

A Stochastic Uncertainty-Propagation Framework for Forward-Looking Expected Credit Loss Under IFRS 9: An Empirical Illustration from an Emerging-Market Banking Portfolio

Charles Otieno Ndede^{1*}, David Chepkonga¹, Cheruiyot Kipkoech²

¹Department of Pure and Applied Mathematics, Jomo Kenyatta University of Agriculture and Technology, Nairobi, Kenya

²Department of Mathematics and Statistics, The Technical University of Kenya, Nairobi, Kenya

*Corresponding Author

DOI: <https://doi.org/10.51244/IJRSI.2026.1309000024>

Received: 11 September 2026; Accepted: 16 September 2026; Published: 30 September 2026

ABSTRACT

IFRS 9 requires expected credit loss (ECL) measurement to incorporate past events, current conditions and reasonable and supportable forecasts of future economic conditions. This study develops a stochastic uncertainty-propagation framework that maps uncertainty in credit-risk inputs and economic scenarios into a distribution of model-implied portfolio ECL. The framework retains the 12-month/lifetime distinction across IFRS 9 impairment stages, represents default timing through marginal default probabilities, and combines scenario-specific credit-risk quantities using probability weights and discounting. The empirical illustration uses an aggregated Kenyan banking portfolio comprising ABF Retail, Check Off, Non-Check Off and Retail Secured exposures, with total exposure of KES 286,404 million and a deterministic ECL benchmark of KES 15,245 million. The computational experiment follows a transparent design using explicit lognormal uncertainty factors, scenario probabilities, dependence assumptions, a fixed random seed and an independently reproducible simulation protocol. A 100,000-run Monte Carlo experiment produces a mean model-implied ECL of KES 15,193.40 million, standard deviation of KES 1,398.67 million and a central 90% interval of KES 13,025.71–17,587.48 million. The mean is 0.34% below the deterministic benchmark. Sensitivity analysis shows that alternative scenario weights and dependence assumptions change dispersion and tail outcomes while leaving the central estimate relatively stable. The results demonstrate uncertainty propagation rather than predictive superiority; account-level longitudinal data are required for calibration, back-testing and comparison with alternative IFRS 9 models. The analysis also documents numerical convergence and scenario-conditional behavior, providing a transparent basis for interpreting central and tail outcomes. The study therefore contributes a reproducible methodological architecture while identifying the empirical requirements for subsequent validation.

Keywords: IFRS 9; expected credit loss; stochastic modelling; Monte Carlo simulation; credit risk.

JEL Classification: G21; G28; C15; C63

INTRODUCTION

The global financial crisis of 2007–2009 exposed weaknesses in delayed recognition of credit losses under the incurred-loss approach. IFRS 9 replaced that logic with an expected-credit-loss framework that incorporates historical experience, current conditions and reasonable and supportable forecasts of future economic conditions (IASB, 2014). The change transformed impairment measurement from a predominantly backward-looking recognition problem into a forward-looking estimation problem in which the distribution of possible credit outcomes matters.

Under IFRS 9, the measurement horizon depends on changes in credit risk. Stage 1 exposures are generally subject to 12-month ECL, whereas Stage 2 and Stage 3 exposures require lifetime ECL. ECL is intended to be

unbiased and probability-weighted, reflect a range of possible outcomes and incorporate the time value of money. The familiar $PD \times LGD \times EAD$ decomposition is therefore useful but incomplete when the three components vary across time, economic states and stages.

The uncertainty is particularly relevant in emerging-market banking systems. Monetary policy, inflation, exchange rates, borrower income and sectoral conditions can change materially, while borrower-level default and recovery histories may be less readily available to researchers. For a portfolio with KES 286,404 million of credit exposure, a point estimate can summarize the central impairment assessment but does not reveal how the estimate changes when uncertainty in the underlying risk inputs is propagated through the ECL calculation.

