Operations research in credit risk management: models, optimisation techniques and practical applications

A comparative critical review of fourteen studies

[1]Abhiraj Rastogi, [2]Abhay Kumar, [3]Manthan Pareek, [4]Pragyan Singh Rathore, [5]Arihant Baid
[1–5]B.Sc. Finance
[1–5]NMIMS (Deemed-to-be University), Bengaluru, Karnataka
[1]abhiraj.rastogi335@nmims.in, [2]abhay.kumar318@nmims.in, [3]manthan.pareek310@nmims.in,
[4]pragyan.rathore559@nmims.in, [5]arhiant.baid358@nmims.in

***Abstract*—Credit-risk decisions combine uncertainty, scarce resources and competing financial objectives, which makes them natural candidates for Operations Research (OR). This structured narrative review consolidates five group reports, comprising fifteen article reviews of fourteen distinct studies published between 2006 and 2026. The studies cover portfolio-loss simulation, dependence and tail risk, robust and dynamic investment, loan pricing, large-scale loan selection, collateral allocation and one related bank-service application. They are classified by decision problem, compared on objectives, constraints and solution methods, and linked to the B.Sc. Finance OR syllabus. Three tensions emerge: realistic dependence against simplified portfolio assumptions, exact solutions against computational feasibility, and model complexity against empirical validation. The findings favour matching the method to the decision. All numerical results are attributed to the original studies; no independent replication was performed.**

***Index Terms*—credit risk; operations research; conditional value-at-risk; importance sampling; stochastic programming; robust optimisation; semi-Markov models; mixed-integer programming**

I. Introduction

A. Background and motivation

Credit risk is the possibility of financial loss when a borrower, bond issuer or counterparty fails to meet a contractual obligation. It is central to lending, loan pricing, bond investment, collateral management and portfolio decisions, which involve scarce capital, competing return and loss objectives, and uncertain future conditions. OR supplies methods for formulating such trade-offs: linear and mixed-integer programming for constrained allocation, quadratic programming for mean–variance models, simulation and stochastic programming for uncertain outcomes, robust optimisation for model misspecification, dynamic programming for sequential decisions and queueing theory for bank operations. Markowitz [1] formalised mean–variance portfolio selection, but variance alone describes credit portfolios poorly because losses are asymmetric, defaults cluster and ratings change, which motivates credit-specific risk measures.

The topic is interesting because one toolkit addresses problems that look unrelated: pricing a loan, hedging a bond, selecting loans for securitisation, allocating collateral and staffing a bank counter are all constrained decisions, and the technique follows their size, uncertainty and timing. Small modelling errors also compound into large losses, and regulators increasingly expect risk-adjusted decisions.

B. Scope and objectives

This paper consolidates five reports prepared by Group 1 for the Operations Research course, covering 15 article reviews. One study, Pasricha et al. [2], was reviewed by two members, leaving 14 distinct studies. The objectives are: (1) to classify the studies by the credit-risk problem each addresses; (2) to explain the OR technique each uses and why it suits the problem; (3) to compare the techniques on uncertainty handling, scalability and assumptions; and (4) to identify benefits, drawbacks and gaps, and what a stronger study would add. A supporting aim is to connect the techniques to course topics (Section III-F). Section II reviews the

studies by problem type, Section III analyses and compares the techniques, and Section IV presents conclusions, drawbacks and research gaps.

II. Literature Review

The fourteen studies fall into seven themes built around the decision being optimised. Each subsection states the model used, its contribution and the assumptions it requires; Table I in Section III-B compares the technique families.

A. Integrating market and credit risk through scenarios

Jobst et al. [3] study fixed-income asset–liability management under uncertain interest rates, spreads, default and rating migration, feeding simulated scenarios into multistage stochastic programmes with recourse and testing CVaR-based index tracking in-sample and out-of-sample; results depend on the scenario generator and risk constraints. Masala et al. [4] generate joint credit-loss scenarios from uncertain default times, recoveries and rating transitions using a Student-t copula, which captures joint adverse outcomes that linear correlation misses. Minimising conditional value-at-risk (CVaR) is then a linear programme when losses are affine in portfolio weights, although joint-tail dependence is hard to estimate from sparse crisis data.

