Just in Time

Quantification of Margin of Conservatism Category C - Correlations and Quantification Levels

October 2026



iason ESSENTIAL SERVICES FOR FINANCIAL INSTITUTIONS

Executive Summary

- This paper reviews the quantification of **Margin of Conservatism Category C (MoC C)** for probability of default (PD) estimates under the IRB regulatory framework. MoC C addresses general estimation error and is therefore distinct from margins linked to specific data or methodological deficiencies.
- The paper focuses on three practical questions: how to quantify MoC C when long-run average default rates (LRADRs) are based on overlapping one-year observations; how to reconcile MoC C calculated at calibration-segment level with the ECB expectation of grade-level quantification; and how to address downward bias in asset-correlation estimates derived from historical default-rate time series.

iason

Copyright © 2026 - All rights reserved

At a Glance

01	Regulatory Context and Core Problem	4
02	MoC C for Overlapping One-year Default Rates	6
03	Harmonising Calibration-segment and Grade-level MoC C	8
04	Downward Bias in Asset-correlation Estimates	10
05	Conclusions and Takeaways	13

Keywords: Credit risk, IRB

Copyright © 2026 - All rights reserved

01

Regulatory Context and Core Problem

Copyright © 2026 - All rights reserved

Regulatory Context and Core Problem

The regulatory objective is to **avoid underestimation of credit risk**.

Institutions first determine best-estimate PDs and then add appropriate margins of conservatism.

Under the EBA Guidelines and the ECB Guide to Internal Models, **MoC C captures general estimation uncertainty**.

The paper's contribution is to show that this uncertainty can be materially affected by the structure of default-rate observations, the aggregation level at which MoC is computed, and the estimation of default dependence through asset correlations.

This distinction is important because MoC C is not meant to compensate for known and identifiable model flaws, which should instead be addressed through Categories A or B or through remediation of the model itself. It captures the **residual uncertainty** that remains even after the institution has used appropriate data, estimation techniques and calibration processes.

*For that reason, the paper treats MoC C as a quantitative add-on that should be proportional to the uncertainty of the relevant estimator and should decrease when the **evidence base becomes stronger**.*

The practical challenge is that a statistically sound MoC C must be both defensible and operationally usable.

- If the margin is too mechanical, it may fail to capture the actual uncertainty of the estimate.
- If it is too judgemental, it may be difficult to evidence, replicate or challenge.

The paper therefore seeks a middle ground: formal methods grounded in estimator uncertainty, but expressed in a way that can be connected to supervisory expectations, model validation and capital impact analysis.

Copyright © 2026 - All rights reserved

02

MoC C for Overlapping One-year Default Rates

Copyright © 2026 - All rights reserved

MoC C for Overlapping One-year Default Rates

Just in Time

The first methodological contribution concerns LRADRs calculated from **overlapping annual windows**.

Annual default rates observed monthly share most of their underlying performance period: two consecutive one-year rates have 11 months in common. Treating them as independent would therefore ignore a relevant source of estimation uncertainty. The paper addresses this by expressing the LRADR as a weighted sum of monthly default counts and deriving its variance while incorporating **serial covariance** between those counts.

This result matters because **positive dependence between monthly defaults** increases the variance of the LRADR and therefore the appropriate MoC C. The paper also notes a definitional difference between its mathematical construction and paragraph 73(b) of EBA/GL/2017/16: the derivation may capture defaults of obligors entering the portfolio after the beginning of the one-year period. Provided these additional defaults are not negatively correlated with regulatory defaults, the resulting variance estimate is **conservative**.

Under stationarity, mixing and finite long-run variance assumptions, a Central Limit Theorem remains applicable despite overlapping observations.
This allows confidence-based MoC C quantification without assuming independence. The approach is also consistent with the supervisory intuition that shorter historical series should imply higher MoC C, because estimation uncertainty declines as the observation period grows.

A useful feature of the treatment is that it preserves the intuition of confidence-interval-based conservatism while adapting it to the actual information structure of the data. If overlapping default rates were handled as if they were independent, the effective sample size would be overstated and the resulting uncertainty would be understated. By contrast, using the covariance structure between monthly counts links the MoC directly to the persistence of default experience across time.

From a **validation perspective**, the key point is that more frequent observations do not automatically create more independent information. Monthly calculation of annual default rates increases the number of data points, but the economic information embedded in those points remains highly overlapping. The paper's framework makes this explicit and prevents the MoC from being understated simply because the data set appears longer after monthly rolling windows are introduced.

Copyright © 2026 - All rights reserved

03

Harmonising Calibration-segment and Grade-level MoC C

iason

Copyright © 2026 - All rights reserved

Harmonising Calibration-segment and Grade-level MoC C

Just in Time

The second contribution addresses the gap between **institutional calibration practice** and supervisory expectations.

