Model Risk Diversification in Bank-wide Risk-Weighted Assets

Themis Rallis*†

August 19, 2025

Abstract

Basel III imposes strict guidelines on large banks that develop statistical models to estimate their risk-weighted assets (RWAs). Banks are required to quantify additional buffers, called a margin of conservatism (MoC), to protect against model risk. Setting the appropriate confidence level for the MoC is, however, non-trivial, and is frequently the topic of debate between banks and regulators. In this paper, we introduce the quantile scaling factor (QSF), a novel tool for regulatory model risk management. The QSF is derived under an asymptotic single risk factor (ASRF) model applied to estimated credit portfolio RWAs. An empirical study on a realistic, synthetic bank loan book applies the QSF framework and the results are discussed. We find that errors in RWA estimates driven by idiosyncratic model risk factors and loan books with more uniform portfolio RWA concentrations benefit from high model risk diversification. This lowers the confidence level required per portfolio's MoC to attain a bank-wide confidence level set by the bank's own model risk appetite. Model risk diversification strengthens further the more portfolios a bank has in its loan book. The presented QSF framework provides risk managers with a simple tool to ensure their bank's model risk appetite is appropriately met. However, given that no two loan books are the same, it complicates regulatory benchmarking exercises and challenges current practices by supervisory authorities like the ECB.

Keywords

Risk-Weighted Assets; Model Risk Management; Statistical Coverage; Quantile Scaling Factor; Margin of Conservatism; Systemic Model Risk; Diversification; Credit Risk Management; Model Risk Appetite

*Credit Risk Model Developer, Quantitative Methods and Bankwide Models, ING | PhD Candidate, School of Business and Economics, VU Amsterdam

†Corresponding author. email: Themis.Rallis@ING.com | T.Rallis@VU.nl, phone: +31 6 21 28 49 96, address: Bijlmerdreef 106, 1102 CT Amsterdam, the Netherlands | de Boelelaan 1105, 1081 HV Amsterdam, the Netherlands.

1 Introduction

Banks that fall under ECB supervision are required to hold capital to cover for unexpected losses in their portfolios. For a given portfolio, this capital, derived from its risk-weighted assets (RWA), is a product of three primary risk parameters estimated via the means of internally developed statistical models. As the robustness and stability of the financial system is of paramount importance to the regulators, the Basel framework envisions that these estimates have applied to them a margin of conservatism (MoC) to ensure RWAs are sufficiently conservative and protected against model risk (**?**).

Defining what "sufficiently conservative" is, however, not trivial. While a 95% confidence level may be appropriate for a single portfolio, aggregating MoCs across multiple portfolios complicates the picture at the total bank level. Bank-wide model risk strategy should instead be anchored to a target confidence level for the total loan book, aligned with the bank's model risk appetite. This target, together with the composition and risk profile of the portfolios, determines the necessary MoC coverage per portfolio to ensure consistency and robustness at the aggregate level.

Within industry, the impact of arbitrary MoC coverage levels across portfolios is well understood. The larger the coverage, the more bolstered the bank's capital buffer, the more protected the bank is from model risk. However, if this coverage is unnecessarily wide, the bank hinders its own business practices by limiting its liquidity and supply to the credit market. With no official regulatory guidance on setting this interval, while subtle, the best choice of coverage level for a credit lending institution should receive more attention in both academia and industry.

The ECB defines model risk as "*the risk of error due to inadequacies in financial risk measurement and valuation models*" (**?**). While there exists adequate academic research on the topic of model risk, this has almost exclusively been focused market risk measures such as the Value-at-Risk (VaR) and Expected Shortfall (ES) (**??????????**). Some focus in the literature has directed itself to model risk for credit portfolios (**??????**), however these typically still use tail risk measures of portfolio losses to assess model inaccuracies.

Literature on MoC quantification and the impact of model risk on IRB models remains scarce and underdeveloped. **?** proposes a structured framework that decomposes model risk into two components, risk differentiation and calibration error, arguing for the integration of MoC directly into model outputs to avoid over-conservatism from uncoordinated parameter-level adjustments. **?** reinforces this perspective by addressing the challenge of data representativeness, particularly under crisis conditions, and advocates for MoC adjustments rather than data exclusions to ensure robustness in model calibration. The AIFIRM Position Paper No. 13 (**?**) defines MoC as the product of a calibrated factor and the standard deviation of the estimator. The paper explores multiple estimation techniques and calibrates the factor via simulation to ensure that model risk is addressed proportionately, without duplicating conservatism already embedded in regulatory capital formulas. **?** focus specifically on the asymptotic single risk factor (ASRF) framework and demonstrate that estimation error in PD leads to systematic underestimation of risk measures. Their proposed correction is based on adjusting the PD using the upper bound of a confidence interval and ensures the probability of observing exceptions aligns with the targeted confidence level. **?** proposes a probabilistic method for quantifying an MoC addressing missing data in calibration samples, and is the only research that discusses the arbitrary choice of confidence level for the MoC. While his approach satisfies key statistical properties (e.g., unbiasness, monotonicity) and is applicable at both score and factor levels, Biche's approach avoids the need to define the significance level of the MoC. Such an approach thus provides limited utility for a bank's general risk management practices.

While existing literature offers various methodologies for quantifying an MoC for risk parameters, none directly address the challenge of setting the confidence level within a rigorous framework that enables banks to align MoC estimation with their overarching model risk appetite across portfolios. This gap limits the ability to systematically balance regulatory conservatism with internal risk tolerance. In this paper, we develop a consistent mathematical framework that formalises interval selection by linking it directly to a bank's model risk appetite. Specifically, we derive a closed-form quantile scaling factor (QSF) between a target bank-wide coverage level and the implied coverage level required per portfolio to achieve it. Through simulation, we show that the QSF is driven by three key characteristics of a bank's loan book: (i) correlations in RWA estimation errors; (ii) RWA concentration across portfolios; and (iii) the total number of portfolios. The central claim of this paper is that model risk exhibits a degree of diversification across portfolios, meaning that individual model uncertainties tend to offset each other when aggregated at the bank-wide level. As a result, the confidence level required for the MoC at the portfolio level can be lower than the target confidence level set for the bank-wide RWA. This gives risk managers a simple tool to ensure RWA coverage levels best reflect the bank's own unique exposure to model risk. However, this challenges current practices by regulatory authorities when benchmarking the adequacy of coverage levels across banks under their supervision. Current model inspection exercises conducted by authorities such as the ECB tend to impose arbitrary and subjective expectations of capital conservatism levels, which, as this work shows, is too far detached from the model risk to which each bank is exposed.

The paper is structured as follows. Section **??** introduces the main concepts of model risk and how this may diversify across a bank's set of portfolios. Section **??** defines the variables with their statistical properties necessary for deriving the QSF in Section **??**. Section **??** utilises simulations to investigate the main characteristics that influence the QSF. Section **??** applies the QSF framework to a synthetic bank loan book with set model risk appetites. The loan book is generated using aggregate statistics published by supervisory authorities, and serves as a realistic example of how to estimate the QSF in practice. Section **??** discusses the results, provides further guidance for practitioners, and highlights topics for future research. Section **??** closes the paper with concluding remarks.

2 Model Risk Diversification in the Loan Book

Credit portfolios contain exposures of specific asset classes or geographies, each influenced by their own unique systemic credit risk factors, e.g., mortgage portfolios in the US versus Australia tied to each country's economic cycle, or corporate loans tied to industrial cycles. On the other hand, idiosyncratic credit risk refers to borrower-specific risk, such as firm-level distress or default, which is largely uncorrelated with broader economic conditions. It is essential to distinguish this credit risk from risks arising due to the models themselves. Model risk specifically refers to the possibility that RWA estimates are inaccurate due to flaws in model design, data, and/or calibration.

The traditional credit risk literature defines idiosyncratic and systemic credit risk in a way that links them through an asset correlation. This is known as the asymptotic single risk factor (ASRF) model, whereby all borrowers in a portfolio share this fixed correlation with the systemic factor. In this work, we will stick to the same terminology for model risk, and the framework developed will encapsulate an underlying ASRF model. Put bluntly, this implies that all main results are derived by keeping the asset correlation "submerged" within each RWA model error. Explicitly, model risk can be decomposed into:

- Idiosyncratic model risk, such as missing risk drivers or regressors in development datasets, or poor data quality affecting specific portfolios.

- Systemic model risk, such as a consistent bias-correction methodology across all portfolios, or statistical uncertainty in risk parameter calibration, especially when all portfolios share the same calibration window or modeling assumptions.

A summary of the distinctions between the types of risk with some examples are provided in Table **??**. Below we discuss several tangible examples of model risks and how they potentially interact across portfolios. Unlike credit risk, which reflects the economic exposure of borrowers and sectors, model risk affects the measurement of that exposure through the modeling process. While model risk introduces uncertainty in RWA estimates, this uncertainty does not necessarily accumulate at the bank-wide level. As will be argued and mathematically shown in later sections, model risk errors across portfolios diversify, meaning that individual estimation uncertainties tend to offset each other when aggregated. This diversification effect implies that the confidence level required for the MoC at the portfolio level tends to be lower than the target confidence level set for the bank-wide RWA. This allows for a more efficient and risk-sensitive allocation of conservatism, consistent with the bank's overall model risk appetite and capital adequacy goals.

Consider a bank with multiple credit portfolios, three of which are: mortgages in Spain, corporate loans in Germany, and SME exposures in Poland. Each portfolio is developed using its own dataset, and while all models follow a common framework, the specific risk drivers used (e.g., regional unemployment rates, sector-specific indicators) vary. Suppose that in the mortgage model for Spain, a relevant risk driver (e.g., housing price index) is omitted due to data limitations. This introduces an idiosyncratic bias in the RWA estimate for that portfolio. The corporate loan model in Germany might omit a different risk driver (e.g., industrial production index), while the SME model in Poland might suffer from noise or outdated information in borrower-level financial ratios. Such model errors are independent, meaning the individual biases tend to average out on the aggregated bank-wide RWA level. This is analogous to the law of large numbers, i.e., independent errors with zero mean and

Dimension	Credit Risk	Model Risk
Idiosyncratic	Borrower-specific default/loss risk (e.g., firm-level distress)	Portfolio-specific model flaws (e.g., missing risk driver values, poor data quality)
Systemic	Economy-wide or sector-wide downturns affecting many borrowers	Shared modelling assumptions or calibration uncertainty across portfolios
Nature of Risk	Economic reality of borrower exposures	Uncertainty in the model measurement of losses
Impact on RWA	Drives actual credit losses and portfolio risk	Affects accuracy and reliability of RWA estimates
Regulatory Treatment	Countercyclical buffer, stress testing	Margin of Conservatism (MoC), model validation and benchmarking

Table 1: Comparison of credit and model risk dimensions.

finite variance will converge toward a stable aggregate estimate. As will be shown in Section **??**, when assuming independence this aggregate uncertainty decreases with the number of portfolios.

Suppose now that the bank uses a consistent methodology to backscore default events across all its wholesale corporate portfolios, a common industry practice. This is then applied to historical data where default triggers, such 90 days-past-due (DPD) or insolvency filings, are not consistently recorded. The bank is thus required to infer default statuses retrospectively. If the backscoring algorithm systematically mis-identifies defaults due to, say, conservative thresholds or missing auxiliary indicators, then this bias will affect all corporate portfolios in a similar way. This shared source of error induces a systemic model risk. In reality, these backscored default rates will cover different historical observation windows, as no two portfolio development datasets will exactly overlap. Thus, not only are systemic model errors present, but each probability of default (PD) estimate will have a varying degree of exposure to the bank's systemic model risk factor. Nonetheless, these correlated biases in the estimation of default rates will propagate into correlated errors in the RWA calculations. Unlike idiosyncratic model risk, which may cancel out when aggregated, these correlated errors will reinforce each other at the bank-wide level. This dampens the diversification effect and implies that the confidence level required for the MoC per portfolio be closer to the bank-wide confidence level to maintain robustness.

The degree of model risk diversification across portfolios is also influenced by the concentration of RWA across a bank's loan book. When RWAs are heavily concentrated in a few portfolios, model risk errors from those portfolios dominate the aggregate uncertainty, weakening the diversification effect. Conversely, when RWAs are more evenly distributed, individual model errors contribute more uniformly, allowing their effects to offset each other more effectively. As demonstrated in Section **??**, maximum diversification occurs when RWAs are uniformly concentrated across portfolios.

In the following section, we combine all concepts introduced in the above discussions into a mathematical framework. This leads us to the derivation of the quantile scaling factor (QSF) in Section **??**, which allows for a one-to-one mapping of the portfolio MoC confidence level to the bank-wide confidence level, and vice versa.

3 Theoretical Framework

The Basel framework envisions that banks should hold capital to cover for unexpected losses of their loans (**?**), known as the Common Equity Tier (CET) 1 capital, whereas expected losses are meant to be captured by the bank's provisions (**?**). Given a loss distribution for a portfolio, unexpected losses are defined here as the difference between the loss at 99.9^{th} percentile and the expected loss. The formula used to compute this is directly provided by the regulatory risk-weighted assets (RWA) formula (**?**).