B. Rare-event simulation of portfolio losses

Bassamboo et al. [5] study extreme portfolio-loss probabilities under dependence allowing simultaneous defaults. Their asymptotics show that linear correlation cannot determine tail behaviour, and their importance sampling makes rare losses frequent in simulation before correcting the distribution; the guarantees need not hold for small or heterogeneous portfolios. Glasserman et al. [6] develop importance sampling under a multifactor Gaussian-copula model, with sampling shifts chosen by combinatorial and convex optimisation. Tests use 1,000 obligors, 100 obligor types and 21 or 22 factors, so the efficiency shown holds only under the tested configurations.

C. Uncertainty about the model itself

Bo and Capponi [7] develop dynamic robust strategies for defaultable bonds under possible credit-model misspecification. The max–min formulation penalises deviation from the reference model through expected log-likelihood ratios, and the value functions satisfy coupled Hamilton–Jacobi–Bellman equations. Chi et al. [8] combine instance-based credit-risk assessment with a relative-entropy robust mean–variance formulation for peer-to-peer lending; the tractable problem is quadratic, not linear. On Prosper data the robust model earns an average test-subset return of 6.62% against 5.53% for its non-robust benchmark (20 test subsets), an experimental result and not a guaranteed advantage. Negligible loan correlation is assumed, a weakness in a downturn.

D. Dynamic credit investment

Pasricha et al. [2] model rating migration with semi-Markov processes that retain time spent in the current state. Simulated rating paths generate bond-return scenarios for a min–max absolute-deviation linear programme with budget, non-negativity and minimum-return constraints; correlated migrations remain an extension. Bo and Capponi [9] ask how much of a defaultable bond an investor should hold when borrowing costs more than lending and both parties may default. Dynamic programming gives a feedback rule in wealth and remaining time; unequal rates make the value function non-differentiable, so the problem is split into two smooth pieces. Calibration to a real portfolio would be demanding (one risky asset, constant rates, CRRA utility).

E. Large-scale selection and approximation

Sirignano et al. [10] study loan selection among thousands of loans, an intractable nonlinear integer programme. Their asymptotically optimal approximation exploits the distribution of loan characteristics, converges as portfolio size grows, and shows large computational advantages on loan data. Germain and Vrins [11] address loan selection for collateralised loan obligations, an NP-hard mixed-integer nonlinear problem, through a large-pool approximation, linearisation and clustering; the abstract reports better objective values and shorter computation than a derivative-free baseline on tested cases. Goblet et al. [12] allocate limited eligible collateral while constraining under-allocation and the number of adjustment movements, using an LP

relaxation with an L1 penalty followed by MILP on a reduced search space; tests on 20 real-world collateral givers show a useful trade-off among quality, runtime and movement count.

F. Pricing and multi-objective portfolio selection

Chun and Lejeune [13] treat the offered interest rate as an endogenous decision shaped by borrower acceptance, default risk and the loan's marginal contribution to portfolio VaR/CVaR. The stochastic mixed-integer nonlinear framework has convex continuous relaxations for specified acceptance functions, improving tractability without guaranteeing global optimality for non-convex multiloan problems. Wang et al. [14] propose a bivariate default-intensity-based risk measure and use a multi-objective genetic algorithm to trade a default-risk proxy against return. The example uses eight single-name CDS contracts (Markit iTraxx Series 31) observed weekly from January 2017 to December 2018, which limits generalisability, particularly to stressed markets.

G. Operational efficiency in banking

Jhala and Bhathawala [15] analyse a bank cash-counter configuration with M/M/1 and two-server M/M/s queues. For a four-week observational setting they calculate ₹1,575 per hour for separate withdrawal and deposit counters and ₹1,307 for pooled counters, a modelled difference of ₹268 per hour (17.0%). This is a service-operations example and not a default-risk study, and the totals rest on assumed waiting costs.

III. Analysis and Methodology

A. Method of analysis

This study is a structured narrative review, not a systematic review or meta-analysis. Five student reports containing three article assessments each, with the semi-Markov article counted twice, yield fourteen distinct publications. Citation metadata were checked against journal, publisher and university records where available, and selected technical and numerical descriptions were tested against accessible article text. No optimisation model, backtest or dataset analysis was rerun, so outcomes remain author-reported, and student-summary figures that could not be confirmed against primary publications are omitted. Each study was coded against three questions: what is uncertain (future scenarios, the credit model, rating migrations or default dependence); whether the decision is made once or sequentially; and how large the problem is and whether an exact solution is reachable. For each article we also record the decision variables, objective, constraints, uncertainty representation, solution method, evidence type and major limitation, which prevents misleading comparisons between, for example, a tail-probability estimator and a loan-selection solver.