Institutions often determine the LRADR at calibration-segment level, while the ECB Guide to Internal Models expects MoC C to be quantified at **grade level**. Since the regulatory RWEA function is nonlinear in PD, grade-level and segment-level MoCs cannot be assumed to have the same capital effect merely because they use the same confidence level.

The paper compares two ways of constructing conservative grade-level PDs. In the first, a calibration-segment MoC is allocated proportionally to grades. In the second, grade-specific MoCs are calculated directly. To make the two approaches approximately equivalent in aggregate capital terms, the paper derives a scaling parameter for grade-level MoCs. The derivation relies on a first-order Taylor expansion of the RWEA function and produces grade-specific sensitivity weights reflecting PDs, exposures, LGDs, regulatory correlations and each grade's relationship with the LRADR.

The resulting scaling factor reflects the relationship between the variance of the calibration-segment default rate and a non-normalized weighted combination of grade-level default-rate variances. Where the analytical approximation is not sufficiently accurate, it can be used as an initial value for numerically solving the exact RWEA-equivalence equation. The practical message is clear: harmonisation requires attention to both **statistical uncertainty** and **capital sensitivity.**

This section is particularly relevant for implementation because it avoids treating the ECB grade-level expectation as a purely mechanical reporting requirement. The grade at which conservatism is applied affects regulatory capital through the shape of the risk-weight function. A MoC concentrated in high-sensitivity grades can have a different aggregate RWEA effect from the same total MoC distributed across low-sensitivity grades. The methodology therefore asks not only whether the PDs are conservative, but whether the conservatism has the intended capital effect.

This also clarifies why grade-level MoC C should not be viewed as a simple decomposition of the calibration-segment margin. A grade with a relatively small exposure share may still have a disproportionate RWEA impact if its PD lies in a region where the regulatory formula is highly sensitive. Conversely, a large grade may have limited incremental capital sensitivity. The proposed scaling approach is designed to preserve the intended aggregate prudential effect while respecting the supervisory preference for grade-level quantification.

Copyright © 2026 - All rights reserved

04

Downward Bias in Asset-correlation Estimates

Copyright © 2026 - All rights reserved

Research Questions and Methodology 1/2

The third contribution concerns **default dependence**.

Asset correlations are often estimated from historical default-rate time series, especially where market data are unavailable or inappropriate.
The paper examines the maximum-likelihood estimator of Düllmann and Trapp, under which asset correlation is an increasing function of the variance of transformed default rates. In practice, however, the estimator may be **biased downward**.

Two drivers are particularly relevant.
First, **default-rate time series** are often short, so asymptotic properties may not materialize.
Second, the **theoretical estimator** assumes a homogeneous-PD portfolio, while real portfolios typically contain obligors with heterogeneous PDs.
Combining grades can increase sample size but violates that homogeneity assumption and can generate structural bias.

The simulation framework makes this bias observable because the true correlation is known by construction. The paper varies asset correlation, unconditional PD and discriminatory power, generates obligor-specific PDs through credit-score distributions and logistic regression, and simulates defaults in a Vasicek-type model. The results show that estimated correlations generally fall below true correlations, with underestimation increasing as **discriminatory power** and **true correlation** rise.
Conventional 95% and 99% upper confidence bounds do not necessarily eliminate the bias, because they address sampling uncertainty around an estimator rather than the structural mismatch caused by **PD heterogeneity**.

Copyright © 2026 - All rights reserved

Research Questions and Methodology 2/2

Just in Time

The paper therefore proposes **simulation-based correction factors**, calculated as the ratio between the theoretical asset correlation and the average correlation recovered by the estimator in the simulation. These factors are **not universal constants**: they depend on portfolio and parameter characteristics such as discriminatory power, unconditional PD, true or estimated correlation, portfolio size and historical observation length.

The supervisory implication is that a **bank should be able to justify not only the point estimate of asset correlation, but also any additional conservatism applied to compensate for estimator bias**. Where the estimator is known to be biased under realistic portfolio heterogeneity, simply replacing the variance by an upper confidence bound may give a false sense of prudence. A simulation-based benchmark can help distinguish sampling uncertainty from structural bias and can support targeted, evidence-based adjustments.

A further merit of the simulation design is that it separates several effects that are often intertwined in empirical data. In a real portfolio, the observed default history reflects the joint influence of macroeconomic conditions, rating-system quality, portfolio composition and sample length. By controlling the true correlation and the discriminatory power of the rating system, the paper can isolate how PD heterogeneity affects the estimator. This gives **supervisors and validators a clearer basis for deciding whether an empirical estimate should be accepted, challenged or adjusted**.