Using historical data, a bank develops regulatory credit risk models for its loan portfolios, the final estimates of which are applied to current and future portfolio snapshots. These models predict a given loan's probability of default (PD), loss given default (LGD), and exposure at default (EAD). Specifically, a loan is scored by a ranking function, followed by a mapping to a calibrated through-the-cycle (TTC)[1] estimate for all three risk parameters (**?**, **?**). These estimates then feed into the Basel formula, producing an RWA for the loan. The portfolio's RWA is taken as the sum of its individual loans' RWAs, which then implies that the bank's total RWA is the sum of the RWAs per portfolio.

3.1 An ASRF Model for Portfolio RWA Estimates

Fix a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Index portfolios in the bank's loan book by $i = 1, \ldots, N$. Let the idiosyncratic and systemic model risk factors be defined respectively as

$$\xi_i, \psi \sim \mathcal{N}(0,1), \quad \forall i$$

where each portfolio has its own systemic model risk factor loading ϕ_i.

Definition 3.1 (Portfolio Risk-Weighted Assets)**.** The modelled RWA for portfolio i is defined as

$$R_i = \mu_i + \sigma_i \eta_i, \tag{1}$$

$$\eta_i = \sqrt{1-\phi_i}\xi_i + \sqrt{\phi_i}\psi, \tag{2}$$

with $R_i \in (0, \infty)$ and $\eta_i \sim \mathcal{N}(0,1)$ $\forall i$.

The model error η_i is described by an ASRF model where ξ_i is the idiosyncratic model risk factor and ψ is the systemic model risk factor, both standard normally

[1]The terminology "through-the-cycle" is one commonly referred to in the regulation and industry and is the topic of many debates about its interpretation. For the purposes of this work, we adopt the notion of TTC understood as a risk parameter estimate that is "unconditional of the PIT (point-in-time) state of the economy".

distributed with asset correlation ϕ_i. Hence, the η_i are identically distributed standard normal errors with

$$\mathbb{E}[\eta_i] = 0, \quad \mathbb{V}[\eta_i] = 1, \quad \mathrm{Cov}[\eta_i, \eta_j] = \sqrt{\phi_i \phi_j} \equiv \rho_{ij} > 0.$$

This is not a typical ASRF setup. Each portfolio's RWA estimate will compose of a mixture of idiosyncratic and systemic model errors, but the dependence of each model to the systemic model risk factor is unique to each portfolio. This captures the nature of the problem. As such, the measured errors in portfolio RWAs due to ψ are scaled per portfolio by σ_i, as well as each portfolio's co-dependence on the systemic model risk factor through ϕ_i. Naïvely, the ρ_{ij} can be thought of as a parameter that captures the overlap in systemic model risk exposure of two modelled RWA estimates for portfolios i and j. As we will see in later sections, we can derive all necessary relationships without explicitly relying on the underlying ASRF model.

3.2 Asymptotic Normality

While a consequence of the ASRF model, it is necessary to explicitly state the asymptotic distribution of modelled portfolio RWAs.

Theorem 3.1 (Portfolio-level asymptotic normality). *Let $\theta_0 = (PD_0, LGD_0, EAD_0, \xi_0)$ be the true parameter vector for a portfolio, and $h(\theta)$ the IRB map returning portfolio RWA. By the central limit theorem (CLT), as the number of obligors in the portfolio $n \to \infty$,*

$$\sqrt{n}\,(\hat{\theta}_n - \theta_0) \to \mathcal{N}(0, \Xi)$$

with h continuously differentiable at θ_0 with gradient $\nabla h(\theta_0)$. Then,

$$\sqrt{n}\left(h(\hat{\theta}_n) - h(\theta_0)\right) \to \mathcal{N}\left(0,\ \nabla h(\theta_0)^\top \Xi\, \nabla h(\theta_0)\right).$$

Proof. ■

Corollary 3.1.1 (Bank-level asymptotic normality). *Stack portfolios $i = 1, \ldots, N$ with parameter vectors $\theta_{0,i}$ and estimators $\hat{\theta}_{n,i}$ satisfying Theorem* **??** *jointly. Denote ι the $N \times 1$ vector of ones. Let $RWA_i = h(\hat{\theta}_{n,i})$ and $RWA_{\text{bank}} = \sum_{i=1}^{N} RWA_i := \mu_{\text{bank}}$. Then, as $n \to \infty$*

$$\sqrt{n}\left(\mu_{\text{bank}} - \sum_{i=1}^{N} h(\theta_{0,i})\right) \ \to\ \mathcal{N}\left(0, \sigma^2_{\text{bank}}\right),$$

where

$$\sigma^2_{\text{bank}} = \iota^\top \nabla H(\theta_0)^\top\, \Xi_{\text{bank}}\, \nabla H(\theta_0)\, \iota \quad \text{with} \quad H(\theta_0) := (h(\theta_{0,1}), \ldots, h(\theta_{0,N}))^\top,$$

and Ξ_{bank} is the covariance of the stacked influence vector across portfolios.

Proof. ■

Theorem **??** is proven naturally via the delta method. The bank-level statement follows by linearity after stacking portfolios and applying the multivariate delta method with the joint covariance $\boldsymbol{\Xi}_{\text{bank}}$ of the stacked influence vector. The vector of ones sum the marginal (co)variance contributions per portfolio into a scalar variance for the loan book. The gradient components $\partial h/\partial PD$, $\partial h/\partial LGD$, $\partial h/\partial EAD$ admit closed forms, while correlations among parameter estimators are preserved through $\boldsymbol{\Xi}$. Theorem **??** is also shown to hold true numerically in Appendix **??**.

3.3 Margin of Conservatism

Although there is no explicit regulatory guidance on how to quantify the MoC for a given risk parameter, it widely accepted that this represents an upward stress with respect to the portfolio's best-estimate RWA.

Definition 3.2 (Stressed portfolio RWA)**.** Let the confidence level for portfolio i be κ_i. Given estimates for the mean and variance of R_i, define the stressed RWA, $\widetilde{\mu}_i$ as

$$\widetilde{\mu}_i = \mu_i + \Phi^{-1}\left(\kappa_i\right)\sigma_i. \tag{3}$$

Definition 3.3 (Portfolio Margin of Conservatism)**.** Define the margin of conservatism (MoC) on the RWA of portfolio i as the relative increase in RWA resulting from the upward stress given in Equation (**??**):

$$\beta_i(\kappa_i) = \frac{\widetilde{\mu}_i - \mu_i}{\mu_i} = \Phi^{-1}\left(\kappa_i\right)\frac{\sigma_i}{\mu_i}. \tag{4}$$

The volatility of the portfolio RWA estimate can be expressed in terms of the best-estimate RWA, the MoC, and the coverage level of the portfolio,

$$\sigma_i = \frac{1}{\Phi^{-1}\left(\kappa_i\right)}\mu_i\beta_i(\kappa_i). \tag{5}$$

Definition 3.4 (Portfolio Regulatory Capital)**.** The regulatory capital for portfolio i is defined as a stressed best-estimate,

$$\widetilde{\mu}_i = \mu_i\left(1 + \beta_i(\kappa_i)\right). \tag{6}$$

Consider that the bank wishes that its total regulatory capital has coverage equivalent to the quantile level κ_{bank}, aligning with its model risk appetite framework. A similar definition for the MoC at the total bank-level can thus be defined.

Definition 3.5 (Bank-wide Margin of Conservatism)**.** Using Corollary **??**, the bank-wide MoC is defined as

$$\beta_{\text{bank}} = \Phi^{-1}\left(\kappa_{\text{bank}}\right)\frac{\sigma_{\text{bank}}}{\mu_{\text{bank}}}. \tag{7}$$

Definition 3.6 (Bank-wide Regulatory Capital)**.** The bank's total regulatory capital is defined as a stressed best-estimate,

$$\widetilde{\mu}_{\text{bank}} = \mu_{\text{bank}}\left(1 + \beta_{\text{bank}}\right). \tag{8}$$

3.4 Local Linearity of the RWA

Regulation expects the MoC percentage leading to conservative RWAs to be estimated per risk parameter, not on the RWA directly. Hence, while the IRB Basel formula is linear in LGD and EAD, it is nonlinear in PD, i.e., errors in PD estimates do not propagate linearly through to the estimated RWA. We now state and show that the estimated RWA is locally linear to first-order for any asset class. For this we use the general Basel IRB formula for RWA. Numerical verification of the resulting error bound is provided in Appendix **??**.

Theorem 3.2 (Local linearity of RWA in PD)**.** *Fix an asset class which is C^2, correlation $R(\cdot)$, and maturity adjustment $M(\cdot)$ on a compact set $K = [\underline{p}, \overline{p}] \subset (0,1)$. Let*

$$RWA(PD) = 12.5 \times \, M(PD)\, LGD \left[\Phi(A(PD)) - PD\right] \times EAD,$$

where $A(PD) = \frac{\Phi^{-1}(PD)}{\sqrt{1-R(PD)}} + \sqrt{\frac{R(PD)}{1-R(PD)}}\, G$ with $G = \Phi^{-1}(0.999)$, and $M(PD)$ is the IRB maturity adjustment. Then $RWA \in C^2(K)$ and, for any $PD_0 \in K$,

$$RWA(PD) = RWA(PD_0) + RWA'(PD_0)(PD - PD_0) + \tfrac{1}{2} RWA''(\xi)\,(PD - PD_0)^2,$$

for some ξ between PD and PD_0. In particular, the linearization error is $\mathcal{O}\big((PD - PD_0)^2\big)$ uniformly on K.

Proof. See Appendix **??**. ■

Theorem **??** allows us to condense the MoCs estimated per risk parameter into one multiplicative factor on the best-estimate RWA for portfolio i:

$$\widetilde{\mu}_i = \mu_i(1 + \beta_{i,PD}(\kappa_i))(1 + \beta_{i,LGD}(\kappa_i))(1 + \beta_{i,EAD}(\kappa_i)) \approx \mu_i(1 + \beta_i(\kappa_i)), \tag{9}$$

where $\beta_{i,PD}$ is the MoC of the best-estimate PD, and respectively for LGD and EAD.

3.5 Fair Allocation of Conservatism

Recalling Equation (**??**), each portfolio receives an upward stress to its estimated RWA depending on the chosen quantile κ_i. With no restrictions applied the problem has infinite solutions. In other words, banks may arbitrarily choose any combination of quantiles κ_i that collectively reach κ_{bank} allowing for RWA arbitrage. Thus, from a policy perspective, equal quantiles for all portfolios should be enforced, $\kappa_i \equiv \kappa \; \forall i$. This notion of absolute fairness in allocation of RWA conservatism can be formalised. First, we define the RWA concentration as the share of the total bank RWA predicted in portfolio i.

Definition 3.7 (Portfolio concentration of RWA)**.** The concentration of RWA in portfolio i is defined as:

$$\gamma_i = \frac{\mu_i}{\sum_{j=1}^{N} \mu_j}. \tag{10}$$

Proposition 3.1 (Absolute fairness of equal quantiles). *Let $\mu_i > 0$ denote the long-run mean RWA of portfolio i and $\sigma_i > 0$ the standard deviation of its model-error (or chosen dispersion scale). Define the* coefficient of variation *$m_i := \sigma_i/\mu_i$ and denote the RWA concentration as γ_i. Define portfolio weights*

$$w_i \;:=\; \gamma_i \, m_i \;>\; 0.$$

For standardised MoC burdens $z_i := \Phi^{-1}(\kappa_i) \geq 0$ and a standardised bank-wide target $Z = \Phi^{-1}(\kappa_{bank}) > 0$ for $\kappa_{bank} \in (0,1)$, consider the feasible set

$$\mathcal{F} \;=\; \Big\{ z \in \mathbb{R}_+^N : \; \sum_{i=1}^{N} w_i z_i \;\geq\; Z \Big\}.$$

The minimax fairness problem

$$\min_{z \in \mathcal{F}} \; \max_{1 \leq i \leq N} z_i$$

has the unique solution $z_i^\star = t^\star$ for all i, where

$$t^\star \;=\; \frac{Z}{\sum_{j=1}^{N} w_j}.$$

Equivalently, the fair allocation in quantile space is $\kappa_i^\star = \Phi(t^\star) \equiv \kappa$ for all i.

Proof. See Appendix **??**. ∎

Proposition **??** is fundamentally a fairness statement, thus formalising a policy stance a bank can take. It is not a capital efficiency claim. With $w_i = \gamma_i m_i$ and $z_i = \Phi^{-1}(\kappa_i)$, the max–min problem chooses the allocation that makes the most stringent standardised requirement as small as possible. In practical terms, this means that no portfolio is asked to carry a higher standardised MoC threshold than any other; the unique way to achieve this is to set equal quantiles, $\kappa_i = \kappa \, \forall i$.

This result does not imply that alternative allocations would use more (or less) total capital. Under local linearity of the RWA estimate (Theorem **??**), the bank-wide constraint binds at $\sum_i w_i z_i = Z$, so any allocation on this boundary has the same first-order aggregate MoC capital. Proposition **??** therefore governs the distribution of "conservatism burden" across portfolios instead of the aggregate amount. Differences in business efficiency arise only from second-order effects such as regulatory constraints (e.g., floors/caps) or managerial preferences about which portfolio should bear relatively more conservatism.