Recent research has addressed macroeconomic adjustment of default risk, lifetime default-risk modelling, survival approaches, benchmark IFRS 9 ECL models and forward-looking credit-loss distributions (Schutte et al., 2020; Kvaal et al., 2024; Moodley et al., 2025; Jakubik & Teleu, 2025; Gross & Millischer, 2025; Botha & Verster, 2026). Elashmawy and Kallunki (2026) further provide empirical evidence that forward-looking IFRS 9 provisioning improves the predictive relevance of loan-loss provisions and reduces the gap between expected and subsequent realized credit losses. Their evidence establishes the predictive value of forward-looking provisioning but does not specify the stochastic mechanism by which economic uncertainty is transmitted into ECL.

This study addresses that methodological issue by formalizing ECL as a stochastic mapping from economic scenarios and uncertain credit-risk inputs to model-implied portfolio loss. The contribution is deliberately narrower than a claim of forecasting superiority. The study (i) formalizes uncertainty propagation through PD, LGD and EAD while preserving the 12-month/lifetime distinction; (ii) makes the dependence and scenario assumptions explicit and reproducible; and (iii) illustrates the framework using an aggregated Kenyan banking portfolio.

The research questions are:

- How can uncertainty in PD, LGD, EAD and economic scenarios be formally propagated through IFRS 9 ECL measurement?
- How does stochastic uncertainty propagation change the distribution of portfolio ECL relative to a deterministic benchmark?
- Which assumptions concerning scenario probabilities and dependence materially affect the ECL distribution?
- What additional data would be required to move from methodological uncertainty propagation to empirical predictive validation?

LITERATURE REVIEW AND RESEARCH GAP

IFRS 9 and forward-looking ECL

IFRS 9 requires ECL measurement to use reasonable and supportable information about past events, current conditions and forecasts of future economic conditions (IASB, 2014). The probability-weighted requirement means that alternative economic outcomes should be considered rather than relying exclusively on a single forecast. The post-implementation review of IFRS 9 also recognizes continuing practical and measurement challenges in impairment estimation (IASB, 2023).

Default risk, survival and the term structure of PD

PD is a principal determinant of expected loss, but lifetime ECL requires attention to the timing of default. Discrete-time survival approaches allow default risk to evolve over the life of an exposure and provide a natural basis for distinguishing conditional default probabilities from cumulative probabilities (Botha & Verster, 2026). This distinction is important because multiplying a cumulative PD by LGD and EAD at every future date can double-count default probability.

Macroeconomic uncertainty and scenario analysis

Macroeconomic conditions can affect repayment capacity, refinancing costs, collateral values and borrower liquidity. Recent work emphasizes the use of macroeconomic information to adjust default risk, particularly in developing-country settings (Moodley et al., 2025; Jakubik & Teleu, 2025). Forward-looking scenarios provide a practical way to represent uncertainty when a sufficiently long borrower-level panel is unavailable, but scenario weights and the mapping from macroeconomic states to credit-risk parameters must be disclosed.

LGD, collateral and EAD

Uncertainty in ECL is not confined to PD. LGD depends on recovery rates, collateral and workout timing, while EAD can change before default as borrowers draw down, repay or otherwise alter utilization. Consequently, a framework that propagates uncertainty through only PD can understate portfolio uncertainty. Schutte et al. (2020) provide a benchmark modular approach to IFRS 9 ECL, while subsequent work has continued to emphasize the importance of macroeconomic and loss-severity components.

Research gap

The literature provides important components of forward-looking ECL modelling, but the methodological gap addressed here is the explicit representation of the ECL calculation as an uncertainty-propagation system. A Monte Carlo procedure applied to a fixed deterministic ECL is not equivalent to a stochastic ECL model: the latter requires uncertain inputs, their dependence structure and scenario states to be specified before portfolio ECL is aggregated. The present study therefore focuses on the architecture and reproducibility of that propagation mechanism.

METHODOLOGY

Conceptual framework

The framework treats ECL as the terminal output of a sequence linking economic conditions, credit-risk parameters, default timing and discounted loss. Economic scenarios define the forward-looking environment; scenario states condition PD, LGD and EAD; default timing converts conditional PD into marginal default probability; and the resulting expected loss is discounted and probability weighted. The resulting workflow is illustrated in Figure 1.

Stochastic uncertainty-propagation framework for IFRS 9 ECL




Figure 1: Conceptual framework for stochastic uncertainty propagation in IFRS 9 ECL measurement.

Notation

Table 1: Notation and symbols used in the stochastic ECL framework.

