduces bias. This regulatory circularity has not been resolved in either the DP or the CCF CP.

In addition, the facility definition issue connects directly with the consistency requirement between modelling approaches established in EBA GL/2017/16, which governs the component LGD model (direct estimation vs. component approach). It requires that, where institutions estimate intermediate components separately and combine them to obtain the final LGD estimate at facility level, empirical evidence be provided of the absence of bias introduced by the combination. This requirement is recalled and elaborated in EGIM, which specifies the ECB's supervisory expectations for component-based LGD models. The critical point is that both requirements presuppose the comparability of the underlying data between the component model and the wholelife approach, in terms of the facility used as the unit of observation. Where this comparability is structurally absent — as occurs in the revolving-to-instalment product transformation case described above, where banking systems register the exposure according to the classification applicable at the time rather than the original facility — the divergence between the two approaches does not represent a measurable methodological difference, but an artefact of the stratification of facility definitions between operational systems and regulatory ex-post reconstructions. This generates a direct and problematic consequence for MoC treatment: the systematic difference between the results of the component model and those of the wholelife approach, in this case, cannot be classified as a data deficiency (category A) or a policy change (category B), and is therefore not quantifiable through the MoC framework established in EBA GL 2017/16. Applying a MoC to such a difference does not reduce it — rather, it introduces a further layer of conservatism on a quantity that expresses not estimation risk but structural incomparability between approaches operating on different representations of the same obligor. The regulatory framework should therefore explicitly allow for greater flexibility; this does not mean abandoning the consistency requirement— which remains fully valid—but rather acknowledging that its mechanical application, in the

presence of structural discontinuities in the data, produces implicit and unquantifiable distortions whose impact on capital adequacy could be greater than what was originally intended to be corrected.

Q10. Should a consistent and single facility definition be applied across all risk parameters?

A single, uniform facility definition across all IRB risk parameters is not considered desirable.

In fact, on one hand, a strict application of the same aggregation level for PD, LGD and CCF would imply a huge operational effort, as existing IRB model architectures were not designed around a uniform facility definition. On the other hand, different level of aggregation may be justified by the different economic meaning of the parameters and empirical foundations of the parameter estimates.

For these reasons, it would be important to retain flexibility in selecting the appropriate aggregation level for each risk parameter, as long as choices are well justified, transparent, and aligned with the institution's underlying processes and data.

Q11. Are adjustments proposed in the representativeness requirement for the CCF parameter also suited for PD and LGD risk parameters? Which amendments would be needed to accommodate PD and LGD specificities?

See Q12.

Q12: Do you consider further simplification of the representativeness requirement, as proposed for the CCF parameter, as necessary for PD and LGD and if so, what kind of simplification?

The adjustments proposed to the representativeness requirements in the context of CCF estimation are considered broadly suitable for extension to the PD and LGD parameters, as they address methodological challenges common to all IRB risk parameters. However, their application to PD and LGD requires specific amendments, as outlined below.

A key concern in applying representativeness requirements to LGD and PD models relates to the structurally asymmetric nature of the samples involved in the comparison. The representativeness analyses prescribed under the Guidelines require a distributional comparison between the model application portfolio - encompassing all exposures, including performing ones - and the development and calibration sample, which by construction consists exclusively of defaulted exposures. This structural asymmetry implies that systematic distributional differences between the two samples may arise not as a genuine signal of a representativeness deficiency, but as a physiological consequence of the default selection process itself. Obligors entering default systematically exhibit characteristics that differ from the performing population across multiple dimensions - exposure size, relationship seniority, product type, economic sector - regardless of the quality of the development sample. Interpreting such distributional differences as indicators of non-representativeness would require the application of additional Margins of Conservatism that do not reflect genuine model uncertainty, but rather a methodological bias inherent in the comparison framework. It is therefore considered necessary to introduce an explicit distinction between structural non-representativeness - attributable to the nature of the data generating process - and model non-representativeness -

attributable to genuine deficiencies of the development sample relative to the target portfolio. Only the latter should mandatorily trigger MoC requirements, while the former should be subject to documentation and justification without automatic capital consequences.

It is recommended to explicitly extend to PD and LGD the differentiation between data used for model development and data used for performance testing. Greater flexibility in representativeness requirements should be permitted at the development stage, acknowledging that the availability of sufficiently granular historical data may require the use of samples not perfectly aligned with the current portfolio. At the testing stage, however, representativeness remains essential, as its absence may undermine the reliability of back-testing outcomes.

It is further noted that current requirements focus predominantly on cross-sectional representativeness, distributional comparisons on static variables, without explicitly addressing temporal representativeness. In the presence of structural changes in the portfolio over time, such as modifications to the definition of default, changes in underwriting policies, or large-scale disposal transactions of non-performing exposures, the historical sample may be statistically representative of the past portfolio while failing to adequately reflect the characteristics of the current one.