B. Comparison of OR techniques

TABLE I. OR technique families across the studies

Technique family	Studies	Uncertainty handling	Main strength	Main limitation
Linear and mixed-integer programming	[12], [2], [4]	Deterministic inputs or simulated scenarios	Structured allocation under explicit constraints; mature solvers	Integer variables slow large instances; inputs must be estimated
Stochastic programming and simulation	[3], [4], [13]	Explicit scenarios or simulated defaults	Joint treatment of risk sources; CVaR tail control	Quality depends on scenario generator and estimates
Rare-event simulation	[5], [6]	Importance sampling of extreme losses	Accurate tail estimates with far fewer trials	Asymptotic results; costly with many factors
Robust optimisation	[7], [8]	Ambiguity sets or likelihood-ratio penalties	Protects against misspecification	Heavy mathematics; assumptions on the uncertainty set
Dynamic programming, control and Markov models	[9], [2]	State-dependent policy; rating migration as a stochastic process	Decisions adapt to new information	Simplifying assumptions: one asset, constant rates, no cross-bond correlation
Approximation and metaheuristics	[10], [11], [14]	Loan-type distributions, clusters or population search	Scales to thousands of loans; works without convexity	Guarantees differ: asymptotic, numerical or heuristic
Queueing theory	[15]	Poisson arrivals, exponential service	Transparent staffing cost trade-off	Strong distributional assumptions

C. Representative formulations

These formulations help interpret the studies and do not reproduce every model's full constraint set. First, the scenario approximation to CVaR minimisation [16] for S equally weighted scenarios and confidence level α is

$$\min \zeta + [1/((1-\alpha)S)]\,\Sigma_s\, u_s,\ \text{s.t. } u_s \geq L_s(x) - \zeta,\ u_s \geq 0,\ x \in X \tag{1}$$

where ζ is an auxiliary threshold, $L_s(x)$ is the loss in scenario s and X is the feasible allocation set. The programme is linear only if scenario losses and constraints are linear in x. Figure 1 illustrates the distinction between VaR and CVaR.

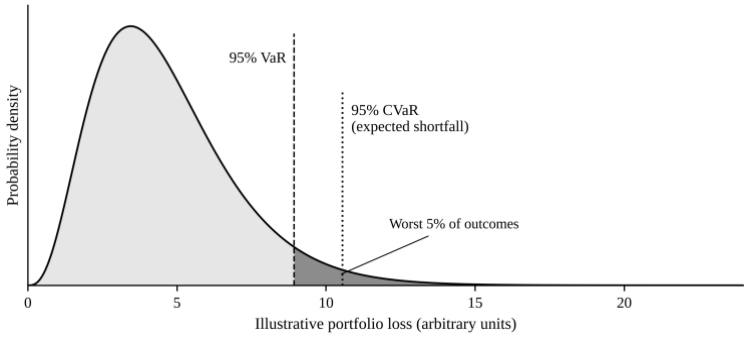


Fig. 1. The distinction between Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR) at the 95% confidence level. Illustrative loss distribution (not empirical data), based on the definitions in [16].

Second, the max–min structure of robust optimisation [7], where π is the investment strategy, Q ranges over plausible models and U is the utility of terminal wealth X, is

$$\max_{\pi \in \Pi} \min_{Q \in \mathcal{Q}} E_Q[U(X_T(\pi))]. \tag{2}$$

The ambiguity set $\mathcal{Q}$ and its penalties determine how conservative the decision becomes; this is a generic robustness template, not the full HJB system. Third, the binary selection problem behind Sirignano et al. [10] and Germain and Vrins [11], with $y_i = 1$ if loan i is chosen, risk measure ρ, returns r_i and required return R, selects K of N loans:

$$\min \rho(y)\ \text{ s.t. }\ \Sigma_i\, r_i\, y_i \geq R,\ \Sigma_i\, y_i = K,\ y_i \in \{0, 1\}. \tag{3}$$