3.6 Bank-wide RWA

We now describe the properties of the bank-wide RWA estimate and discuss its implications. Collect $R = (R_1, \ldots, R_N)^\top$, $\mu = (\mu_1, \ldots, \mu_N)^\top$, $\sigma = (\sigma_1, \ldots, \sigma_N)^\top$, $\gamma = (\gamma_1, \ldots, \gamma_N)^\top$, and $\eta = (\eta_1, \ldots, \eta_N)^\top$. Let ι denote a vector of ones of conformable dimension. Let $\mathbf{P} = [\rho_{ij}]$ denote the symmetric, positive semidefinite error correlation matrix such that $\mathbf{\Sigma} = \operatorname{diag}(\sigma)\, \mathbf{P} \operatorname{diag}(\sigma)$. Write $\beta = (\beta_1, \ldots, \beta_N)^\top$, such that $\mathbf{B} := \operatorname{diag}(\beta)$. Define

$$\mathbf{\Psi} := \mathbf{B}^\top \mathbf{P} \mathbf{B}.$$

Where $\mathbf{\Sigma}$ and $\mathbf{\Psi}$ differ is that the former represents the unknown covariance matrix for all portfolio RWAs, whereas the latter is a scaled estimate expressed as a "coefficient of variation matrix".

Theorem 3.3 (Bank-wide RWA moments and MoC aggregation)**.** *Under* (**??**) *and* (**??**),

$$\mu_{\text{bank}} = \sum_{i=1}^{N} \mu_i, \qquad \sigma^2_{\text{bank}} \;=\; \frac{1}{\left(\Phi^{-1}(\kappa)\right)^2}\, \mu^\top \mathbf{\Psi}\, \mu,$$

and the bank-level MoC aggregates as

$$\beta_{\text{bank}} \;=\; \sum_{i=1}^{N} \gamma_i\, \beta_i.$$

Proof. See Appendix **??**. ■

Theorem **??** is an important result. We can estimate the unknown variance of the bank-wide RWA estimate through a scaled sum of each portfolio's estimated RWA variance. The theorem also establishes how the MoC estimated per portfolio contributes to the bank-wide MoC through the proportion of the bank-wide estimated RWA contained within each portfolio. In the next section, this result will be utilised to link the bank-wide confidence level, κ_{bank}, to the confidence level set for each portfolio, κ.

4 The Quantile Scaling Factor

The previous section introduced all the variables necessary for the derivation of the quantile scaling factor (QSF), which we will call q. The RWA estimate and its corresponding MoC at both portfolio and total bank level were defined along with their statistical properties. In this section, we combine these results to define the QSF and investigate its properties. The main results will be presented first with their interpretations, thereafter followed by examples to showcase how the QSF can be used in practice.

Definition 4.1 (Quantile Scaling Factor)**.** For two quantiles $0.5 < \kappa \leq \kappa_{\text{bank}} < 1$ of the standard normal distribution Φ, the quantile scaling factor (QSF) is defined as:

$$q = \frac{\Phi^{-1}(\kappa)}{\Phi^{-1}(\kappa_{\text{bank}})} \in (0, 1], \tag{11}$$

The QSF is a ratio of two statistical coverage levels. Assuming the target bank quantile, κ_{bank}, is set by the bank's model risk appetite framework and the bank has an estimate for its QSF, $\hat{q}$, the implied portfolio quantile required to attain κ_{bank} is thus given as:

$$\hat{\kappa} = \Phi\left(\hat{q}\,\Phi^{-1}(\kappa_{\text{bank}})\right). \tag{12}$$

Figure **??** shows how κ varies for different combinations of the QSF and target bank-wide quantile, κ_{bank}. Consider the limiting cases where $\hat{q} \to 0$ and $\hat{q} \to 1$. In the

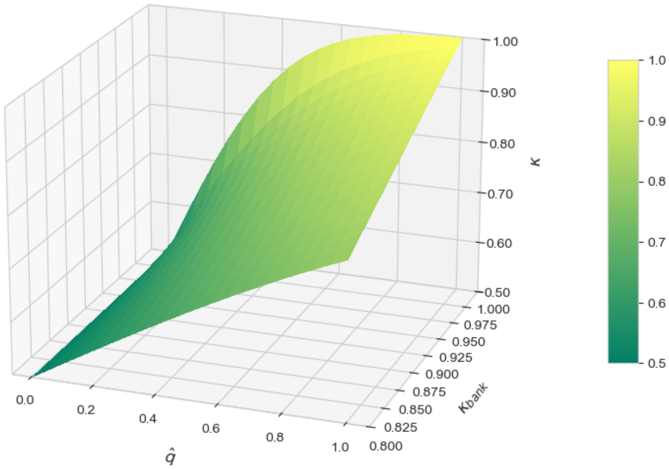

Figure 1: Implied portfolio quantile, κ, for different combinations of $\hat{q}$ and κ_{bank}. The target bank-wide quantile varies from 80% to 99.9%, whereas the QSF is in $(0, 1)$.

former, regardless of how risk averse the target bank-wide quantile is, the portfolio quantile remains low at $\hat{\kappa} \approx 0.5$. This is verified by inspecting Equation (**??**). In other words, the portfolio best-estimate RWA provides full statistical coverage, or a full diversification of risk. This also implies zero uncertainty in the total bank-wide RWA; an unrealistic situation. For the latter case Equation (**??**) becomes $\kappa \to \kappa_{\text{bank}}$. To reach κ_{bank} each portfolio must have an MoC quantified at that same quantile. This situation corresponds to no model risk diversification. In essence, the QSF represents the strength of model risk diversification present in the bank's loan book.

Theorem 4.1 (QSF)**.** *Let $\kappa_{\text{bank}} \in (0, 1)$ be the bank-level target and denote the Hadamard product by $\odot$. The QSF as given in Definition* **??** *satisfies*

$$q = \frac{\sqrt{(\gamma \odot \beta)^\top \mathbf{P} (\gamma \odot \beta)}}{\iota^\top (\gamma \odot \beta)} = \frac{\left\|\mathbf{P}^{1/2}(\gamma \odot \beta)\right\|_2}{\left\|\gamma \odot \beta\right\|_1}, \tag{13}$$

Proof. See Appendix **??**. ■

Theorem 4.2 (Bounds and extremal structures)**.** *Let $a = \gamma \odot \beta \in \mathbb{R}^N_+$ and $\mathbf{P}$ a correlation matrix. Then*

$$0 \leq q \leq 1, \qquad \sqrt{\lambda_{\min}(\mathbf{P})}\,\frac{\|a\|_2}{\|a\|_1} \leq q \leq \sqrt{\lambda_{\max}(\mathbf{P})}\,\frac{\|a\|_2}{\|a\|_1}.$$

Moreover:

- *(Fully systemic model errors) If $\mathbf{P} = \iota\iota^\top$, then $q = 1$.*
- *(Fully idiosyncratic model errors) If $\mathbf{P} = \mathbf{I}_N$, then $q = \|a\|_2/\|a\|_1$.*

Proof. See Appendix **??**. ■

Corollary 4.2.1 (Schur-convexity of q under uniform RWA concentration)**.** *Under* $\mathbf{P} = \mathbf{I}_N$*, if* γ_i *are also all equal with* $\sum_i \gamma_i = 1$*, then*

$$\gamma_i = \frac{1}{N}\,\forall i \;\Rightarrow\; q = \frac{\|\beta\|_2}{\|\beta\|_1} \quad \textit{with} \quad \frac{1}{\sqrt{N}} \le q \le 1.$$

The QSF q *is thus Schur-convex and decreases toward* $1/\sqrt{N}$ *as MoC estimates become more similar. In other words, under the conditions stated above, less dispersed MoC values across portfolios maximise the bank's model risk diversification.*

Proof. See Appendix **??**. ■

Theorem **??** provides the main result of this work. The QSF (**??**) is a dispersion ratio: an ℓ_2 norm (under $\mathbf{P}$) over an ℓ_1 norm of $\gamma \odot \beta$. It equals 1 when model risk is fully systemic, while it falls toward $1/\sqrt{N}$ as errors in RWA estimates become more idiosyncratic and γ equalises. More precisely, under fully idiosyncratic model errors, uniform RWA concentrations, and equivalent portfolio MoC estimates[2], $q \to 0$ as $N \to \infty$.

The dependence on the portfolio quantile κ through β also drops in Equation (**??**). Thus, only information about each portfolio's RWA concentration and error correlations is required to understand the level of model risk diversification in the bank's loan book. From this, with a set model risk appetite, the bank can directly determine the quantile required per portfolio, κ.

5 Understanding the QSF and its Properties

In this section we provide several simulated examples aimed at better understanding how the QSF in Equation (**??**) behaves. We will set and vary three key characteristics of the bank's loan book to do this, namely, the number of portfolios, RWA error correlations, and RWA concentrations. MoC values will be sampled based on typical values observed in practice following certain basic principles that must be met. Note that information about neither the bank's model risk appetite nor the portfolio RWA distribution is required to conduct these analyses. Thereafter, in Section **??** we will construct a more realistic loan book that will estimate correlations of RWA errors and feed these into the QSF formula to produce an implied portfolio quantile usable by the bank.

5.1 MoC Sampling

There are two main conditions that the sampled β_i should satisfy to ensure these represent realistic MoC percentages. They must first be strictly postive, i.e., nonzero and non-negative. This is also a regulatory expectation: MoCs must represent a conservative stress on RWA estimates. Secondly, the majority of the density should

[2]It should also be noted that Schur-convexity of q no longer holds if γ is not uniform. This is due to the asymmetry of the resulting dispersion in $\gamma \odot \beta$.

sit between 1-40%. While excessively high MoCs > 40% are rare, they are still possible.

To satisfy the first condition, the chosen distribution must have support on $\mathbb{R}$. Satisfying the the second condition requires a right-skewed and long-tailed distribution, i.e., unbounded above. To retain interpretability of the hyperparameters the lognormal distribution is chosen,

$$\beta_i \sim \text{LogN}(\nu, \tau).$$

For the analyses in the remaining sections we will set $\nu = 0.2$ and $\tau = 0.1$. It should be noted that independent samples with the same parameters are drawn for all portfolios. Hence, on average, β will be equivalent across portfolios, inducing a minor amount of model risk diversification.

5.2 RWA Error Correlation and Concentration

Consider a bank with $N = 2$ portfolios. We define two correlation structures between their RWA errors, η_i,

$$\mathbf{P}_0 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{P}_1 = \begin{bmatrix} 1 & 0.95 \\ 0.95 & 1 \end{bmatrix}.$$

These define the fully idiosyncratic (maximal diversification) and strongly systemic model errors, respectively. Let the RWA concentration vector vary as

$$\gamma = \begin{bmatrix} \varepsilon \\ 1 - \varepsilon \end{bmatrix}, \; \varepsilon \in (0, 1).$$

Consider two scenarios. One where RWA concentration is uniform, $\varepsilon = 0.5$, and the other where RWA is concentrated in portfolio $i = 1$, $\varepsilon = 0.95$. Taking 1000 MoC samples per β_1 and β_2 as described earlier and using Equation (**??**), we get distributions for q for each combination of scenarios. These are shown in Figure **??**.

As expected, fully idiosyncratic model errors lead to a low QSF and vice versa, with little diversification present for strongly systemic model errors. The diversification effect is dampened when RWA concentration in a single portfolio is high. To better understand the effect of concentration, we let ε vary in $(0, 1)$ and take the mean of the simulated distribution for q, which we will call $\bar{q}$. The resulting plot is given in Figure **??**.

We make the following observations. The QSF attains its minimum when RWA concentration is uniform regardless of the error correlation between the two portfolios, confirming Theorem **??**. For strongly systemic RWA errors the QSF stays close to 1 regardless of the RWA concentration. This is offset when RWA concentration is high. Finally, the peaked distributions for q when concentration is uniform suggests that RWA concentration across portfolios has a larger effect on model risk diversification than the error correlations between them.