Symbol	Meaning
I	Exposure or portfolio account
T	Future time period
S	Economic scenario
M	Monte Carlo simulation index, $m = 1, \dots, M$
T	Terminal maturity
$q_{i, s, t}$	Conditional probability of default in period t given survival to t-1
$S_{i, s, t-1}$	Probability of survival to the beginning of period t
$m_{i, s, t}$	Marginal probability of default in period t
$LGD_{i, s, t}$	Loss given default
$EAD_{i, s, t}$	Exposure at default
D_t	Discount factor
ω_s	Scenario probability
ECL_s	Scenario-specific portfolio ECL
$ECL^{\wedge}(m)$	Model-implied ECL in Monte Carlo run m
F_PD, F_LGD, F_EAD	Mean-one stochastic uncertainty factors

Default-Timing Formulation

To avoid ambiguity between cumulative and marginal PD, let $q_{i, s, t}$ denote the conditional probability that exposure i defaults in period t given survival to the end of period $t - 1$ under scenario s . The survival probability is:

$$S_{i, s, t-1} = P(\tau_i \geq t | s) \quad (1)$$

The marginal probability that default occurs in period t is therefore:

$$m_{i, s, t} = S_{i, s, t-1} q_{i, s, t} \quad (3)$$

If an available PD input is cumulative, it must first be transformed into the corresponding conditional or marginal default probability before it is multiplied by LGD and EAD over successive periods.

Discounted Scenario-Specific ECL

For exposure i under scenario s , the discounted expected loss over the relevant horizon is:

$$ECL_{i, s} = \sum_t m_{i, s, t} \times LGD_{i, s, t} \times EAD_{i, s, t} \times D_t \quad (4)$$

$$D_t = 1 / (1 + r_t)^t \quad (5)$$

Portfolio ECL under scenario s is:

$$ECL_s = \sum_i ECL_{i, s} \quad (6)$$

Stage-Specific Measurement Horizons

Stage 1 exposures are measured over a 12-month horizon, whereas Stage 2 and Stage 3 exposures are measured over the remaining expected lifetime. The empirical portfolio supplied for this study is aggregated by segment and stage and does not contain account-level migration histories. Consequently, the model does not estimate a dynamic stage-migration process. Instead, stage is retained as an observed impairment-state stratified, and stage-specific base ECL amounts are propagated stochastically.

$$ECL_{i, Stage1} = \Sigma_t = 1^{12} m_{i, s, t} \times LGD_{i, s, t} \times EAD_{i, s, t} \times D_t \quad (7)$$

$$ECL_{i, Stage2/3} = \Sigma_t = t + 1^T m_{i, s, t} \times LGD_{i, s, t} \times EAD_{i, s, t} \times D_t \quad (8)$$

Scenario Weighting

Let $s = 1, \dots, S$ denote alternative economic scenarios and ω_s their probability weights. The weights satisfy:

$$\Sigma_{s=1}^S \omega_s = 1, \omega_s \geq 0 \quad (9)$$

The probability-weighted portfolio ECL is:

$$ECL = \Sigma_{s=1}^S \omega_s ECL_s \quad (10)$$

Scenario weights therefore enter before the final portfolio ECL is interpreted. In the reconstruction, each simulation draw first selects an economic scenario and then applies the scenario-specific PD, LGD and EAD multipliers.

Stochastic Input Distributions and Dependence Structure

Stochastic uncertainty is represented by mean-one lognormal multiplicative factors applied to the retained segment-stage ECL components. These factors are an uncertainty-propagation device rather than independently estimated borrower-level PD, LGD or EAD distributions.

Table 2: Stochastic input distributions and parameterization used in the reproducible reconstruction.

