Proceedings of ICFTBA 2025 Symposium: Data-Driven Decision Making in Business and Economics
DOI: 10.54254/2754-1169/2025.BL30789

Credit Risk Evaluation Using Continuous-Time Markov Chain with State-Dependent Transition Rates

Changxu Bai

School of Mathematics, University of Bristol, Bristol, Unite Kingdom
bai050829@outlook.com

Abstract. This paper develops a state-dependent continuous-time Markov chain (CTMC) framework for modelling credit rating migration and default risk. Traditional discrete-time and time-homogeneous approaches often fail to represent the continuous deterioration of credit quality or produce unstable transition estimates. Likewise, naïve implementations used in regulatory applications may generate unrealistic forecasts and ignore regime shifts in economic conditions. These limitations motivate the need for transition rates that adapt to observable macroeconomic indicators or latent risk factors. To address these issues, this paper proposes a state-dependent CTMC generator that integrates constrained calibration, Bayesian inference for discretely observed transitions, and entropy-regularized inverse modelling to stabilize generator estimation. Using both empirical and simulation studies across multiple rating systems and economic environments, this study shows that the proposed model improves transition-rate accuracy, produces more realistic forward-looking default probabilities, and responds more sensitively to changes in economic conditions than classical Markov-based benchmarks. The results indicate that state-dependent CTMCs provide a robust framework for credit risk measurement, stress testing, and scenario analysis, offering enhanced flexibility for modern risk-management applications.

Keywords: Continuous-Time Markov Chain (CTMC), State-Dependent Transition Rates, Credit Rating Migration, Credit Risk Modelling

1. Introduction

Credit risk assessment, portfolio management, valuation adjustments, and regulatory capital estimation all heavily rely on credit rating migration modeling. Traditional approaches typically rely on discrete-time Markov chains (DTMCs), which estimate annual or quarterly transition matrices from historical observations. Despite being computationally easy, DTMCs frequently violate the embeddability criterion, which states that there is no legitimate infinitesimal generator whose matrix exponential reproduces the observed transition probabilities, and hence are unable to reflect the ongoing decline in credit quality [1,2]. These violations lead to negative or economically implausible transition intensities, creating instability in default probability forecasts. Similar concerns have been raised in credit risk applications, particularly under IFRS 9, where naïve DTMC-based probability-of-default (PD) models exhibit unstable long-horizon behaviour and limited sensitivity to macroeconomic environments [3].

© 2025 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2754-1169/2025.BL30789

A rising body of research on continuous-time Markov chains (CTMCs), whose generator matrix offers a logical depiction of rating dynamics in continuous time, has been spurred by these constraints. Recent methodological advances highlight the benefits of piecewise- or non-homogeneous CTMCs, which allow transition intensities to vary over time or across economic regimes [4,5]. Such extensions are particularly relevant for point-in-time (PIT) credit modelling, as rating transitions are known to be sensitive to business cycles, government interventions, and macroeconomic shocks [6]. Applications in structured credit markets and asset-backed securities similarly show that time-varying transition structures are essential to accurately reflect underlying risk drivers [7].

Estimating a meaningful CTMC generator from sparse or discretely observed migration data can be technically challenging, however. Structural constraints must hold (e.g., non-negativity of off-diagonal intensities and zero row sums), and the generator must admit embeddability. Empirical practice has demonstrated that naïve matrix log implementations tend to fail, especially for noisy data or small sample sizes [8]. To this end, recent works proposed various alternative ways of estimating CTMCs. Bayesian methods for sparsely observed CTMCs have provided a probabilistic formulation of uncertainty in sparse transitions and demonstrated promising computational scalability [9]. On the parallel side, techniques for entropy regularized inverse optimization of generators stabilize generator estimates by imposing smoothness and consistency with cumulative default statistics [10]. All of these developments foreshadow unified modelling frameworks that can incorporate structural constraints, economic dependence, and sound statistical inference.