5.3 Number of Portfolios

We now turn to the total number of portfolios. For this analysis we will increase N from 2 until 50, while restricting RWA errors to be fully idiosyncratic and portfolios

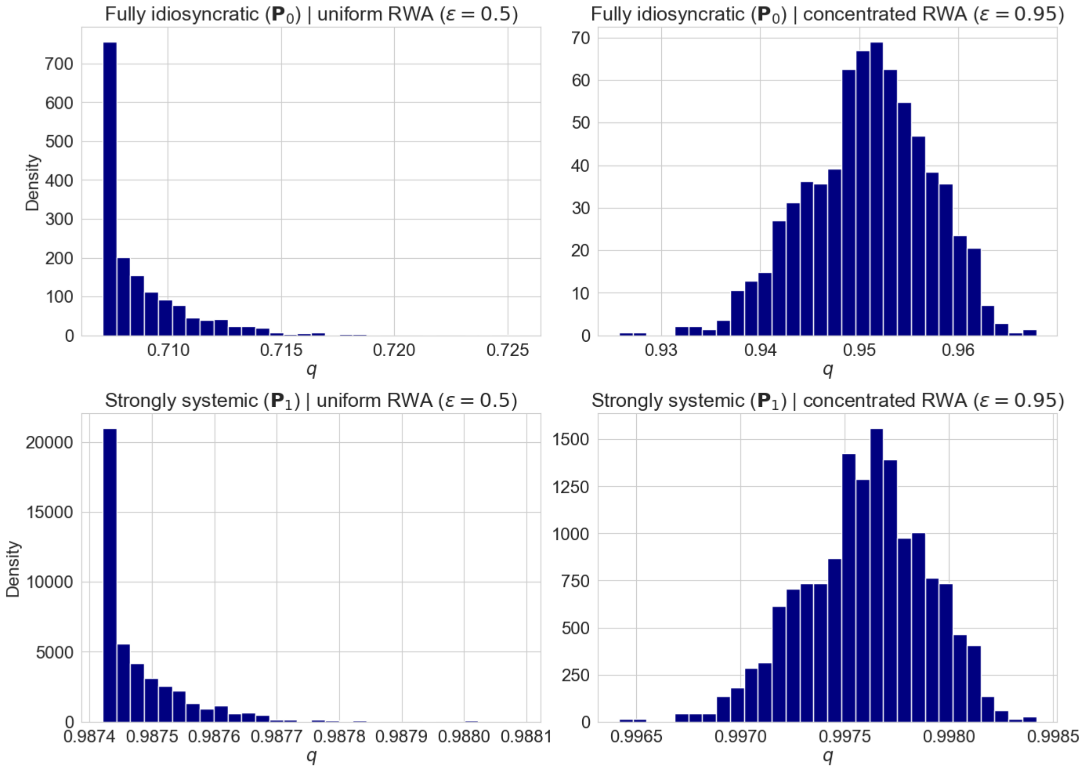

Figure 2: Simulated distributions for the QSF given four combinations of RWA concentrations and error correlations.

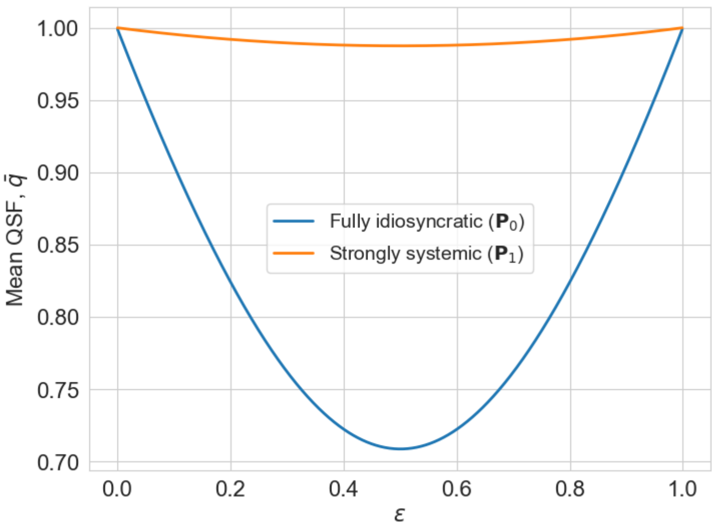

Figure 3: Effect of RWA error correlation and concentration on the QSF for a bank with $N = 2$ portfolios.

to have uniform RWA concentrations,

$$\mathbf{P} = \mathbf{I}_N, \quad \gamma = \begin{bmatrix} 1/N \\ 1/N \\ \vdots \\ 1/N \end{bmatrix}.$$

Sampling β_i for $i = 1, \ldots, N$ as before and taking the mean of the simulated QSF distribution, we get the relationship shown in Figure ??. We first note that the resulting function is exactly equivalent to $1/\sqrt{N}$, as expected by Corollary ??. Holding all else fixed and increasing the total number of portfolios leads to a sub-

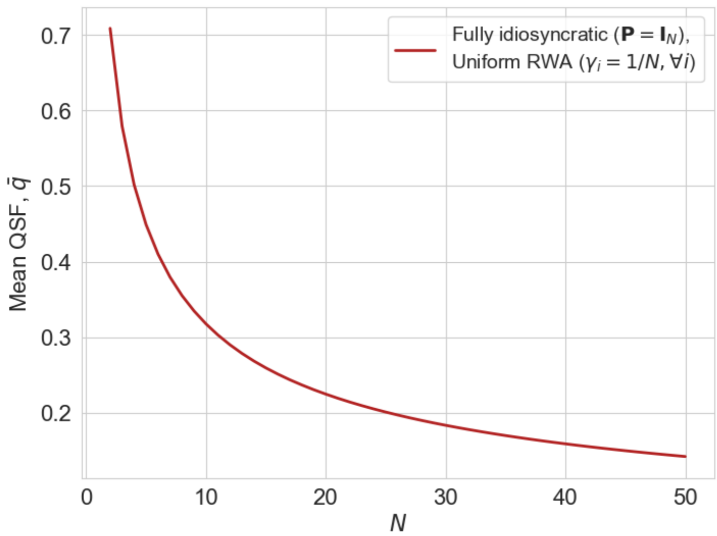

Figure 4: Effect of the total number of portfolios on the QSF.

stantial model risk diversification effect. The effect on the QSF begins to dissipate fairly quickly, with the majority of the benefit realised up to $N = 10$ portfolios.

6 Portfolio Quantile Estimation

We now investigate a realistic scenario and showcase how the above analyses lead to an implied quantile that is usable by a bank. To illustrate empirical applicability, we apply the QSF framework to a synthetic loan book calibrated to public regulatory disclosures, designed to mimic a large diversified bank's RWA profile while containing no proprietary information. It is given that the bank's model risk appetite dictates its bank-wide coverage, κ_{bank}, and utilises the result in Equation (??). We compute the implied portfolio quantile, κ, assuming three levels of risk appetite: $\kappa_{\text{bank}} = 85\%,\ 95\%,$ and 99%[3].

6.1 Synthetic Loan Book

A synthetic loan book comprising of $N = 10$ representative portfolios is constructed with yearly observations from 2014–2024. Portfolio sizes, PD ranges, LGD means

[3]These percentages are chosen because they are roughly equidistant on $\mathbb{R}$ when Φ^{-1} is applied to them, representing equally large "jumps" in model risk appetite.

and EAD distributions are calibrated to publicly available supervisory aggregates and disclosure templates. Specifically, these are the EBA EU-wide Transparency Exercise and the ECB's Supervisory Banking Statistics, which report aggregate RWA density and sectoral statistics. The Basel risk-weight function is used to compute true RWAs. Realised defaults per year are drawn via binomial sampling using the true yearly PDs, realised LGDs are drawn from Beta distributions with portfolio-specific means and variances, and EADs per obligor are drawn from log-normal distributions calibrated to produce realistic exposure concentration. An ASRF model is employed to induce variability in the true PDs, while we will assume LGDs and EADs to be deterministic through time. Specifically, each portfolio PD has an idiosyncratic credit risk factor and, through different asset correlations, is also shocked via a systemic macro factor. All details of how the loan book was generated are provided in Appendix **??**.

6.2 Estimating RWAs and Model Error Correlations

To produce estimates of the QSF, create modelled RWA estimates for each portfolio–year in the loan book. We consider only a PD model while fixing LGDs to the portfolio mean and EADs to the generated totals. We will assume that modelled PDs are calibrated to the long-run average default rate (LRA DR) for each portfolio:

$$\widehat{PD}_i \equiv LRADR_i = \frac{1}{T}\sum_{t=1}^{T} ODR_{i,t}.$$

We induce model risk variability in the predicted PDs by applying an ASRF model, whereby each portfolio has a unique systemic model risk factor loading. Each portfolio's predicted PD is thus subject to unique idiosyncratic and systemic model risk volatility. Modelled portfolio PDs are sampled from a Beta distribution with unique parameters (α_i, β_i) that match the predicted PD and the pre-specified model risk factor volatilities and loadings. Beta parameters are first computed as:

$$\alpha_i = \left(\frac{1-\widehat{PD}_i}{\sigma^2_{\text{idiosync}}} - \frac{1}{\widehat{PD}_i}\right)\widehat{PD}_i^2, \quad \beta_i = \frac{\alpha_i}{\frac{1}{\widehat{PD}_i} - 1},$$

which are then used to produce a sample of the modelled PD via:

$$\widehat{PD}_i^{(k)} \sim \text{logit}^{-1}\left(\text{logit}\left(\text{Beta}(\alpha_i, \beta_i)\right) + \sqrt{c\,\phi_i}\psi,\right)$$

where $\psi \sim \mathcal{N}(0,1)$ is the systemic model risk factor with portfolio-specific factor loadings ϕ_i and $c \in \mathbb{R}^+$. The specific values used are found in Table **??**, which also contain the chosen hyperparameters used per portfolio's MoC sample distribution. The sets of (ν, τ) are chosen based on what is typically observed in practice and to allow for unequal MoC distributions.

The final PD samples are then fed into the Basel formula corresponding to the portfolio's asset class, using the mean portfolio LGD and total EAD at each year to simulate predicted RWAs:

$$\hat{R}_{i,t}^{(k)} = 12.5 \times g\left(\widehat{PD}_i^{(k)}, M = m\right) \times \overline{LGD}_{i,t} \times EAD_{i,t}^{\text{tot}},$$

Portfolio	σ_{idiosync}	ϕ_i	ν_i	τ_i
Residential mortgages	0.002	0.05	0.07	0.03
Retail unsecured	0.006	0.08	0.15	0.05
SMEs	0.009	0.12	0.12	0.06
Large corporates	0.011	0.01	0.30	0.10
Commercial real estate	0.004	0.04	0.18	0.07
Sovereigns	0.008	0.06	0.32	0.09
Financial institutions	0.024	0.09	0.22	0.08
Auto loans	0.019	0.13	0.11	0.04
Project finance	0.001	0.17	0.26	0.08
Specialty consumer finance	0.017	0.03	0.16	0.04

Table 2: Values for the idiosyncratic model factor volatilities, systemic model factor loadings, and MoC sample hyperparameters per portfolio.

where the maturity m is set to 1 for retail asset classes and to 2.5 for corporate asset classes. While the project finance and commercial real estate portfolios technically fall under specialised lending and thus use supervisory slotting, for the purpose of this analysis they are treated as corporate asset classes. Asset correlations in the IRB formula are also retained for each portfolio's asset class.

The true RWA, $R_{i,t}$, is given in the synthetic loan book by using the true, latent PD series. RWA errors are thus estimated as:

$$\hat{\eta}_{i,t} = \hat{R}_{i,t}^{(k)} - R_{i,t},$$

from which we can estimate the error correlations, $\hat{\rho}_{ij}$, at every year. Spearman rank correlations are used to negate scale issues with large RWA values. An example heatmap for the year 2024 is provided in Figure ??. It should be noted that these errors will contain some residual correlations due to the systemic macro shocks. In reality, it is true that modelled portfolio RWAs will exhibit procyclicality, but as we assume TTC calibration of risk parameters, model errors are expected to be dominated by model risk factors. As RWA concentrations in a bank's loan book is determined by predicted RWAs, these are estimated at each year, t, by averaging across all samples, k:

$$\hat{\gamma}_{i,t} = \frac{\frac{1}{K}\sum_{k=1}^{K} \hat{R}_{i,t}^{(k)}}{\sum_{i=1}^{N}\left(\frac{1}{K}\sum_{k=1}^{K} \hat{R}_{i,t}^{(k)}\right)}.$$

Finally, to compute the QSF at each year, MoC samples are produced following the method described in Section ??. The estimated QSF in year t is then taken as the mean of the resulting QSF distribution.

6.3 Results

The above setup is run on the synthetic loan book for $K = 1000$ modelled PD samples and $J = 1000$ MoC samples at each year. The resulting time series for the estimated QSF, $\hat{q}_t$, is given in Figure ??. The bank may then choose to take the time-weighted average of the QSF time series to estimate its QSF.