Input	Distribution	Parameterization	Role
PD factor	Mean-one lognormal	CV = 0.057	Common credit-risk uncertainty
LGD factor	Mean-one lognormal	CV = 0.037	Recovery-severity uncertainty
EAD factor	Mean-one lognormal	CV = 0.0255	Exposure-at-default uncertainty
Stage 1 factor	Mean-one lognormal	CV = 0.030	Stage-specific residual uncertainty
Stage 2 factor	Mean-one lognormal	CV = 0.022	Stage-specific residual uncertainty
Stage 3 factor	Mean-one lognormal	CV = 0.015	Stage-specific residual uncertainty

For a mean-one lognormal factor F with coefficient of variation c , the implementation uses:

$$F \sim \text{Lognormal}(\mu, \sigma^2), \sigma^2 = \ln(1 + c^2), \mu = -\sigma^2/2 \quad (14)$$

The dependence among the three common credit-risk factors is represented through a Gaussian copula construction with the correlation matrix:

$$R = [[1.00, 0.15, 0.10], [0.15, 1.00, 0.05], [0.10, 0.05, 1.00]] \quad (15)$$

The matrix is positive definite and is transformed through a Cholesky factor before the latent normal draws are mapped into the lognormal factors. This structure allows PD, LGD and EAD uncertainty to co-move rather than assuming independence.

Economic Scenarios and Macroeconomic Mapping

Table 3: Economic scenarios and stochastic multipliers used in the Monte Carlo experiment.

Scenario	Probability	PD multiplier	LGD multiplier	EAD multiplier
Baseline	0.72	1.000	1.000	1.000
Upside	0.18	0.930	0.995	1.000
Downside	0.08	1.040	1.020	1.010

Scenario	Probability	PD multiplier	LGD multiplier	EAD multiplier
Severe	0.02	1.100	1.050	1.020

The scenario multipliers are explicit model outputs. They are directional rather than econometrically estimated coefficients. The baseline represents the central economic state; the upside scenario reduces credit-risk pressure; and the downside and severe scenarios increase PD, LGD and EAD. Because the available macroeconomic history contains only ten annual observations, the study does not estimate a causal regression linking GDP, inflation, the policy rate or exchange rate directly to realized portfolio losses.

The scenario environment is characterized using the Kenyan policy rate, inflation, real GDP growth and KES/USD exchange rate. These variables provide economic context for the scenario states but are not interpreted as independently identified causal determinants of ECL.

Monte Carlo Algorithm

The stochastic implementation proceeds as follows: The principal stochastic inputs are summarized in Tables 2 and 3.

1. Retain the observed segment-level ECL benchmark and stage balances.
2. Allocate each segment's deterministic ECL across Stage 1, Stage 2 and Stage 3 using explicit stage-intensity weights of 0.035, 0.180 and 0.650, respectively, normalized to the observed segment ECL.
3. Select an economic scenario according to the probabilities in Table 3.
4. Generate correlated latent normal shocks for PD, LGD and EAD and transform them into mean-one lognormal uncertainty factors.
5. Apply the scenario multipliers to the three common credit-risk factors.
6. Generate stage-specific lognormal factors and multiply each segment-stage base ECL by the PD, LGD, EAD and stage factors.
7. Aggregate all segment-stage components to obtain one model-implied portfolio ECL draw.
8. Repeat independently for $M = 100,000$ runs using random seed 20,260,904.
9. Summarize the resulting distribution by its mean, standard deviation, percentiles and scenario-conditional statistics.

Worked Numerical Example

Consider the ABF Retail Stage 1 component. Its observed Stage 1 balance is KES 10,474 million. Applying the stage-intensity rule gives a base Stage 1 ECL of KES 422.66 million after normalizing the three stage components to the segment's reported ECL of KES 1,690 million. In a particular simulation draw, suppose the selected scenario is Baseline and the generated PD, LGD, EAD and Stage 1 factors are denoted by FPD , $FLGD$, $FEAD$ and $FS1$. The simulated component is then:

$$ECL^m_{ABF,S1} = 422.66 \times FPD \times FLGD \times FEAD \times FS1 \quad (16)$$

The same operation is applied to every segment-stage component and the results are summed. The worked example demonstrates exactly how uncertainty enters the empirical reconstruction; it does not imply that the underlying borrower-level PD, LGD and EAD parameters have been independently estimated from account-level observations.