When ρ is nonlinear this is a nonlinear integer programme whose candidate portfolios grow combinatorially with N, which is why the studies in Section II-E aggregate, relax or cluster instead of solving it directly. Fourth, the semi-Markov min–max LP of Pasricha et al. [2] for budget C and required return rate d is

$$\min z\ \text{ s.t. }\ -z \leq \Sigma_j\, (r_{jt} - \bar{r_j}) x_j \leq z\ \forall t,\ \ \Sigma_j\, \bar{r_j}\, x_j \geq dC,\ \ \Sigma_j\, x_j = C,\ \ x_j \geq 0. \tag{4}$$

The added variable z linearises the maximum absolute deviation and links the study to the course's LPP, simplex and shadow-price topics. The course's LP assumptions (proportionality, additivity, divisibility, known coefficients) give useful baselines; loan inclusion may need binary variables, dependence-aware risk is generally nonlinear, and unknown parameters add uncertainty. Dependence among losses does not by itself invalidate LP when scenario losses enter as given linear coefficients.

D.Cross-study findings

The studies use three concepts of uncertainty. Scenario models [3], [4], [13] treat a probabilistic model as given and optimise over simulated outcomes. Robust methods [7], [8] allow plausible alternative distributions or misspecified dynamics. Dynamic methods [9], [3] choose actions conditional on future states; Pasricha et al. [2] use a time-dependent rating model to generate scenarios, but their allocation is not a multistage policy.

Tractability competes with exactness. Sirignano et al. [10] prove asymptotic convergence under their model conditions, Germain and Vrins [11] report numerical gains from approximation, Goblet et al. [12] combine LP with a reduced MILP, and Wang et al. [14] explore a genetic algorithm. Asymptotic convergence, finite-instance bounds and heuristic performance are different types of evidence, and a timely approximate decision may beat a costly exact search if its accuracy is evaluated.

Dependence assumptions vary widely. Masala et al. [4] use a Student-t copula, Bassamboo et al. [5] study extremal dependence and Glasserman et al. [6] study factor-dependent defaults under a Gaussian copula, whereas Chi et al. [8] assume negligible P2P-loan correlations and Pasricha et al. [2] identify joint rating migration as an extension. External validation is also uneven: Glasserman et al. use 1,000 obligors with structured factor loadings, Goblet et al. test 20 real collateral givers, Jhala and Bhathawala [15] observe one bank over four weeks, and Wang et al. use eight CDS contracts. The evidence is variously simulated, empirical or based on assumed costs, and sample size alone neither establishes nor invalidates it.

The studies can be read as a credit-decision cycle: loan pricing [13], loss estimation [4], [5], [6], portfolio choice [3], [8], [2], [10], [11], [14], dynamic investment [9], [7], collateral management [12], and a related branch-service application [15]. This integration is our conceptual synthesis, not an architecture tested jointly by the cited authors, and cross-paper superiority cannot be inferred without re-estimation on common data.

E. Technique selection guide

The comparison supports a practical rule. Small, well-specified allocations with linear constraints suit linear or mixed-integer programming; uncertainty described by plausible scenarios with tail loss to control suits stochastic programming with CVaR; doubt about the model itself suits robust optimisation; sequential decisions with a changing state suit dynamic programming or stochastic control; very small loss probabilities call for importance sampling; thousands of binary decisions call for aggregation, relaxation or clustering before the integer step; problems with no closed-form objective or a non-convex feasible set suit metaheuristics, accepting near-optimality; and congestion and service-cost trade-offs suit queueing models. Implementation begins with formulation and data validation: is the task estimation or optimisation, can decisions be revised, do dependence and distribution error matter, and can the model be solved in time?

F. Link to the techniques taught in the course

The course covers transportation and assignment problems, graphical and simplex linear programming, Big-M, shadow prices, Excel Solver and project-network scheduling through CPM and PERT [17], [18], [19]. Table II separates directly used methods from structural analogies and possible classroom extensions. The mapping is our interpretation; the authors did not necessarily use classroom algorithms. Collateral allocation resembles an extended transportation problem, loan selection is a binary subset-choice problem, and not every OR method reduces to an LP.