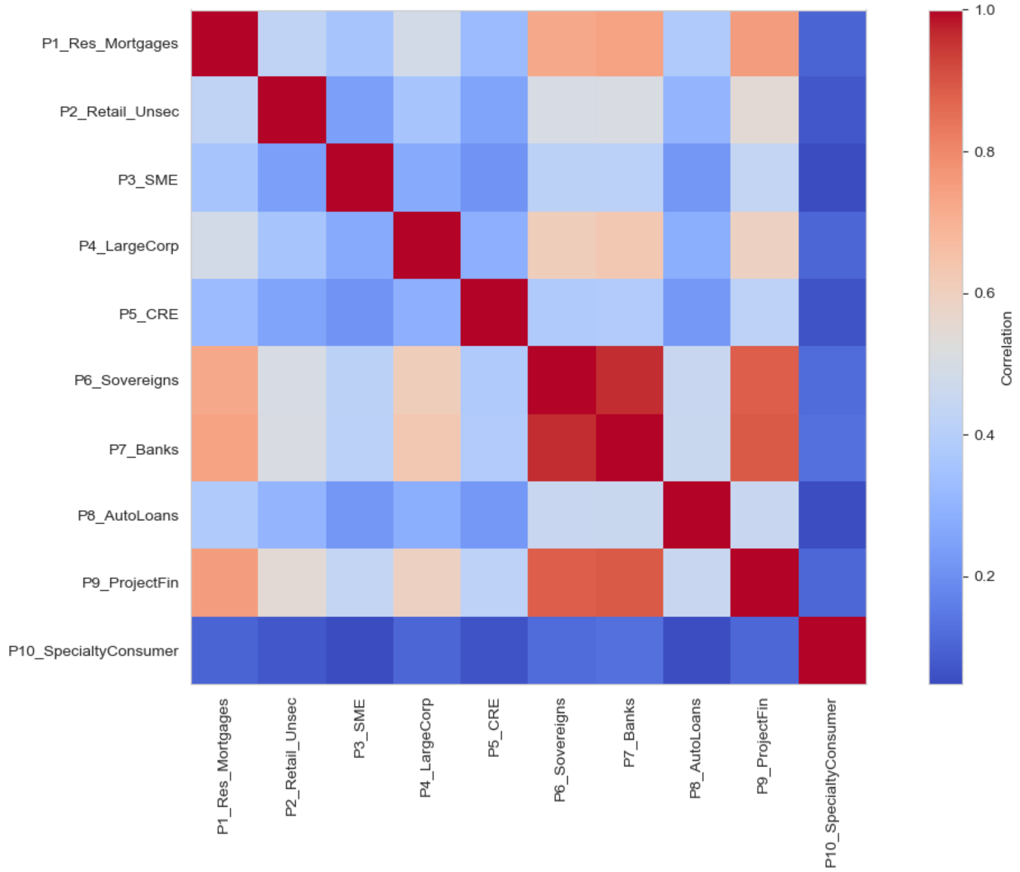

Figure 5: Estimated RWA error correlations for the synthetic loan book in 2024.

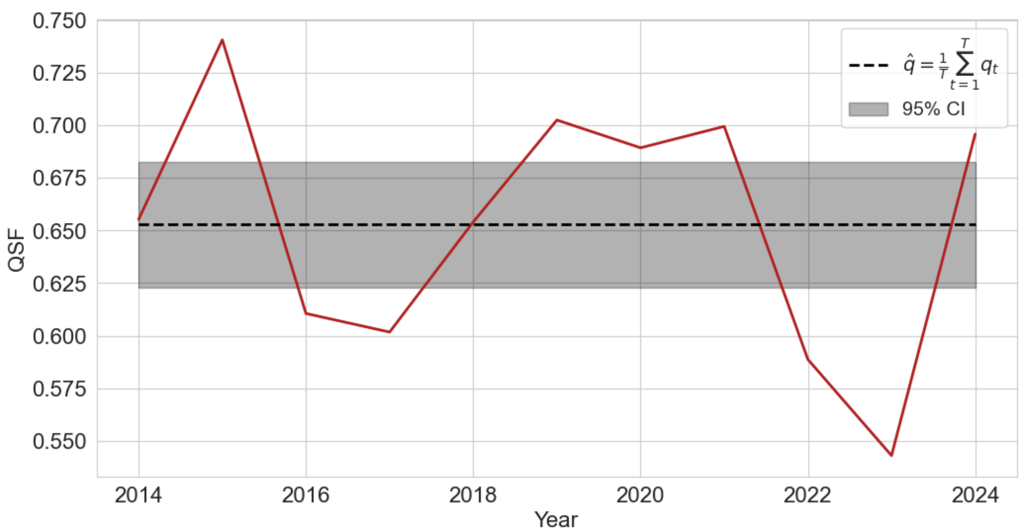

Figure 6: Estimated QSF time series with a long-run average and a 95% confidence band for the synthetic loan book.

A robustness analysis is conducted by scaling the systemic model factor loadings by the value c and computing standard errors on the resulting implied portfolio quantiles, κ, given three model risk appetites. Standard errors on $\hat{\kappa}$ are approximated by an application of the delta method to Equation (**??**):

$$s_{\hat{\kappa}}\big|_{\kappa_{\text{bank}}} \approx \left|\frac{\partial \hat{\kappa}}{\partial \hat{q}}\right| s_{\hat{q}} = \Phi\left(\kappa_{\text{bank}}\right)\, \phi\left(\hat{q}\, \Phi\left(\kappa_{\text{bank}}\right)\right)\, s_{\hat{q}}.$$

The results are shown in Table **??** and Figure **??**. The sensitivity of the estimated $\hat{\kappa}$ to scaling of the systemic model risk factor loadings is shown in Figure **??**.

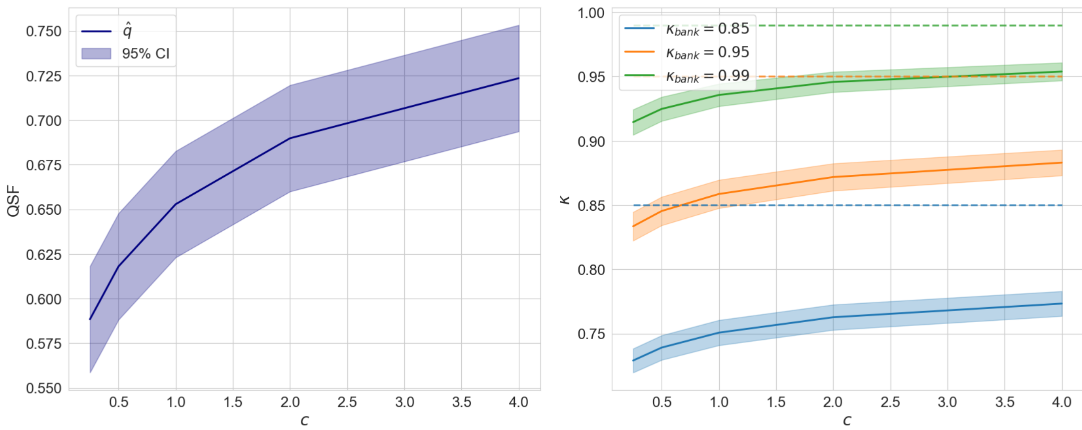

Figure 7: Estimated QSF for scaled model risk factor loadings (left); and implied portfolio quantile for three model risk appetites (right). Standard errors use for the 95% confidence bands around κ are estimated by applying the delta method to Equation (**??**).

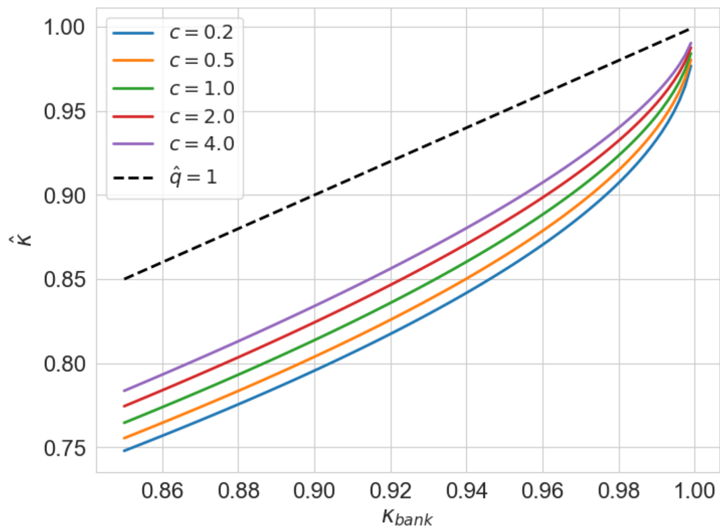

Figure 8: Estimated portfolio quantiles for different scalings of the model systemic risk factor loadings. Fully systemic model errors (dashed) shown for comparison.

As expected, model risk diversification is present, lowering the necessary confidence level per portfolio. Portfolios with higher systemic model risk asset correlations predictably have higher RWA error correlations, but its effect is offset by lower

RWA concentrations. The QSF is generally robust to the systemic model risk factor loadings, and its dependence on systemic errors in model estimates diminishes the more averse the bank's model risk appetite.

		κ		
c	$\hat{q}$	$Q_{\text{bank}}(0.85)$	$Q_{\text{bank}}(0.95)$	$Q_{\text{bank}}(0.99)$
0.25	0.645	0.748	0.855	0.933
0.50	0.667	0.755	0.864	0.940
1.00	0.696	0.765	0.874	0.947
2.00	0.727	0.774	0.884	0.955
4.00	0.757	0.778	0.893	0.961

Table 3: Implied portfolio quantiles necessary to reach the bank-wide confidence level corresponding to three model risk appetites, Q_{bank}.

7 Discussion

The previous sections proposed the QSF, investigated its properties, and showcased a simulated example of how it can be applied in practice. We now discuss the implications of the results, provide practical guidance on how to apply the framework, and highlight current gaps for future research.

The results of the previous sections show how correlated RWA estimation errors, RWA concentration, and total number of portfolios in the loan book contribute to the level of model risk diversification. The level of diversification determines the appropriate statistical coverage necessary per portfolio for a bank's model risk appetite strategy to be met. While the aforementioned are the key identified drivers, they effectively boil down to one single characteristic: the primary source of model uncertainty faced by the bank. More specifically, how much RWA mis-estimation across models is driven by systemic error sources as opposed to idiosyncratic.

Recalling Theorem **??**, should model errors be fully driven by a known systemic bias, then no diversification is present and portfolio MoC quantification should match the bank-wide model risk appetite strategy, $\kappa = \kappa_{\text{bank}}$. This is an unrealistic scenario. Every portfolio will have its own sources of uncertainty that are unique. In fact, it is more reasonable to assert that the majority of RWA estimation error will stem from a collection of idiosyncratic factors rather than a single systemic factor that permeates the entire loan book. This is more likely the larger and more global the bank is. On the other hand, portfolios in smaller banks that operate in local markets and/or on centralised data systems are more likely to be exposed to pertinent systemic model errors. Regardless, the QSF framework is applicable to all banks that produce statistical estimates of its portfolio RWAs.

The first step to determining the appropriate MoC confidence level is a model risk appetite strategy that defines a target bank-wide confidence level. Estimates of γ can be taken as the historical proportion of the RWA in each portfolio on non-overlapping observation windows. MoCs can be sampled using the method described in Section **??**, where the hyperparameters ν and τ can be chosen based on the bank's own historical MoC estimates for its PD, LGD and EAD models. As regulatory MoCs are computed per risk parameter, an appropriate estimate for the portfolio MoC, β_i, is thus approximated via Equation (**??**). Should a more informed choice be desired, an empirical distribution of quantified portfolio MoCs provide information about the location and spread of these historically. If this information is not available, either the values used in this work or expert opinion from model owners and developers can set ν and τ. Preliminary analyses of the sensitivity of the QSF to these hyperparameters indicate that the QSF is not dependent on ν, but there is a linear dependence on τ. Specifically, higher values for τ will result in a more conservative portfolio quantile κ. This is expected when considering the strict Schur-convexity of the QSF under certain idealised conditions: all else fixed, a higher dispersion in MoCs will increase q.

What remains is how to estimate the error correlations between estimated portfolio RWAs. Under the ASRF framework developed here for estimated portfolio RWAs, it is argued that, by virtue of overall bank size, idiosyncratic model risks likely dominate RWA estimate errors. Model errors due to missing regressor values, modelling choices, loan policy changes, or changes in the composition of the portfolio through time are all examples of commonly known idiosyncratic model risks. The framework developed in this paper implies that large banks with many portfolios globally are justifiably in a position to set $\mathbf{P} = \mathbf{I}_N$. In reality, systemic model risks are present even in large banks, albeit to a much lesser degree. For example, ambiguity as to which portfolio a specific borrower belongs is not uncommon for large corporate portfolios. This uncertainty in the scope of RWA application is an example of a systemic source of model risk that results in dependent RWA biases. Should a bank wish to assume dominant idiosyncratic model errors, it must then be shown that variability in RWA estimates from identified sources of systemic model errors are negligible.

For banks wishing to err on the side prudence, accurately estimating each portfolio's systemic model risk factor loading ϕ_i presents a challenge. However, given the results in Figures **??** and **??**, one may conclude that accurate estimation is not necessary, and setting a fixed value for each portfolio based on qualitative arguments is sufficiently prudent. If this is not adequate for the bank's risk management, directly estimating ϕ_i may be bypassed by estimating ρ_{ij} instead, but this is again not immediately straightforward. Models that estimate RWA are not frequently developed, and each time a bank creates only one realisation of quantified model errors. The bank may choose to adopt the method described in Section **??** and extend it to their LGD and EAD models using beta and lognormal distributions, respectively, should such portfolios receive internal estimates. The bank's QSF is by definition insensitive to the volatility of the idiosyncratic model risk factor per portfolio, so this can be chosen qualitatively based on expert model opinion. Generating model error distributions then requires simulating from existing models and subtracting the true historical RWAs implied by historically observed defaults, economic losses/write-offs, and total exposures. Spearman rank correlations may then be estimated to avoid any bias due to portfolio size. This is neither the only nor the best approach, but it is a tractable and sound solution, with other possibilities left to future research.