Alternative Representations of ECL

Two additional representations are retained for conceptual completeness. A capital-market formulation can derive a default intensity from an observable credit spread:

$$\lambda_t \approx s_t / LGD_t \quad (17)$$

$$S(t) = \exp(-\int_0^t \lambda(u) du) \quad (18)$$

A discounted-cash-flow representation is:

$$ECL_i = \sum_t [C_t^c - C_t^e] / (1 + EIR)^t \quad (19)$$

These formulations are not separately estimated or averaged with the stochastic PD–LGD–EAD reconstruction. They clarify the broader modelling landscape and identify directions for future empirical comparison.

Empirical Setting and Data

Kenyan Macroeconomic Context

Year	CBK policy rate (%)	Inflation (%)	GDP growth (%)	KES/USD
2015	10.19	6.6	5.0	98.2
2016	10.67	6.3	4.2	101.5
2017	10.00	8.0	3.8	103.4
2018	9.33	4.7	5.6	101.3
2019	8.92	5.2	5.1	102.0
2020	7.19	5.3	−0.3	106.5
2021	7.00	6.1	7.6	109.4
2022	7.67	7.6	4.9	117.9
2023	10.25	7.7	5.6	139.8
2024	12.50	5.0	5.0	140.0

Table 4: Kenyan macroeconomic indicators used to characterize the forward-looking environment, 2015–2024.

The policy rate declined from 10.19% in 2015 to 7.00% in 2021 before increasing to 12.50% in 2024. Inflation varied materially, GDP growth fell to −0.3% in 2020 before recovering to 7.6% in 2021, and the KES/USD exchange rate moved from 98.2 in 2015 to approximately 140.0 in 2024. These movements provide contextual information for scenario construction. The ten annual observations are not used to estimate dynamic macro-credit elasticities. The macroeconomic series are reported in Table 4.

Source: Central Bank of Kenya policy-rate information and World Bank Kenya macroeconomic data as supplied for the study.

Portfolio composition and impairment stage

Segment	Stage 1	Stage 2	Stage 3	Total exposure
ABF Retail	10,474	589	1,528	12,590
Check Off	193,214	10,192	7,770	211,176
Non-Check Off	37,210	1,788	3,145	42,143
Retail Secured	17,638	1,035	1,823	20,495
Total	258,536	13,604	17,266	286,404

Table 5: Credit exposure by lending segment and IFRS 9 impairment stage (KES million).

Check Off accounts for KES 211,176 million, approximately 73.8% of total exposure, and therefore dominates the portfolio by balance. The concentration is important because aggregate ECL is jointly determined by exposure size and loss intensity. The detailed segment-stage composition is presented in Table 5.

Deterministic ECL benchmark

Segment	Total exposure	Deterministic ECL	ECL rate
ABF Retail	12,590	1,690	13.42%

Segment	Total exposure	Deterministic ECL	ECL rate
Check Off	211,176	8,063	3.82%
Non-Check Off	42,143	4,155	9.85%
Retail Secured	20,495	1,337	6.52%
Total	286,404	15,245	5.32%

Table 6: Portfolio exposure and deterministic ECL benchmark (KES million, except ECL rate).

The deterministic benchmark is KES 15,245 million, equal to 5.32% of total exposure. The benchmark is treated as the retained portfolio-level comparison point rather than as a borrower-level ECL model re-estimated in this study. The benchmark composition is reported in Table 6.

RESULTS

Stage and segment heterogeneity

The portfolio contains KES 258,536 million in Stage 1, KES 13,604 million in Stage 2 and KES 17,266 million in Stage 3. Under the reconstructed stage-allocation rule, the stage-specific ECL rates vary substantially across segments, illustrating why a single portfolio-wide loss ratio can conceal heterogeneous loss intensity. Figures 2 and 3 visualize the segment-stage ECL structure.




Figure 2: Observed deterministic ECL contribution by lending segment and impairment stage.




Figure 3: Reconstructed stage- and segment-specific ECL rates.

Monte Carlo distribution of portfolio ECL

Statistic	Model-implied ECL (KES million)
Number of simulations	100,000
Mean	15,193.40
Standard deviation	1,398.67
Minimum	10,081.84
1st percentile	12,246.03
5th percentile	13,025.71
25th percentile	14,233.30
Median	15,120.56
75th percentile	16,069.38
95th percentile	17,587.48
99th percentile	18,891.43
Maximum	22,636.39
Monte Carlo standard error	4.42
Deterministic benchmark	15,245.00
Relative deviation	-0.34%

Table 7: Monte Carlo distribution of model-implied portfolio ECL.