Copyright © 2026 - All rights reserved

05

Conclusions and Takeaways

iason

Copyright © 2026 - All rights reserved

Conclusions and Takeaways 1/2

Just in Time

The paper provides a coherent treatment of three sources of complexity in MoC C quantification: **temporal dependence** in overlapping default-rate observations, reconciliation between calibration-segment and grade-level calculations, and **bias in asset-correlation estimation**. Its main supervisory implication is that apparently conservative methods may still be insufficient where the underlying estimator or aggregation method embeds **structural weaknesses**.

For practical model validation, the paper suggests several challenge points. Validators should check whether overlapping observations have been treated consistently with their dependence structure, whether the grade-level allocation of MoC produces a capital impact aligned with the calibration objective, and whether asset-correlation estimates have been tested for bias under realistic PD heterogeneity. They should also assess whether any correction factors are calibrated to the institution's own portfolios rather than imported from **generic benchmarks**.

The remaining limitations are mainly implementation-related. The derivations rely on assumptions that may be only approximately satisfied in real portfolios, including stationarity, sufficiently weak dependence over time and stable portfolio composition. The RWEA-equivalence approach is also approximate unless the nonlinear equation is solved directly. Finally, the simulation-based correction of asset correlation is necessarily portfolio-specific.

Copyright © 2026 - All rights reserved

Conclusions and Takeaways 2/2

Just in Time

These limitations do not undermine the contribution of the paper, but they mean that the proposed methods should be applied with documented judgement rather than as automatic formulas.

1. Overlapping default-rate observations require explicit treatment of **serial dependence**. Positive covariance between monthly defaults increases **LRADR uncertainty** and therefore **MoC C**. For banks using rolling annual default rates, the key risk is treating monthly observations as **independent evidence**; this may lead to an understated margin, insufficiently conservative PDs and supervisory challenge on the **effective sample size**.

2. Grade-level MoC C cannot be reconciled with calibration-segment MoC C through nominal confidence levels alone. What matters is the combined effect of **capital sensitivity**, **default-rate variance** and **grade allocation**. Banks should therefore demonstrate that grade-level conservatism produces an aggregate **RWEA impact** consistent with the calibration objective, rather than simply redistributing a segment-level add-on across grades.

3. Asset-correlation estimates from default-rate series may be **structurally downward biased**. **PD heterogeneity** and short histories can make conventional confidence intervals insufficient. Banks relying on internally estimated correlations may therefore need additional **validation**, **bias testing** and portfolio-specific overlays, especially where heterogeneous rating grades are pooled to increase the available sample.

4. Simulation-based adjustments are useful, but should be calibrated to **portfolio-specific dimensions** rather than applied as generic uplifts. Any correction factor should reflect the bank's own **PD distribution**, **portfolio size**, observation period and discriminatory power. Generic uplifts may be difficult to defend in **model validation** or **supervisory review**.

5. The main operational impact is stronger **evidence**, **documentation** and **governance** around MoC C. Banks should be prepared to explain the statistical basis of the margin, the level at which it is quantified, the resulting **capital impact** and the validation evidence supporting any correlation adjustment. This may increase modelling effort, but it also improves the **defensibility of IRB estimates** and reduces the risk of ad hoc supervisory add-ons.

Copyright © 2026 - All rights reserved

ESSENTIAL SERVICES FOR FINANCIAL INSTITUTIONS

iason is an international consulting firm that has been supporting both financial institutions and regulators in topics related to Risk Management, Finance and ICT since 2008

Strategy

Strategic advisory on the **design** of **advanced frameworks** and **solutions** to fulfil both **business** and **regulatory needs** in Risk Management and IT departments

Methodology & Governance

Implementation of the designed **solutions** in bank departments **Methodological support** to both **systemically important financial institutions** and **supervisory entities**

Solution

Advanced **software solutions** for **modelling, forecasting, calculating** metrics and **integrating** risks, all on cloud and distributed in Software-as-a-Service (**SaaS**)

KEEP IN TOUCH

iason

Company Profile

iason is an international firm that consults Financial Institutions on Risk Management. **iason** integrates deep industry knowledge with specialised expertise in Market, Liquidity, Funding, Credit and Counterparty Risk, in Organisational Set-Up and in Strategic Planning.

Dario Esposito

Lorena Corna

Bianca Ghilardi

This is an **iason creation**.
The ideas and the model frameworks described in this presentation are the fruit of the intellectual efforts and of the skills of the people working in iason. You may not reproduce or transmit any part of this document in any form or by any means, electronic or mechanical, including photocopying and recording, for any purpose without the express written permission of iason.

www.iasonltd.com

Copyright © 2026 - All rights reserved