Lastly, we briefly discuss the level at which risk parameters are calibrated and how this impacts the applicability of the methods described in this work. Regulation allows for the calibration of risk parameters at more granular levels than portfolio or segment. Specifically, these may be at rating grade level for PD or pool level for LGD and EAD, which then also extends to the level at which MoCs are quantified. Since these more granular levels are independently calibrated with their own MoCs, a similar diversification argument holds. Stressed grade or pool level estimates aggregate to a higher coverage level at the portfolio or segment level. As such, grade or pool level calibrations may require even lower quantiles than the estimated κ to ensure κ_{bank} is attained. In practice, instead of optimising for this quantile, it is advised to simply set the quantile for more granular MoC quantifications to a fixed value $\widetilde{\kappa} \leq \kappa$. Once the final parameters are estimated, the actual quantile attained at the portfolio or segment level can be checked to ensure at least κ is reached. This can be done via an application of the central limit theorem (CLT) to the portfolio or segment level long-run average PD, LGD or EAD. A number-weighted average of stressed grade/pool level risk parameters gives the stressed risk parameter at the portfolio level. Using the CLT, one can check the quantile the stressed value reach with respect to the implied distribution of the risk parameter at this level.

8 Conclusion

In this paper we introduced the QSF, a simple formula usable by credit risk managers to set an appropriate, non-trivial confidence level for portfolio MoCs. Specifically, the QSF links a bank-wide confidence level set by the bank's risk appetite framework to the confidence level required per portfolio's estimated RWA. The QSF itself is a ratio of volatilities: the total bank-wide RWA volatility with respect to the sum of the marginal RWA volatilities per individual portfolio. We investigated the characteristics of a bank's loan book that influence the required statistical coverage per portfolio's estimated RWA to satisfy the bank's model risk appetite. These are, namely, RWA error correlations, RWA concentration, and the total number of portfolios in the loan book. For a given model risk appetite, more uniform RWA concentrations and/or primarily idiosyncratic model errors result in a high degree of model risk diversification, allowing for a narrower MoC confidence level per portfolio. High RWA concentration in one or more portfolios and/or a dominant systemic model error weakens this diversification effect. This is expected, since the volatility of the total bank RWA will be mostly driven by either a single portfolio or the systemic model bias. Additionally, the more portfolios in the bank's loan book, the stronger the model risk diversification. The QSF was estimated on a synthetically generated loan book using real-world aggregate statistics published by supervisory authorities, and the implied portfolio quantiles estimated. In the discussion section we provide practical guidance on how to implement the QSF framework for any bank wishing to enhance its model risk management framework. It should be noted, however, that doing so might lead to questions by regulators, who often apply benchmark conservatism values to banks under their supervision. This notwithstanding, the work in this research shows that no benchmark is universally applicable. Loan book characteristics that influence a bank's QSF are all unique, as no two banks experience the exact same model risks.

Acknowledgments

The author would like to thank Wilco den Dunnen, Andrea Tiseni, and Fotios Dimitrakopoulos of ING for helpful discussions.

Data Availability Statement

The data used in this research is generated using online public sources. The EBA transparency dataset and ECB supervisory statistics are used solely to calibrate aggregate scale and concentration parameters; no bank-specific or confidential data are used. Data and code are available upon request.

Disclaimer

This research is not directly funded by any third party. The author reports no conflicts of interest, and all views and conclusions presented here are the views of the author alone. This manuscript benefited from productivity tools for editing and code generation; the author remains responsible for all content.

A Proofs

A.1 Proof of Theorem ??

Proof. All building blocks (Φ, ϕ, Φ^{-1}, algebra, square-roots, $R(\cdot), M(\cdot)$) are C^2 on $K = [\underline{p}, \overline{p}]$ with $\underline{p} > 0$. Hence $RWA \in C^2(K)$. Now, compute the first two derivatives of RWA with respect to PD. Let $z(PD) = \Phi^{-1}(PD)$, hence we have that

$$z' = \frac{1}{\phi(z)}, \quad z'' = \frac{z}{\phi(z)^2}.$$

Write $r = R(PD)$, $r' = R'(PD)$, $r'' = R''(PD)$, and

$$A(PD) = \frac{z}{\sqrt{1-r}} + G\sqrt{\frac{r}{1-r}}, \qquad G = \Phi^{-1}(0.999).$$

Then

$$A' = \frac{z'}{\sqrt{1-r}} + \frac{z\, r'}{2(1-r)^{3/2}} + \frac{G\, r'}{2(1-r)^{3/2}\sqrt{r}}, \qquad A'' \text{ exists and is continuous on } K.$$

For the maturity adjustment with $b(PD) = (a - c\ln PD)^2$ and $M(PD) = \frac{1+\alpha b}{1-1.5b}$,

$$b' = -\frac{2c(a - c\ln PD)}{PD}, \quad b'' = \frac{2c(a + c - c\ln PD)}{PD^2}$$

$$M' = \frac{(\alpha+1.5)\, b'}{(1-1.5b)^2}, \quad M'' = (\alpha+1.5)\Big[\frac{b''}{(1-1.5b)^2} + \frac{3(b')^2}{(1-1.5b)^3}\Big].$$

Define $K(PD) = LGD\,[\Phi(A) - PD]$. Then

$$K' = LGD\,\big[\phi(A)A' - 1\big], \qquad K'' = LGD\,\phi(A)\,\big[A'' - A(A')^2\big].$$

Finally, with $Y = 12.5\,EAD$,

$$RWA' = Y\{M'K + MK'\}, \qquad RWA'' = Y\{M''K + 2M'K' + MK''\}.$$

Taylor's theorem with a Lagrange remainder yields

$$RWA(PD) = RWA(PD_0) + RWA'(PD_0)(PD - PD_0) + \tfrac{1}{2}RWA''(\xi)(PD - PD_0)^2$$

for some ξ between PD and PD_0. Continuity on compact K implies $L_K := \sup_{p\in K}\|RWA''(p)\| < \infty$, giving a uniform quadratic error bound. ∎

A.2 Proof of Proposition ??

Proof. Let $t^\star := Z/\sum_{j=1}^{N} w_j$ and $z^\star := (t^\star, \ldots, t^\star)$. Then $\sum_i w_i z_i^\star = t^\star \sum_i w_i = Z$, so $z^\star$ is feasible, and its maximum coordinate equals $t^\star$. For any other feasible z, we have the elementary bound $\sum_i w_i z_i \le (\max_i z_i)\sum_i w_i$. Because feasibility requires $\sum_i w_i z_i \ge Z$, this bound implies $\max_i z_i \ge Z/\sum_i w_i = t^\star$. Hence no feasible vector can achieve a strictly smaller maximum coordinate than $z^\star$.

To see uniqueness, suppose z is feasible with $\max_i z_i = t^\star$ but some component satisfies $z_k < t^\star$. Then

$$\sum_i w_i z_i \;\le\; t^\star \sum_i w_i - (t^\star - z_k)\, w_k \;<\; t^\star \sum_i w_i \;=\; Z,$$

contradicting feasibility. Equivalently, if one component were below $t^\star$, some other would have to exceed $t^\star$ to meet the bank-wide constraint, raising the maximum above $t^\star$. Therefore all components must equal $t^\star$, which proves that the unique minimax-fair allocation is $z_i = t^\star$ for all i, i.e., $\kappa_i = \Phi(t^\star)$. ■

A.3 Proof of Theorem ??

Proof. By linearity,

$$\mu_{\text{bank}} := \mathbb{E}\left[R^\top \iota\right] = \sum_{i=1}^{N} \mu_i.$$

With $\mathbf{\Sigma} = \operatorname{diag}(\sigma)\, \mathbf{P} \operatorname{diag}(\sigma)$ and ι the vector of ones,

$$\sigma^2_{\text{bank}} := \mathbb{V}[R^\top \iota] = \iota^\top \mathbf{\Sigma}\, \iota.$$

Substitute $\sigma_i = \mu_i\, \beta_i / \Phi^{-1}(\kappa)$ from (??):

$$\mathbf{\Sigma} = \frac{1}{(\Phi^{-1}(\kappa))^2} \operatorname{diag}(\mu)\, \mathbf{B}^\top\, \mathbf{P}\, \mathbf{B} \operatorname{diag}(\mu),$$

whence

$$\sigma^2_{\text{bank}} = \frac{1}{(\Phi^{-1}(\kappa))^2}\, \mu^\top \mathbf{B}^\top\, \mathbf{P}\, \mathbf{B}\, \mu = \frac{1}{(\Phi^{-1}(\kappa))^2}\, \mu^\top\, \mathbf{\Psi}\, \mu.$$

Finally, by definition $\beta_{\text{bank}} = \sum_{i=1}^{N} \gamma_i \beta_i$. ■

A.4 Proof of Theorem ??

Proof. Let $m_i := \sigma_i/\mu_i$ and $m_{\text{bank}} := \sigma_{\text{bank}}/\mu_{\text{bank}}$ be coefficients of variation. Equating the aggregated bank MoC to the bank-wide MoC at their respective quantiles gives

$$\Phi^{-1}(\kappa) \sum_i \gamma_i m_i \;=\; \Phi^{-1}(\kappa_{\text{bank}})\, m_{\text{bank}},$$

Using Theorem ??, Equation (??), and $\gamma_i = \mu_i/\mu_{\text{bank}}$ yields

$$q = \frac{\Phi^{-1}(\kappa)}{\Phi^{-1}(\kappa_{\text{bank}})} = \frac{\sqrt{(\gamma \odot \beta)^\top \mathbf{P}\, (\gamma \odot \beta)}}{\iota^\top (\gamma \odot \beta)}.$$

The ℓ_2/ℓ_1 form is immediate by setting $a = \gamma \odot \beta$ and writing $a^\top \mathbf{P} a = \|\mathbf{P}^{1/2} a\|_2^2$. ■

A.5 Proof of Theorem ??

Proof. For any $a \geq 0$, $a^\top \mathbf{P} a \leq \sum_{i,j} a_i a_j = (\sum_i a_i)^2$, since all off-diagonal correlations $\rho_{ij} \leq 1$ and diagonal terms equal 1. Hence $q \leq 1$. Nonnegativity is immediate. The eigenvalue bounds follow from the Rayleigh quotient: $\lambda_{\min}\|a\|_2^2 \leq a^\top \mathbf{P} a \leq \lambda_{\max}\|a\|_2^2$, yielding the stated inequalities upon division by $\|a\|_1$. Extremal structures are direct checks. ■

A.6 Proof of Corollary ??

Proof. Fix $N \geq 2$ and let $\gamma = (\frac{1}{N}, \dots, \frac{1}{N}) \in \mathbb{R}^N$. For $\beta \in \mathbb{R}_+^N \setminus \{0\}$ we have that $q(\beta) = \|\beta\|_2 / \|\beta\|_1$. For any $s > 0$, let

$$\mathcal{S}_s := \{\beta \in \mathbb{R}_+^N : \ \|\beta\|_1 = s\}.$$

We will thus show that the restriction of q to $\mathcal{S}_s$ is strictly Schur-convex. In other words, if $\beta, \beta' \in \mathcal{S}_s$ and $\beta \succeq \beta'$ (in the sense of majorisation), then $q(\beta) \geq q(\beta')$, with equality if and only if β is a permutation of β'.

The function q is homogeneous of degree 0, i.e. $q(c\beta) = q(\beta)$ for all $c > 0$. Hence it suffices to work on the simplex

$$\Delta := \Big\{p \in \mathbb{R}_+^N : \ \sum_{i=1}^N p_i = 1\Big\},$$

by setting $p = \beta / \|\beta\|_1$. On Δ, $q(\beta) = \|p\|_2 = \sqrt{\sum_{i=1}^N p_i^2}$. Thus it is enough to show that $p \mapsto \sum_i p_i^2$ is strictly Schur-convex on Δ.

Consider $\phi(p) := \sum_{i=1}^N p_i^2$. This function is symmetric and continuously differentiable. The Schur–Ostrowski criterion states that a C^1 symmetric function f is Schur-convex if $(x_i - x_j)\big(\partial f / \partial x_i - \partial f / \partial x_j\big) \geq 0$ for all x and all i, j. Here

$$\frac{\partial \phi}{\partial p_i}(p) = 2p_i \quad \Longrightarrow \quad (p_i - p_j)\Big(\frac{\partial \phi}{\partial p_i} - \frac{\partial \phi}{\partial p_j}\Big) = 2(p_i - p_j)^2 \ \geq \ 0.$$

Hence ϕ is Schur-convex on $\mathbb{R}_+^N$, and in particular on Δ. Moreover, since $t \mapsto t^2$ is strictly convex, ϕ is *strictly* Schur-convex on Δ: if $p \succeq p'$ and p is not a permutation of p', then $\phi(p) > \phi(p')$.