Figure 4: Monte Carlo distribution of model-implied portfolio ECL.

The mean model-implied ECL is KES 15,193.40 million, which is KES 51.60 million below the deterministic benchmark, a relative deviation of approximately -0.34%. The central 90% interval extends from KES 13,025.71 million to KES 17,587.48 million. This interval is an uncertainty interval for model-implied ECL outcomes conditional on the reconstruction; it is not a confidence interval for an unknown parameter and is not presented as a validated forecast interval for realized losses. The complete simulation summary is reported in Table 7 and illustrated in Figure 4.

Cumulative distribution




Figure 5: Cumulative distribution of reconstructed model-implied ECL.

The cumulative distribution allows thresholds to be interpreted in terms of simulated probability. The interpretation remains conditional on the specified distributions, scenario weights and dependence structure. It does not represent the empirical frequency of future realized losses. The cumulative distribution is shown in Figure 5.

Independent Monte Carlo convergence

M	Mean ECL	SD	Monte Carlo SE
1,000	15,155.50	1,348.51	42.64
5,000	15,214.53	1,387.02	19.62
10,000	15,211.29	1,399.54	14.00
50,000	15,187.24	1,387.40	6.20
100,000	15,193.40	1,398.67	4.42

Table 8: Independent Monte Carlo convergence diagnostics (KES million).




Figure 6: Independent Monte Carlo convergence of the simulated mean ECL.

The independent convergence runs show that the mean stabilizes as the simulation size increases. The mean moves from KES 15,155.50 million at 1,000 draws to KES 15,193.40 million at 100,000 draws, while the Monte Carlo standard error declines from KES 42.64 million to KES 4.42 million. This provides a numerical convergence diagnostic for the reconstructed simulation. The numerical diagnostics are reported in Table 8 and shown in Figure 6.

Scenario-conditional results

Table 9: Scenario-conditional distribution of model-implied ECL (KES million).

Scenario	Runs	Mean ECL	SD	P05	Median	P95
Baseline	72,039	15,259.79	1,220.30	13,332.76	15,213.71	17,346.10
Downside	7,857	16,336.27	1,310.21	14,240.80	16,259.31	18,558.48
Severe	2,024	17,992.16	1,472.45	15,706.52	17,860.82	20,564.79
Upside	18,080	14,118.89	1,132.06	12,331.91	14,066.64	16,071.50

The scenario results display the intended directional structure. The upside scenario produces the lowest conditional mean, the baseline is intermediate, and downside and severe scenarios produce progressively larger conditional losses. The severe scenario has a low probability weight but materially increases the upper tail. This illustrates why scenario uncertainty and scenario probability should be reported separately from the final weighted ECL. The scenario-conditional statistics are reported in Table 9.

Sensitivity and robustness analysis

Specification	Mean ECL	SD	P05	Median	P95
Baseline weights	15,182.95	1,397.03	13,046.94	15,099.93	17,602.28
Higher downside	15,319.50	1,504.02	13,023.32	15,223.78	17,965.50
Lower downside	15,168.77	1,342.51	13,090.35	15,113.38	17,466.59
No cross-factor correlation	15,177.06	1,306.72	13,153.74	15,113.99	17,412.59
50% stronger correlation	15,180.46	1,440.23	12,966.32	15,090.66	17,656.94

Table 10: Sensitivity of model-implied ECL to scenario weights and dependence assumptions (KES million).




Figure 7: Sensitivity of mean ECL and dispersion to scenario and dependence assumptions.

The sensitivity analysis shows that increasing the weight assigned to adverse scenarios raises the central estimate and widens the distribution. Reducing cross-factor dependence slightly narrows the distribution, while stronger dependence increases dispersion. These results are consistent with the interpretation of the framework as an uncertainty-propagation system: the output distribution depends not only on central assumptions but also on the joint structure assigned to uncertain inputs. The sensitivity specifications are reported in Table 10 and illustrated in Figure 7.