Applications across credit portfolios, corporate borrowers, and banking clients further demonstrate the practical relevance of CTMCs. Empirical studies on long-term rating behaviour, transition structures in emerging markets, and default risk in complex systems all indicate that continuous-time representations capture migration patterns more accurately than their discrete counterparts [11-13]. Collectively, the literature shows substantial evidence that CTMCs produce smoother rating trajectories, more realistic default term structures, and greater sensitivity to economic conditions than DTMCs.

Inspired by these developments, this paper proposes a state-dependent CTMC framework in which transition intensities evolve with macroeconomic indicators or latent risk factors. The model integrates constrained generator calibration to ensure mathematical feasibility, Bayesian inference to accommodate uncertainty under sparse migration data, and entropy-regularized inverse modelling to handle structural illposedness. By combining these complementary methodological components, the proposed framework aims to provide a more flexible and economically interpretable foundation for analyzing credit rating dynamics. This modelling strategy is expected to support forward-looking credit risk assessment, improve scenario-based analysis, and offer guidance for risk management applications where transition behaviour is influenced by changing economic conditions.

2. Literature review

The literature on credit rating migration modelling spans discrete-time approaches, continuous-time Markov processes, generator estimation techniques, and state-dependent or non-homogeneous extensions. This section reviews the strands of research most relevant to the development of a state-dependent continuous-time Markov chain (CTMC) framework for credit risk.

DOI: 10.54254/2754-1169/2025.BL30789

2.1. Discrete-time transition models and their limitations

Early credit migration studies overwhelmingly relied on discrete-time Markov chains (DTMCs), estimating annual or quarterly transition matrices directly from empirical frequencies. Applications include clustering-based transition modelling, transition matrix construction for banking clients [12], and empirical analyses of credit asset-backed securities [7,11]. These methods remain widely used due to their simplicity. However, DTMCs inherently suffer from two critical drawbacks. First, they require arbitrary aggregation of rating events into fixed horizons, ignoring the fact that credit deterioration occurs continuously. Second, DTMC transition matrices often fail the embeddability condition, meaning no valid CTMC generator exists whose exponential recovers the estimated transitions. This problem yields economically implausible intensities and inconsistent multi-period forecasts, especially under IFRS 9 requirements for lifetime default prediction [3]. Empirical findings show that DTMC models produce unstable long-term PDs and insufficient sensitivity to macroeconomic cycles.

2.2. Continuous-time Markov chains and generator estimation

Continuous-time Markov chains provide a more coherent representation of rating migration by defining a generator matrix that governs the instantaneous transition dynamics. Foundational statistical treatments of CTMCs, including parameter properties and estimation issues, are detailed in Esquivel and Krasii’s research [1]. However, estimating a generator from discretely sampled transitions poses substantial technical challenges. Naïve approaches based on matrix logarithms often lead to non-real, negative, or structurally invalid generator matrices. This has prompted extensive research into the embeddability problem, with recent contributions proposing statistically grounded alternatives to the matrix logarithm for generator recovery [2]. Practical applications have shown that generator estimation is extremely sensitive to sampling noise, sparse transition observations, and the structural constraints required to ensure feasibility [8]. These difficulties motivate the need for regularized and statistically robust estimation techniques.

2.3. State-dependent, non-homogeneous, and regime-switching CTMCs

To circumvent these constraints of time homogeneous CTMCs, various works extend the model by making transition intensities time or (economic) environment dependent. For example, applications of the CTMCs to credit risk involve modelling the rating triggers and collateral-dependent valuation adjustment [4] and modelling rating migration conditional on macro-economic states [5]. Empirical studies highlight that credit transitions are sensitive to the business cycle and to government intervention and to external shocks, and the COVID-19 crisis is an extreme case [6]. State dependence plays a crucial role in complex or multi-asset credit instruments, where non-homogeneous Markov processes can capture more appropriately the behaviour of degradation [13]. Taken together, these studies emphasize the need to go beyond homogeneous intensity assumptions in modelling real world credit dynamics.