Take $h(u) = \sqrt{u}$ strictly increasing on $[0, \infty)$. Therefore $h \circ \phi$ is (strictly) Schur-convex wherever ϕ is. In particular,

$$p \succeq p' \quad \Longrightarrow \quad \|p\|_2 = \sqrt{\phi(p)} \ \geq \ \sqrt{\phi(p')} = \|p'\|_2,$$

with strict inequality unless p is a permutation of p'. Hence, if $\beta, \beta' \in \mathcal{S}_s$, then $p = \beta / s$ and $p' = \beta' / s$ lie in Δ and satisfy $p \succeq p'$ whenever $\beta \succeq \beta'$. Thus

$$q(\beta) = \frac{\|\beta\|_2}{s} = \|p\|_2 \ \geq \ \|p'\|_2 = \frac{\|\beta'\|_2}{s} = q(\beta'),$$

with equality if and only if β and β' differ by a permutation. This proves the claim. Upper and lower bounds on q are straightforward to show ($\|\beta\|_2 \leq \|\beta\|_1$ for the upper bound and an application of Cauchy-Schwarz for the lower bound). ■

B Numerical Asymptotic Portfolio RWA Distribution

To describe the asymptotic distribution of a portfolio's RWA, a bootstrap is run on a portfolio containing $k = 1, \dots, 10000$ loans. For simplicity we consider the loss distribution of each individual loan to follow a lognormal, but the result does not change when using a Vasicek loss distribution. Each loan's loss is sampled via

$$l^{(k)} \sim \text{Lognormal}\left(\mu^{(k)}, \sigma^{(k)}\right),$$

where the mean and standard deviation are sampled uniformly,

$$\mu^{(k)} \sim U(-60, 10), \quad \sigma^{(k)} \sim U(0.01, 0.50).$$

This bounds the mean of the loan loss, $e^{\mu^{(k)}}$, below by 0 and above by a value with order of magnitude 10^8. The errors on this mean are sampled to reflect typical MoC values between 1-50%. The RWA for the portfolio is then taken to be the difference between the 99.9th percentile and the mean of this loss distribution.

The portfolio loss distribution is bootstrapped 1000 times and the RWA re-computed to produce a distribution for the RWA, given a lognormal loss distribution with varying means and variances for each loan. The resulting RWA distribution is provided in Figure **??** below. It is clear that the RWA indeed converges to a normal distribution.

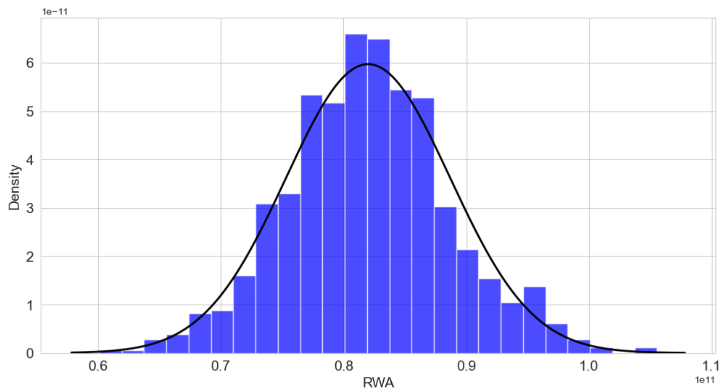

Figure 9: Bootstrapped RWA distribution given Monte Carlo simulated portfolio loss distributions.

C Numerical Diagnostics for Local Linearity of RWA in PD

This appendix quantifies and visualizes the local linearity of the IRB RWA function in PD across common asset classes. For each class we evaluate the curvature

$$L_K = \sup_{p \in K} \left| RWA''(p) \right|$$

over a realistic PD range K, and then compare the empirical first-order linearization error to the quadratic bound $\frac{1}{2} L_K \Delta^2$ implied by Taylor's theorem. Across all asset classes and PD ranges considered, the observed linearization error is dominated by the quadratic bound, and the ratio $|\text{error}|/\Delta^2$ stabilizes for small Δ. This empirically corroborates Theorem **??**: the remainder is indeed second-order small, validating local linearity as a working approximation for MoC aggregation.

Figure **??** shows the magnitude of the second derivative across PD on a logarithmic scale, highlighting curvature at very low PDs. The subsequent figures report, for each asset class, (i) the absolute linearization error versus the quadratic bound, and (ii) the second-order scaling ratio $|\text{error}|/\Delta^2$ around a representative PD_0 within the asset-class range. Parameter settings follow the main text (LGD/EAD typical values; corporate maturity 2.5 years).

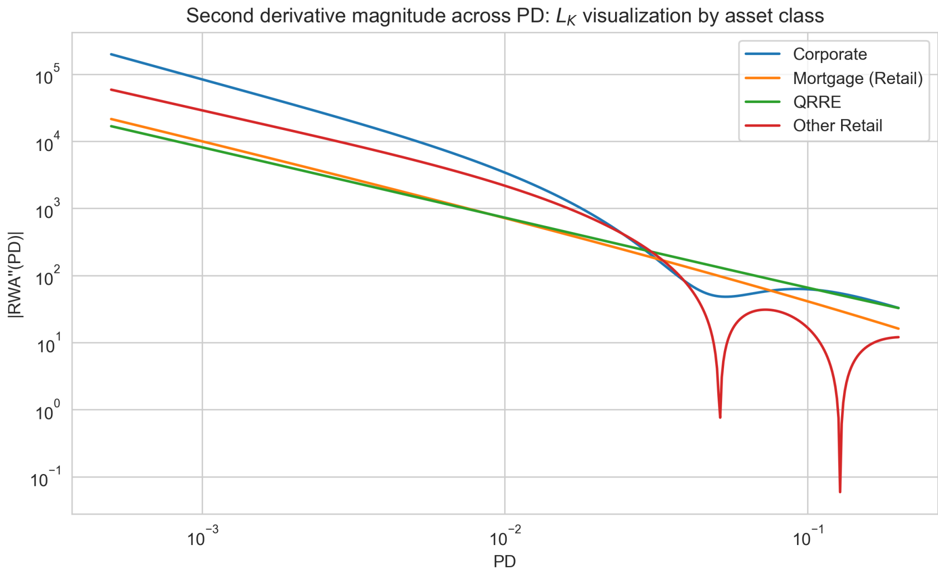

Figure 10: **Curvature diagnostic.** Magnitude of the second derivative $|RWA''(PD)|$ across PD. The supremum over each asset-class PD range is $L_K = \sup_{p \in K} |RWA''(p)|$.

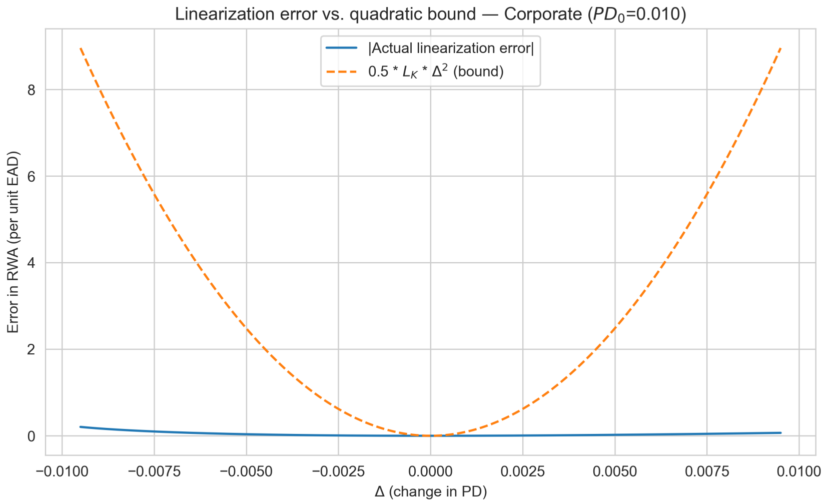

Figure 11: **Linearization error vs. quadratic bound** for Corporate at $PD_0 = 0.010$. The empirical linearization error (solid) is dominated by the quadratic bound $\frac{1}{2}L_K\Delta^2$ (dashed), confirming the second-order remainder.

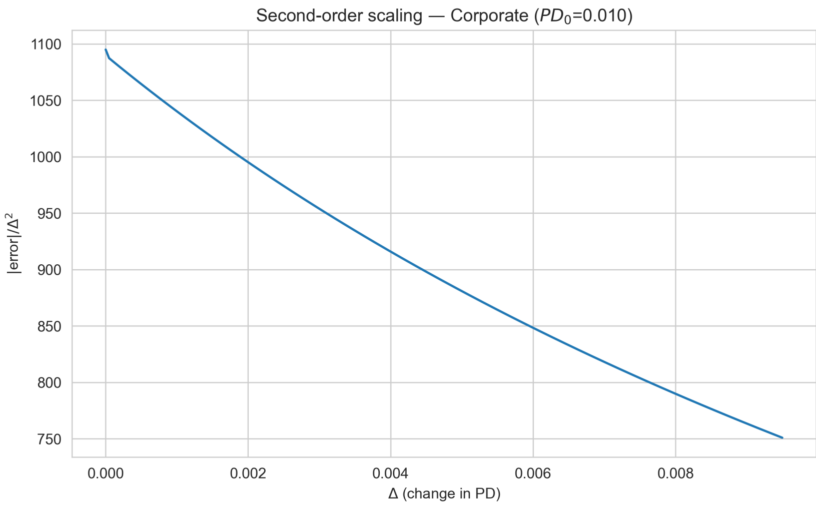

Figure 12: **Second-order scaling check** for Corporate at $PD_0 = 0.010$. The ratio $|\text{error}|/\Delta^2$ stabilizes for small Δ, consistent with a quadratic remainder.

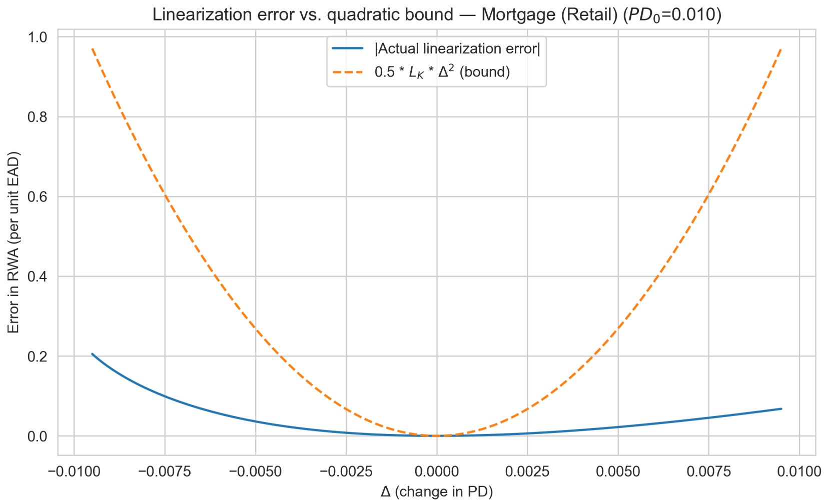

Figure 13: **Linearization error vs. quadratic bound** for Mortgage (Retail) at $PD_0 = 0.010$. The empirical linearization error (solid) is dominated by the quadratic bound $\frac{1}{2}L_K\Delta^2$ (dashed), confirming the second-order remainder.

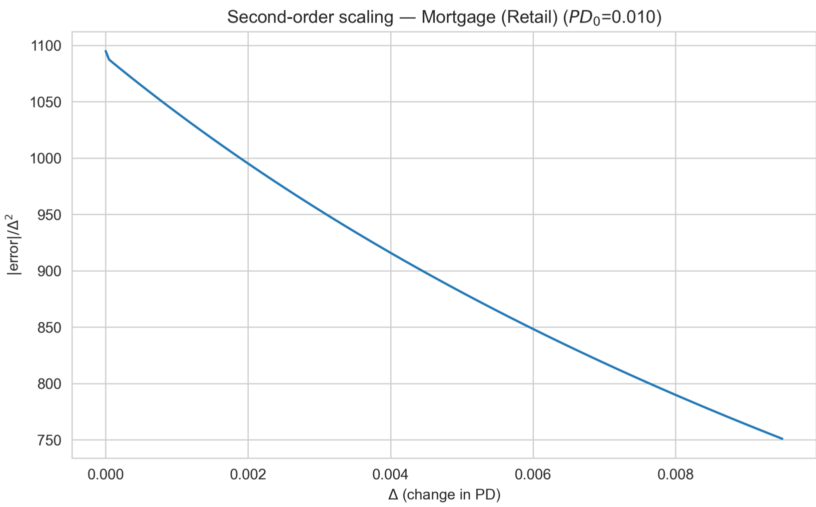

Figure 14: **Second-order scaling check** for Mortgage (Retail) at $PD_0 = 0.010$. The ratio $|\text{error}|/\Delta^2$ stabilizes for small Δ, consistent with a quadratic remainder.

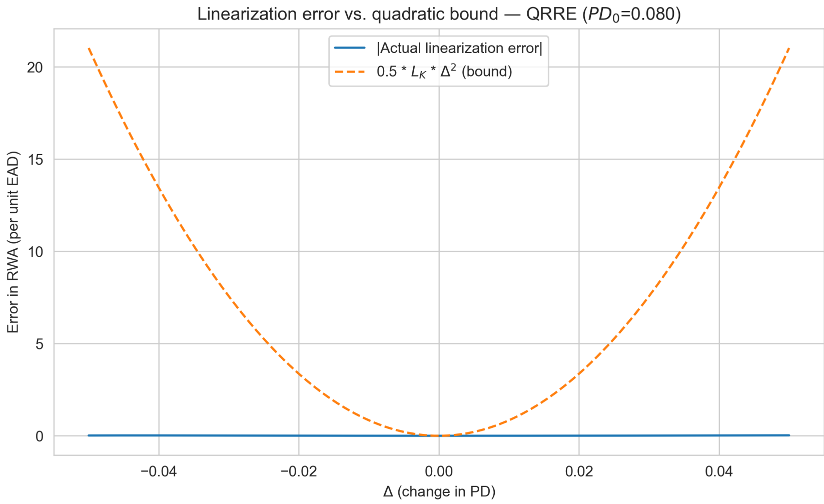

Figure 15: **Linearization error vs. quadratic bound** for QRRE at $PD_0 = 0.080$. The empirical linearization error (solid) is dominated by the quadratic bound $\frac{1}{2}L_K\Delta^2$ (dashed), confirming the second-order remainder.