Consistency of stochastic and deterministic ECL

$$\text{Relative deviation} = [(15,193.40 - 15,245.00) / 15,245.00] \times 100 = -0.34\%$$

$$\text{Monte Carlo variance} = (1,398.67)^2 \approx 1,956,267 \text{ (KES million)}^2$$

The close correspondence between the constructed stochastic mean and the deterministic benchmark is an internal consistency check. It does not establish predictive accuracy, model superiority or regulatory compliance. Those claims require independent realized outcomes and an empirically estimated PD, LGD and EAD system.

DISCUSSION

The results demonstrate the information added by treating ECL as a distribution of model-implied outcomes rather than only as a point estimate. The deterministic benchmark of KES 15,245 million summarizes the retained portfolio impairment assessment, while the stochastic framework shows how uncertainty in risk parameters, economic scenarios and dependence can widen the range around that central estimate.

The scenario-conditional results provide a direct interpretation of forward-looking uncertainty. The upside scenario lowers the conditional mean, while downside and severe scenarios raise it. Importantly, the severe scenario carries only a 2% weight in the reconstruction; its effect therefore operates primarily through the upper tail rather than dominating the central estimate.

The portfolio concentration also matters. Check Off represents approximately 73.8% of total exposure and contributes the largest absolute deterministic ECL despite having the lowest ECL rate. Conversely, ABF Retail has the highest segment ECL rate but a much smaller balance. Aggregate ECL monitoring should therefore consider exposure concentration and loss intensity jointly.

The framework should be distinguished from a fully estimated predictive credit-risk system. The empirical data are aggregated by segment and stage and do not contain borrower-level default dates, recovery histories, collateral values, EAD trajectories or sufficiently long longitudinal observations for estimating dynamic macroeconomic relationships. The reconstructed Monte Carlo experiment therefore demonstrates uncertainty propagation under explicit assumptions rather than predictive model efficacy.

Model Validation and Robustness Framework

Validation should occur at two levels. First, the simulation engine requires numerical validation through convergence diagnostics, fixed-seed reproducibility, positive-definite dependence matrices and sensitivity analysis. The present reconstruction addresses these requirements through independent convergence runs, an explicit random seed, a documented dependence matrix and scenario/dependence sensitivity tests.

Second, substantive credit-risk validation requires account-level or sufficiently granular longitudinal data. PD models should be evaluated using discrimination and calibration measures such as the Brier score, calibration error and, where appropriate, the area under the ROC curve. LGD models should be compared with realized recoveries and recovery timing, while EAD models should be evaluated against realized utilization or outstanding exposure at default.

For the stochastic ECL framework itself, future validation should examine empirical coverage of uncertainty intervals, stability of scenario weights, sensitivity to alternative distribution families, dependence structures and

macroeconomic mappings, and the contribution of each source of uncertainty to total ECL variance. Out-of-sample back-testing should compare the stochastic framework with the deterministic IFRS 9 benchmark and alternative PD, LGD and EAD specifications.

Managerial And Regulatory Implications

For provisioning, the framework allows management to consider a central ECL estimate together with the dispersion generated by model uncertainty. This can inform internal discussions around management overlays, capital planning and risk appetite, provided that the simulated distribution is clearly distinguished from regulatory capital requirements and from independently validated production-model outputs.

For tail-risk monitoring, the reconstructed 95th percentile of approximately KES 17.59 billion provides a quantitative reference within the specified simulation environment. It should not be treated as an externally validated stress-loss forecast. For portfolio concentration, the Check Off results demonstrate why absolute ECL contribution should be monitored alongside ECL rates. From a model-governance perspective, the principal benefit of the framework is transparency about which quantities are uncertain, how scenarios enter the calculation and how uncertainty is aggregated.

Limitations

The principal limitation is data granularity. The empirical material contains aggregated portfolio balances and ECL figures for four lending segments and three impairment stages. It does not contain the borrower-level observations required to estimate dynamic PD, LGD and EAD relationships or to validate simulated outcomes against realized credit events. The ten-year macroeconomic series are also too short to support robust causal estimation of dynamic credit-risk elasticities.

A second limitation is that the capital-market and discounted-cash-flow formulations are not separately estimated and compared with the scenario-weighted framework. They are included as theoretical alternatives. Future research can test whether combining accounting-risk, market-implied and cash-flow information improves forecast performance.