TABLE II. Course techniques and the reviewed studies

Course technique	Where it appears or applies	Role in the credit setting
Linear programming formulation	[2] and the CVaR LP of [4]; contrast: [8] uses QP	LP forms under stated assumptions; quadratic objective in [8]
Simplex, slack, surplus, artificial variables, Big-M	LP stages of [2] and [12]; classroom simplex tables not claimed	Surplus variables and Big-M or two-phase methods are classroom analogies, not reported procedures
Shadow prices and the final table	Dual interpretation of LP models; no simplex final table reported	Local change in the optimal objective per unit of a constraint; values scarce capital or collateral
Transportation problem	[12], structural analogy only	Securities as supply, exposures as demand; ineligible securities are prohibited routes
Assignment problem and Hungarian method	Not implemented	Matching is one-to-one; loan subset selection differs
Network techniques: CPM, PERT, time-cost trade-off	Not used in any of the 14 studies	Loan-approval timelines are scheduling problems; PERT times parallel the uncertainty theme
Beyond the syllabus	Stochastic and robust optimisation, dynamic programming, simulation, metaheuristics, queueing	Extend the deterministic toolkit to uncertainty, sequential decisions and large problems

G. Integrated decision workflow (proposed synthesis)

An institution could connect the reviewed models in six stages: data and validation, dependence-aware scenarios, tail-risk estimation, constrained portfolio and loan-pricing choices, collateral and service implementation, and monitoring with re-optimisation. This conceptual design is proposed by this review and has not been tested jointly. The supervisory alignment in Fig. 2 follows the four credit-risk management areas in [20].

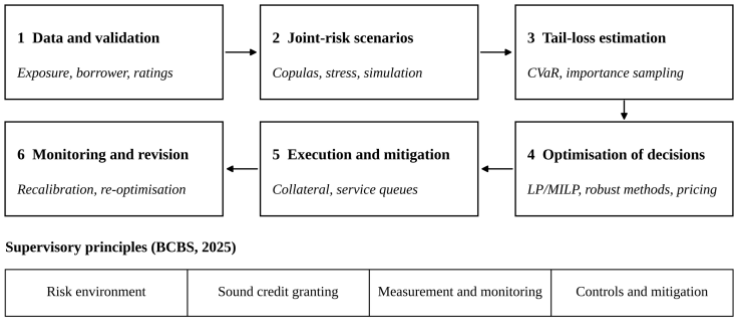


Fig. 2. Integrated OR workflow for credit-risk management and its supervisory context. Original conceptual synthesis of the 14 reviewed studies; the four supervisory domains follow [20], and the six-stage workflow is proposed by the authors and is not an official BCBS model.

IV. Conclusion

A. Conclusions

The conclusions are narrative syntheses, not experimentally proven comparisons. First, no single technique is established as universally superior; LP/MILP, stochastic and robust optimisation, simulation, control, quadratic programming and heuristics suit different structures, and no common-data comparison was run. Second, dependence modelling can materially alter extreme-loss estimates, and correlation alone does not identify tail dependence, although its importance relative to other estimation errors depends on the portfolio. Third, approximation is often a practical design choice, but its error should be evaluated. Fourth, pricing, risk measurement, portfolio choice, dynamic investment and collateral management can be linked into a decision cycle; the bank-queueing application is a related operational problem, not a stage of default management.

B. Benefits, drawbacks and gaps

OR makes assumptions, objectives and constraints explicit so decisions can be examined and challenged. Reported benefits include multistage market–credit modelling, reduced sensitivity to model error, efficient tail simulation, portfolio-aware loan pricing, large-scale selection and an illustrative ₹268 per hour reduction in modelled queueing cost (theoretical, not observed). The main drawbacks are fourfold. Parameter dependence: scenario generation, copula dependence, default intensities, recoveries and acceptance functions must be estimated and periodically reassessed. Simplifying assumptions: negligible loan correlation [8], uncorrelated rating transitions [2], a single risky asset with constant rates [9], and asymptotic results that need large portfolios [5], [10]. Uneven evidence: some studies use real data and backtesting, others simulated examples or a single institution, so results may not generalise across portfolios or market regimes. Fragmentation: studies use different data and metrics, so performance claims cannot be compared directly, and few studies combine robustness, dynamics and dependence in one model.