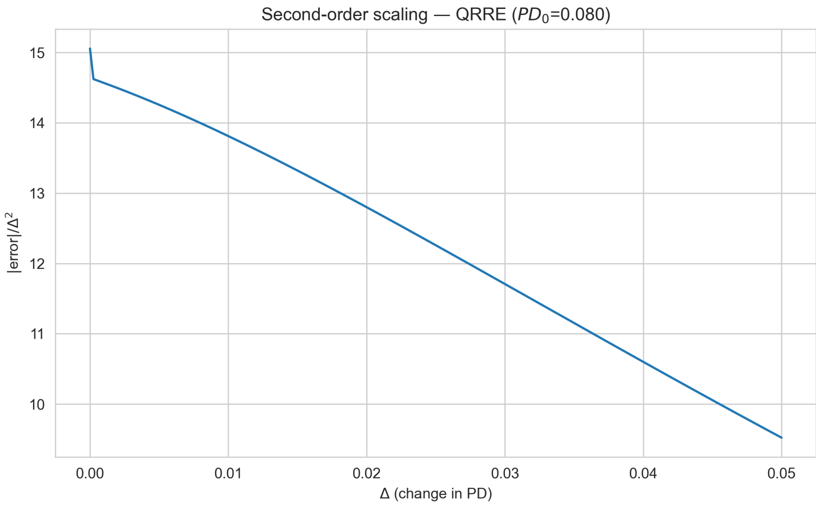

Figure 16: **Second-order scaling check** for QRRE at $PD_0 = 0.080$. The ratio $|\text{error}|/\Delta^2$ stabilizes for small Δ, consistent with a quadratic remainder.

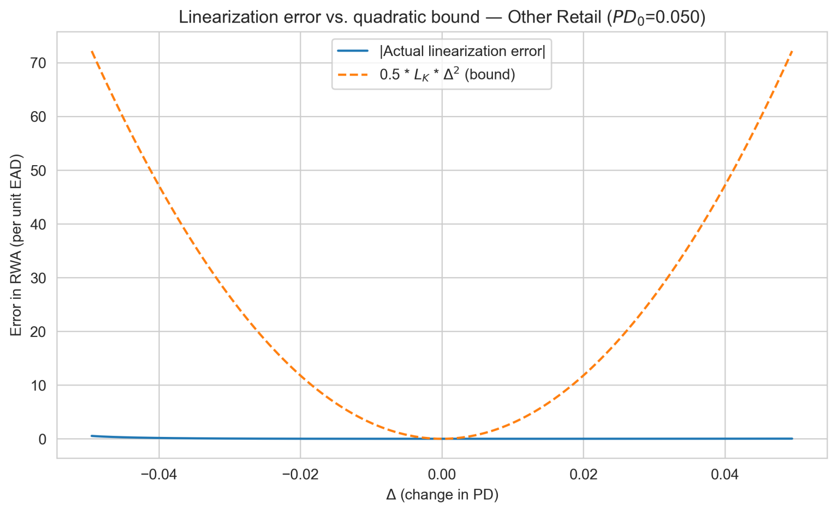

Figure 17: **Linearization error vs. quadratic bound** for Other Retail at $PD_0 = 0.050$. The empirical linearization error (solid) is dominated by the quadratic bound $\frac{1}{2}L_K\Delta^2$ (dashed), confirming the second-order remainder.

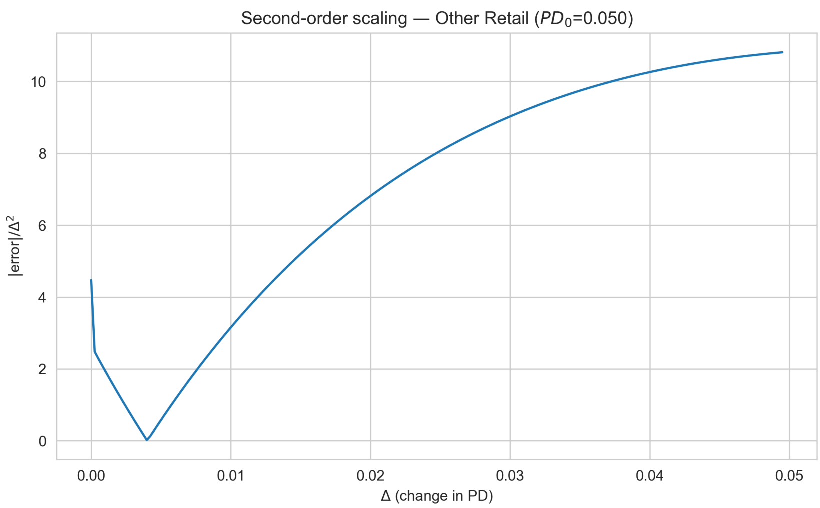

Figure 18: **Second-order scaling check** for Other Retail at $PD_0 = 0.050$. The ratio $|\text{error}|/\Delta^2$ stabilizes for small Δ, consistent with a quadratic remainder.

D Synthetic Loan Book Generation

To enable empirical analysis of the QSF without relying on proprietary bank-level IRB data, we construct a synthetic loan book dataset that replicates the statistical properties of real-world bank portfolios. To do this, the EBA transparency dataset and ECB supervisory statistics were used to calibrate aggregate scale and concentration parameters. The design ensures that the data are realistic enough for meaningful methodological testing, fully reproducible, and free of confidential supervisory information.

Loan book structure. We define 10 representative IRB asset-class portfolios:

Identifier	Portfolio Type
P1	Residential mortgages
P2	Retail unsecured
P3	SMEs
P4	Large corporates
P5	Commercial real estate
P6	Sovereigns
P7	Financial institutions
P8	Auto loans
P9	Project finance
P10	Specialty consumer finance

Each portfolio is observed annually for the period 2014–2024. Borrower counts are fixed per portfolio–year within a range reflecting sector size in EBA transparency datasets: large retail portfolios have tens of thousands of borrowers, while low-default wholesale or sovereign portfolios have fewer than 1,000. Specific details can be found in Table **??**.

Credit risk parameter calibration. Three risk parameters are assigned to each portfolio:

1. **Probability of Default (PD).** Baseline mean PDs range from 0.02% (sovereigns) to 2% (specialty consumer finance), calibrated to EBA Transparency Exercise aggregates. Yearly PDs are generated via:

$$PD_{i,t} = \min\left\{1,\, \max\left[0.0001,\, PD_i^{\text{true}} e^{\sqrt{1-\rho_i}\varepsilon_{i,t}+\sqrt{\rho_i}\lambda_{i,t}}\right]\right\}$$

   where:

   - PD_i^{true} = portfolio baseline PD;
   - ρ_i = asset correlation of portfolio i to the systemic macro factor;
   - $\lambda_{i,t} \sim \mathcal{N}(0, \sigma^2_{\text{macro}})$ is the systemic macro shock;
   - $\varepsilon_{i,t} \sim \mathcal{N}(0, 0.2 \times PD_i^{\text{true}})$ is idiosyncratic portfolio-specific noise.

This induces realistic cross-sectional correlation of default rates across portfolios. The true PD time series can be seen in Figure **??**.

2. **Loss Given Default (LGD).** Mean LGDs range from 20% (prime residential mortgages) to 60% (project finance), consistent with public EBA LGD data. Annual realised LGDs are drawn from a Beta distribution with parameters $(\alpha_{\mathrm{LGD}}, \beta_{\mathrm{LGD}})$ set to match the mean and a variance of 0.02. If no defaults occur in a given year–portfolio, the realised LGD is set to the mean.

3. **Exposure at Default (EAD).** Mean EAD per obligor varies from EUR 10,000 (specialty consumer finance) to EUR 2,000,000 (interbank exposures). EADs are drawn from a lognormal distribution with mean equal to the target average EAD and volatility parameter $\sigma = 0.5$ to introduce skew. Total EAD for each portfolio–year equals the average EAD multiplied by the number of borrowers.

Sampling defaults. Defaults are sampled from a binomial distribution for each portfolio–year, using a varying number of borrowers and the true PD:

$$D_{i,t} \sim \mathrm{Binomial}\left(N_{i,t}, PD_{i,t}\right).$$

An example of the observed default rate (ODR) and corresponding long-run average default rate (LRA DR) for the synthetic residential mortgages portfolio can be seen in Figure **??**.

RWA calculation. The true risk-weighted assets are computed using the Basel IRB formula corresponding to the portfolio asset class:

1. Asset correlation ρ follows the BCBS-prescribed PD-dependent formula for corporate exposures.

2. The Vasicek one-factor model is applied to compute unexpected loss at the 99.9th percentile.

3. RWA is calculated as:

$$RWA = 12.5 \times \text{Capital Requirement} \times \text{EAD}$$

   with LGD fixed at the portfolio mean and the total EAD per portfolio for the capital calculation.

Portfolio	Borrowers (thousands)	PD	LGD	EAD (EUR, thousands)	Asset correlation, ρ	Macro factor vol, σ_{macro}
Residential mortgages	20-50	0.20%	0.20	200	0.05	0.50
Retail unsecured	15-30	0.50%	0.40	50	0.10	0.20
SMEs	5-15	1.00%	0.45	80	0.15	0.30
Large corporates	2-8	0.30%	0.50	500	0.20	0.10
Commercial real estate	4-12	1.50%	0.55	300	0.25	0.25
Sovereigns	0.1-0.5	0.05%	0.20	1000	0.05	0.05
Financial institutions	0.5-2	0.08%	0.25	2000	0.05	0.15
Auto loans	8-20	1.00%	0.50	25	0.10	0.10
Project finance	0.2-3	0.50%	0.60	1500	0.15	0.10
Specialty consumer	1-8	2.00%	0.50	10	0.30	0.40

Table 4: Synthetic loan book characteristics based on published EBA and ECB statistics. PDs, LGDs and EADs are averages per borrower.

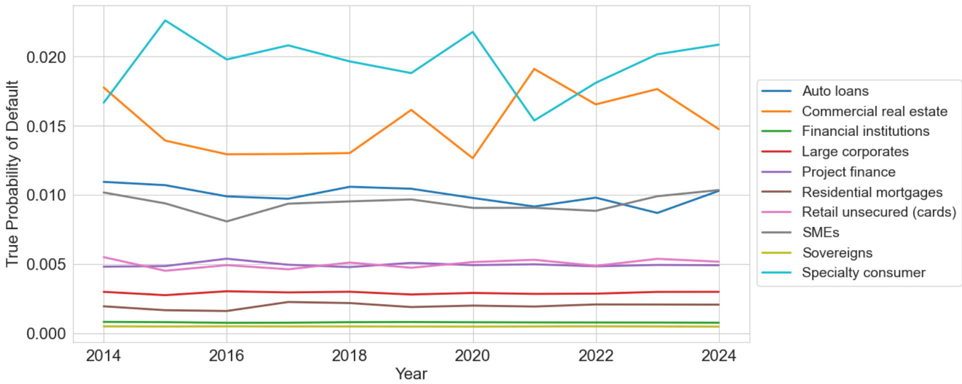

Figure 19: True, latent probability of default time series for the synthetic loan book.

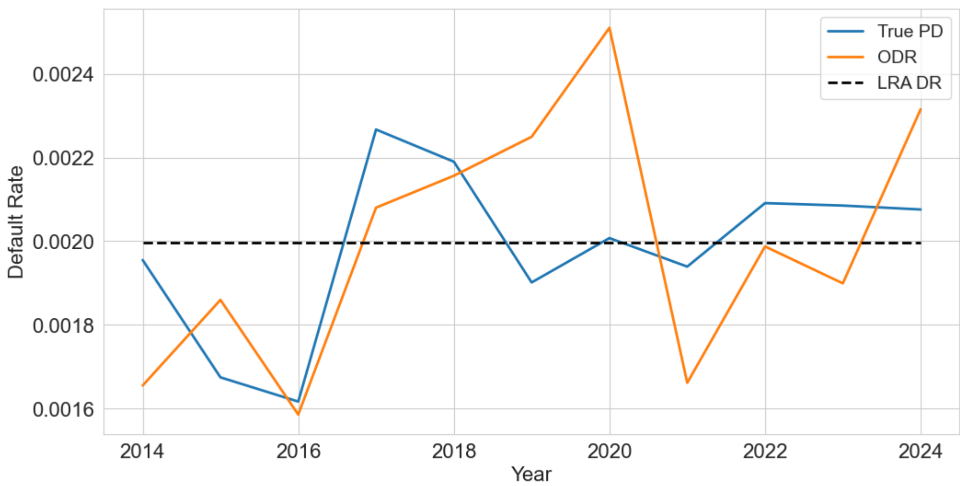

Figure 20: Generated ODR and LRA DR for the synthetic residential mortgages portfolio.