CONCLUSION

This study develops a stochastic uncertainty-propagation framework for forward-looking ECL measurement under IFRS 9 and provides an empirical illustration using an aggregated Kenyan banking portfolio. The framework represents ECL as a scenario-weighted, discounted function of default timing, PD, LGD and EAD while retaining the distinction between 12-month and lifetime measurement horizons.

The portfolio contains KES 286,404 million of exposure and a deterministic ECL benchmark of KES 15,245 million. Under the explicitly reconstructed stochastic specification, 100,000 Monte Carlo simulations produce a mean model-implied ECL of KES 15,193.40 million, standard deviation of KES 1,398.67 million and a central 90% uncertainty interval of KES 13,025.71–17,587.48 million. The mean is approximately 0.34% below the deterministic benchmark. Independent convergence diagnostics show stabilization of the simulated mean and a declining Monte Carlo standard error.

Sensitivity analysis demonstrates that scenario probabilities and dependence assumptions affect both the dispersion and tail of the ECL distribution. The findings therefore support the central methodological proposition that uncertainty should be propagated through the ECL calculation rather than suppressed into a single point estimate. They do not establish that the stochastic framework is more accurate than competing approaches. Such a claim requires account-level longitudinal data, empirical calibration, out-of-sample testing and comparison with realized defaults, recoveries and EAD.

The principal contribution is consequently methodological: the study makes the uncertainty structure of forward-looking ECL explicit, auditable and reproducible while identifying the data requirements for future predictive validation.

Statements And Declarations

Competing interests

The authors declare that they have no relevant financial or non-financial competing interests.

Funding

No external funding was reported for this study.

Data availability

The underlying banking portfolio data are aggregated and subject to confidentiality restrictions. The portfolio-level values reported in Tables 5 and 6 and the macroeconomic series used to characterize the study context are provided in the manuscript. Account-level data are not publicly available.

Code availability

A reproducible Python implementation of the stochastic reconstruction, together with the machine-readable results workbook, is available on request.

Use of generative AI

An AI-assisted tool was used during manuscript preparation for document restructuring, and language refinement. The authors retain responsibility for the mathematical formulation, assumptions, numerical implementation, results, interpretation and final manuscript.

Ethics approval

Not applicable; the study uses aggregated portfolio-level information and does not involve human participants or identifiable personal data.

REFERENCES

1. Botha, A., & Verster, T. (2026). Approaches for modelling the term-structure of default risk under IFRS 9: A tutorial using discrete-time survival analysis. *International Journal of Data Science and Analytics*, 22, 67.
2. Elashmawy, S., & Kallunki, J.-P. (2026). Does forward-looking accounting improve bank asset quality? *Accounting in Europe*, 23(1), 57–86. <https://doi.org/10.1080/17449480.2025.2470719>
3. Gross, M., & Millischer, L. (2025). Credit loss in translation: Informing bank provisions and capital buffer requirements with forward-looking credit loss distributions. *IMF Working Paper* 2025/228.
4. International Accounting Standards Board. (2014). *IFRS 9 Financial Instruments*. IFRS Foundation.
5. International Accounting Standards Board. (2023). *Request for Information: Post-implementation Review of IFRS 9 Financial Instruments—Impairment*. IFRS Foundation.
6. Jakubik, P., & Teleu, S. (2025). Improving credit risk assessment in uncertain times: Insights from IFRS 9. *Risks*, 13(2), 38.
7. Kvaal, E., Löw, E., Novotny-Farkas, Z., Panaretou, A., Renders, A., & Sampers, P. (2024). Classification and measurement under IFRS 9: A commentary and suggestions for future research. *Accounting in Europe*, 21(2), 154–175.
8. Moodley, S., Verster, T., & Raubenheimer, H. (2025). A pragmatic macroeconomic default risk adjustment in developing countries. *South African Journal of Economic and Management Sciences*, 28(1), 1–13.
9. Schutte, W. D., Verster, T., Doody, D., Raubenheimer, H., & Coetzee, P. J. (2020). A proposed benchmark model using a modularised approach to calculate IFRS 9 expected credit loss. *Cogent Economics & Finance*, 8(1), 1735681.