Q13. Should these simplifications be pursued? Do you have any preferred approaches with respect to these simplifications?

The proposed simplifications are mainly focused on the introduction of fall-back approaches as possible substitutes to certain elements of internal models. In this regard, a few principles are key:

- the decision to adopt a fall-back approach should remain a methodological choice of the institution, based on clear internal criteria and embedded in model governance. The burden of proof should not fall on the institution, including control functions, to demonstrate the appropriateness of the fall-back through extensive analysis against the traditional methodology, as this would undermine the intended simplification objective. Likewise, fall-back approaches should not become a benchmark against which traditional methodologies are expected to be justified or reassessed;
- although calibrated in a prudential way, fall-backs should not result in outcomes that are not consistent with the level of risk (e.g. a 100% fallback on CCF plus MOC implies a higher capital charge for a facility that is available to the client but not yet withdrawn than for a drawn exposure).

Q14. Do you have comments and suggestions with reference to the calibration of the fall-back approaches?

As also said in the response to Q13, it is essential that the calibration of the fall-back approaches does not result in outcomes that are not consistent with the level of risk. More generally speaking, given that calibration is a major driver in the choice on whether to apply the fall-back options, overly conservative calibration would make such options ineffective.

Q15. Do you see other potential simplification areas where the modelling burden is not commensurate to the gain in risk sensitivity?

A very important issue that should be addressed in the review of the IRB framework is the introduction in CRR3 Article 182(1) of a zero floor for observed CCFs used in risk quantification.

This is not a mere issue of modelling burden but rather of appropriateness of the regulatory approach. The CRR3 rule is not in line with the Basel text, only requiring that “For on-balance sheet items, banks must estimate EAD at no less than the current drawn amount [...]”. The changes introduced in CRR3 shift this from an output floor (as intended by the Basel framework) to an input floor. This requirement:

- introduces an unnecessary upward bias and implies a significant impact in terms of increase in the estimated EAD, resulting in increased conservatism and reduced risk sensitivity
- limits the ability of the models in reflecting improvements in proactive management practices and processes and
- increases the divergence between CCF estimates used for regulatory capital computations and those used for IFRS9 accounting purposes. The latter is namely required to be as close as possible to a best estimate of the parameter, which for the many reasons reported above is not allowed by the presence of a floor on negative CCFs. More generally speaking, decoupling of the regulatory parameter from actual estimate undermines the use test.

Under a more general perspective, it would be desirable that excessively detailed modelling specification (like the abovementioned floor) were not part of Level 1 rules, since such level of detail appears disproportionate.

On a different page, the flexibility proposed under C15 — allowing institutions to incorporate elements of the cohort approach within the 12-month fixed horizon framework — is welcomed. The 12-month fixed horizon introduces a structural rigidity that may distort CCF estimates in both directions, depending on the product type and the borrower’s behavioural dynamics prior to default. Permitting a hybrid framework, subject to adequate documentation and explanation of material deviations from long-run averages, represents a proportionate and pragmatic step toward enhanced risk sensitivity without abandoning the standardisation benefits of a fixed reference point.

A further and more structural source of distortion, specific to the 12-month fixed horizon approach, relates to the historical reconstruction of the New Definition of Default (New DoD). Under the 12-month fixed horizon, the reference exposure must be sourced exactly 12 months prior to the default date. However, a significant subset of defaults recognised under the New DoD would not have been classified as defaults under the previously applicable criteria. This creates a systematic identification problem: for these observations, the exposure recorded at T-12 reflects the behaviour of a borrower who, under the regulatory and operational framework in place at that historical reference point, was not yet — and would not have been — classified as showing signs of material credit deterioration. As a result, utilisation levels at T-12 tend to be very low or close to zero, not because the borrower was managing credit risk prudently, but simply because no financial stress was observable or actionable under the then-applicable default definition.

When these observations are included in the CCF estimation sample, they generate systematic outliers.

This distortion is fundamentally different in nature from standard data quality issues: it is not addressable through the Margin of Conservatism framework, as it does not reflect a modelling deficiency or a data gap in the conventional sense, but rather an irreconcilable inconsistency between the default classification criteria applicable at the time of the historical observation and those used to define the target variable in the current estimation framework. The cohort approach, by relying on periodically observed exposure snapshots aligned with a consistent default definition, is structurally more robust to this issue and provides an additional operational and conceptual rationale for the flexibility proposed under C15.

A pragmatic solution could consist in a combined framework, where CCFs are estimated using both approaches and subsequently reconciled through selection or aggregation criteria aimed at preserving risk sensitivity.