C. Limitations and future research

This review rests on five supplied student reports, not an exhaustive journal search, and most studies were not re-implemented, so improvements cannot be pooled or ranked. A stronger project would document formal search criteria, verify numerical tables against the originals, benchmark at least two methods on one open loan dataset with a dependence-aware extension, perturb uncertain inputs, and report out-of-sample risk and computation time with reproducible code. The clearest research opportunities are models combining robustness with realistic dependence, multi-asset versions of the dynamic control results, approximation schemes with proven error bounds, and validation on large real portfolios across full economic cycles. Course-linked extensions include small LP models in Excel Solver to read shadow prices, and PERT/CPM applied to loan-approval and project-finance timelines, which none of the 14 studies address.

References

[1] H. Markowitz, “Portfolio selection,” *J. Finance*, vol. 7, no. 1, pp. 77–91, 1952.

[2] P. Pasricha, D. Selvamuthu, G. D’Amico, and R. Manca, “Portfolio optimization of credit risky bonds: A semi-Markov process approach,” *Financial Innovation*, vol. 6, Art. no. 25, 2020.

[3] N. J. Jobst, G. Mitra, and S. A. Zenios, “Integrating market and credit risk: A simulation and optimisation perspective,” *J. Banking Finance*, vol. 30, no. 2, pp. 717–742, 2006.

[4] G. Masala, M. Menzietti, and M. Micocci, "Optimisation of conditional-VaR in an actuarial model for credit risk assuming a Student copula dependence structure," *Investment Manage. Financial Innovations*, vol. 4, no. 3, pp. 39–56, 2007.

[5] A. Bassamboo, S. Juneja, and A. Zeevi, "Portfolio credit risk with extremal dependence: Asymptotic analysis and efficient simulation," *Oper. Res.*, vol. 56, no. 3, pp. 593–606, 2008.

[6] P. Glasserman, W. Kang, and P. Shahabuddin, "Fast simulation of multifactor portfolio credit risk," *Oper. Res.*, vol. 56, no. 5, pp. 1200–1217, 2008.

[7] L. Bo and A. Capponi, "Robust optimization of credit portfolios," *Math. Oper. Res.*, vol. 42, no. 1, pp. 30–56, 2017.

[8] G. Chi, S. Ding, and X. Peng, "Data-driven robust credit portfolio optimization for investment decisions in P2P lending," *Math. Problems Eng.*, vol. 2019, Art. no. 1902970, 2019.

[9] L. Bo and A. Capponi, "Optimal credit investment with borrowing costs," *Math. Oper. Res.*, vol. 42, no. 2, pp. 546–575, 2017.

[10] J. A. Sirignano, G. Tsoukalas, and K. Giesecke, "Large-scale loan portfolio selection," *Oper. Res.*, vol. 64, no. 6, pp. 1239–1255, 2016.

[11] A. Germain and F. Vrins, "Credit selection in collateralized loan obligation: Efficient approximation through linearization and clustering," *Eur. J. Oper. Res.*, vol. 330, no. 1, pp. 326–342, 2026.

[12] J. Goblet, F. Willame, and J. Grauls, "A mixed-integer linear programming approach for collateral management," *Decision Analytics J.*, vol. 10, 2024.

[13] S. Y. Chun and M. A. Lejeune, "Risk-based loan pricing: Portfolio optimization approach with marginal risk contribution," *Manage. Sci.*, vol. 66, no. 8, pp. 3735–3753, 2020.

[14] Z. Wang, X. Zhang, Z. Zhang, and D. Sheng, "Credit portfolio optimization: A multi-objective genetic algorithm approach," *Borsa Istanbul Rev.*, vol. 22, no. 1, pp. 69–76, 2022.

[15] N. Jhala and P. Bhathawala, "Application of queuing theory in banking sector," *IOSR J. Math.*, vol. 12, no. 2, pp. 73–75, 2016.

[16] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *J. Risk*, vol. 2, no. 3, pp. 21–41, 2000.

[17] SVKM's NMIMS School of Commerce, "Operations research: B.Sc. Finance Semester III course teaching plan (2026–2027)," course outline, 2026.

[18] F. S. Hillier, G. J. Lieberman, B. Nag, and P. Basu, *Introduction to Operations Research*, 12th ed. McGraw Hill, 2025.

[19] H. A. Taha, *Operations Research: An Introduction*, 10th ed. Pearson, 2019.

[20] Basel Committee on Banking Supervision, "Principles for the management of credit risk," Bank for International Settlements, 2025.