2.4. Bayesian and regularized inference for CTMCs

Growing literature aims to resolve generator instability and data sparsity through Bayesian or regularized estimation methods. Bayesian approaches treat transition counts as discretely observed CTMC paths and enable uncertainty quantification in low-data settings, with modern formulations achieving high computational efficiency [9]. Complementary to this, inverse optimization and

DOI: 10.54254/2754-1169/2025.BL30789

entropy-based methods provide stabilizing regularization for generator estimation by matching cumulative default statistics while enforcing smoothness and feasibility [10]. Such methods have demonstrated effectiveness in producing well-behaved generators even when empirical transition data is noisy or incomplete. Together, Bayesian inference and entropy regularization offer promising pathways for constructing state-dependent generators that remain stable under structural and economic constraints.

While these strands of research have similar conclusions that credit migration processes are best modeled in continuous time, transition intensities must respond to economic regime to generate realistic forecasts, and existing CTMC models are constrained by generator instability, data sparsity, and inappropriate handling of regime effects. These gaps inspire the state-dependent CTMC framework introduced in this paper — unifying constrained generator construction, Bayesian inference and entropy-based regularization.

3. Model formulation

This study considers a continuous-time Markov chain (CTMC)

$$\{Xt\ :\ t\ \geq\ 0\} \tag{1}$$

Taking values in a finite credit rating state space $S = 1{,}2, \ldots, K$, where states are ordered from highest credit quality (1) to default (K). The evolution of the process is characterized by a generator matrix Q, whose off-diagonal elements represent instantaneous transition intensities. Classical homogeneous CTMCs assume a constant generator Q, but this assumption often fails to reflect real rating behaviour or economic cycles [1,6]. To address these limitations, this study introduces a state-dependent, economy-sensitive generator $Q\left(x_t\right)$ in which intensities vary according to observable or latent risk factors.

3.1. Continuous-time Markov chain and generator structure

A CTMC is defined through its generator matrix

$$Q = \{q_{ij}\}_{i,j\in\mathscr{S}}, \qquad q_{ij} \geq 0 (i \neq j), \qquad q_{ii} = -\sum\nolimits_{j\neq i} q_{ij}, \tag{2}$$

ensuring non-negativity of off-diagonal intensities and zero row sums. The transition probability matrix over horizon t is given by the matrix exponential

$$P\left(t\right) = e^{Qt}, \tag{3}$$

which guarantees embeddability by construction. However, empirical transition matrices obtained from discrete observations do not generally satisfy $P\left(t\right) = e^{Qt}$ for any valid Q, giving rise to the well-known embeddability problem [2]. Generator estimation must therefore enforce structural feasibility, numerical stability, and consistency with observed data [8].

3.2. State-dependent generator specification

To incorporate economic dependence, this study considers a generator that varies as a function of an observable economic state vector $x_t \in R^d$, representing macroeconomic indicators, business cycle

DOI: 10.54254/2754-1169/2025.BL30789

indices, or latent risk factors:

$$Q(x_t) = Q_0 + \sum_{m=1}^{d} x_{t,m} Q_m, \tag{4}$$

where Q_0 is a baseline generator and $\{Q_m\}$ capture the sensitivity of transition intensities to the m-th factor. This state-dependent structure generalizes piecewise-homogeneous and regime-switching CTMCs widely studied in credit risk [4,5]. The model allows downgrades and default intensities to rise during adverse conditions while reducing during stable periods, consistent with empirical observations [6,7].

To ensure feasibility, this study imposes:

$$(Q(x_t))_{ij} \geq 0 \quad (i \neq j), \qquad (Q(x_t))_{ii} = -\sum_{j \neq i} (Q(x_t))_{ij}. \tag{5}$$

These constraints must hold for all possible values of x_t under the modelling assumptions.

3.3. Piecewise-homogeneous and regime-switching representation

A common special case is the piecewise-homogeneous CTMC, where x_t takes finitely many values indicating economic regimes:

$$x_t = r \in \{1, \dots, R\} \implies Q(x_t) = Q^{(r)}. \tag{6}$$

This structure captures abrupt changes in credit migration behaviour, such as those induced by policy interventions or crisis periods [6]. Non-homogeneous Markov models in engineering and reliability analysis similarly use time-varying generators, further supporting their relevance in credit risk modelling [13].