Other areas for possible interventions are:

- **Data representativeness requirements for long-run averages:** the requirement to use long historical observation periods — often spanning multiple economic cycles — can be disproportionately burdensome for institutions that have undergone significant changes in business model, portfolio composition, or risk management practices. A more flexible approach, allowing institutions to use shorter but more representative data series subject to appropriate adjustment and validation, would reduce data management costs while preserving the conceptual integrity of the long-run average;
- **M&A:** a more streamlined approach would equally be warranted in how historical data is handled in the context of mergers and acquisitions. Obligating institutions to incorporate legacy data from acquired entities — frequently built on divergent systems, varying default definitions and differing credit policies — creates undue complexity while contributing little in terms of predictive accuracy. Permitting institutions to instead draw on forward-looking, post-integration data would more accurately reflect the true risk profile of the newly combined entity. Similarly, we advocate for a more balanced application of homogeneity and heterogeneity requirements within segmentation frameworks. An overreliance on extensive statistical testing can produce grade structures that are either unstable or lack economic coherence. A more streamlined framework, focused on a targeted set of key indicators and supported by expert opinion and qualitative analysis, would uphold prudential standards while avoiding unnecessary methodological burden. This consideration is especially pertinent for low-default portfolios, where directly applying testing methodologies originally designed for high-default contexts risks driving excessive consolidation, impairing the model's ability to discriminate effectively, and ultimately undermining its practical value for internal risk management, pricing and business steering.

In conclusion, the principle underlying C15 — that regulatory requirements should be proportionate to their contribution to risk sensitivity — should be applied consistently across the A-IRB framework. The areas identified above represent concrete opportunities to reduce modelling burden without undermining the risk-based nature of internal models

and would contribute to maintaining the competitiveness and attractiveness of the A-IRB approach for European institutions.

Q16. What do you perceive as challenges in your capacity to collect appropriate data, in particular in relation to indirect costs?

[...]

EBA reports to assess L1

Q17. Do you agree with the approach proposed by EBA? Do you see further measures as necessary?

Generally speaking, the proposed criteria are considered reasonable and appropriate. Anyway, a few issues should be highlighted. The first one is the need to specify that the industry should be engaged in the assessment of the calibration as well as of costs and benefits of the proposed simplification measures. A transparent debate on the outcome of the assessment should be envisaged before the final position of the Authority is defined.

Another important aspect, that should be considered in the assessment of the overall framework and that seems not addressed in the Discussion Paper, is the adherence of L2/L3 products and other implementing measures to the Level 1 mandates (and more generally, to Level 1 provisions).

While the Level 1 should not embed highly technical details (e.g. as regards modelling specifications), the simplification process should aim at ensuring that the numerous layers of regulation do not add complexity and burdens that exceed what is necessary to comply with the Level 1.

With specific regard to internal models, the framework could be further enhanced by explicitly incorporating additional dimensions aimed at capturing the behavioural and dynamic nature of credit risk.

In particular, we would suggest considering the following complementary aspects:

- **Representativeness of risk dynamics:** Beyond measuring the gap between regulatory parameters and observed risk, it is important to assess whether the framework adequately captures the underlying dynamics of exposures, especially in the pre-default phase. Empirical evidence (e.g. on CCF behaviour) shows that exposure trajectories may reflect structural mechanisms such as proactive credit management or product-specific features, rather than pure risk signals. A framework that does not account for these dynamics risks misinterpreting economically meaningful behaviours as estimation errors, leading to corrections that reduce rather than enhance risk sensitivity
- **Avoidance of systematic biases:** Certain modelling assumptions or regulatory prescriptions may introduce structural biases — rather than random estimation errors — potentially leading to persistent over- or under-estimation of risk. For instance, the treatment of negative CCFs or the reliance on fixed observation horizons may distort the true risk profile if not properly calibrated (see also Q15)
- **Proportionality and heterogeneity across institutions and products**
A uniform application of simplified rules may have uneven effects across institutions, portfolios, and product types. Ensuring sufficient flexibility to account

for idiosyncratic characteristics is therefore key to avoiding unintended penalisation and preserving risk sensitivity. This is especially relevant for self-liquidating products and revolving facilities, where structural exposure dynamics differ fundamentally from those of term loans and where a single fixed-horizon reference point may be systematically unrepresentative.

These additional dimensions are particularly relevant in areas such as the estimation of CCF under the 12-month fixed horizon approach, where empirical evidence highlights potential mismatches between regulatory assumptions and observed behaviours — including decreasing exposures due to active credit management, increasing utilisation close to default, and historically reconstructed reference exposures affected by the New DoD transition. Taken together, these sources of distortion are not independent: they may compound one another within the same estimation sample, amplifying the divergence between regulatory CCF estimates and the true underlying risk profile.

In this context, complementing the existing framework with criteria that explicitly address behavioural representativeness, structural biases stemming from regulatory transitions, and product-level heterogeneity would strengthen the overall assessment and better support the objective of achieving an appropriate balance between simplicity and risk sensitivity.