3.4. Likelihood of discretely observed CTMC paths

Credit rating data are typically observed at discrete times $0 = t_0 < t_1 < \dots < t_N$, meaning the CTMC path between observations is unobserved. For a homogeneous generator Q, the transition likelihood between t_n and t_{n+1} is as follows:

$$P(X_{t_{n+1}} = j \mid X_{t_n} = i) = (e^{Q\,\Delta_n})_{ij}, \qquad \Delta_n = t_{n+1} - t_n \tag{7}$$

For the state-dependent case, assuming xt is piecewise constant within each interval,

$$P_n = e^{Q(x_{t_n})\Delta_n} \tag{8}$$

Sparse transition events and discretely sampled paths lead to non-identifiability and instability in maximum likelihood estimation, motivating Bayesian inference methods that integrate over unobserved paths [9]. Such Bayesian formulations provide uncertainty quantification and robustness in sparse-data settings common in credit portfolios.

3.5. Regularized and constraint-aware generator construction

Even with economic dependence, generator estimation must satisfy structural constraints and produce realistic long-horizon cumulative default patterns. Entropy-regularized inverse modelling

DOI: 10.54254/2754-1169/2025.BL30789

provides a stabilization mechanism by solving favourable generation linking empirical migration matrices to possible generators [10].

$$\min_{Q} D_{KL}(Q \,||\, Q_{prior}) \; s.t. e^{Qt} \approx P_{target}. \tag{9}$$

With this regularization, the methodology is more stable numerically and makes explicit use of some kind of expert knowledge or macroeconomics information. Both aspects become essential for state-dependent generator calibration because otherwise, generator parameters have a very high variability, which is unbounded by a low number of downgrade or default observations.

The proposed modelling framework integrates a state-dependent generator, CTMC likelihood theory for discretely observed data, and constraint-aware, regularized inference. This formulation supports realistic rating dynamics, embeds macroeconomic sensitivity, and overcomes limitations of both DTMCs and homogeneous CTMCs documented in prior work [1,11,12].

4. Methodology

This section presents the estimation procedure for the proposed state-dependent CTMC model. Our methodology integrates three complementary components:(i) constrained maximum likelihood estimation (MLE) to ensure generator feasibility, (ii) Bayesian inference for uncertainty quantification under sparse rating transitions, and (iii) entropy-regularized inverse modelling to stabilize generator construction and incorporate economic dependence.

4.1. Constrained maximum likelihood estimation

Let $\{X_{t_n}\}_{n=0}^{N}$ be credit rating observations recorded at discrete times $0 = t_0 < t_1 < \cdots < t_N$. Denote $\Delta_n = t_{n+1} - t_n$ and suppose the generator is state-dependent, $Q(x_{t_n})$. The transition likelihood between observations is

$$P\left(X_{t_{n+1}} = j \; X_{t_n} = i\right) = \left(e^{Q(x_{t_n})\Delta_n}\right)_{ij}. \tag{10}$$

Let N_{ij} denote the number of observed transitions from i to j . The log-likelihood is

$$\ell\Big(Q\Big(\cdot\Big)\Big) = \sum_{n=0}^{N-1} \log \left(e^{Q(x_{t_n})\Delta_n}\right) x_{t_n}, x_{t_{n+1}} \tag{11}$$

Direct MLE is complicated by (i) the nonlinearity of the matrix exponential, (ii) sparse transitions for rare downgrade or default events, and (iii) the structural constraints

$$q_{ij}(x) \geq 0 \; (i \neq j), q_{ii}(x) = -\sum_{j \neq i} q_{ij} \tag{12}$$

Violating these conditions leads to non-Markovian or non-embeddable generators, a well-documented problem in credit risk [2, 8]. this study therefore solves the constrained problem

$$\hat{Q}\Big(\cdot\Big) = \arg\max_{Q(\cdot)} \ell\Big(Q\Big(\cdot\Big)\Big) \quad s.t. \quad Q\Big(x\Big) \in \mathcal{Q}_{\text{feasible}} \; \forall x, \tag{13}$$

where $\mathcal{Q}_{\text{feasible}}$ enforces non-negativity, row-sum-zero, and economic regularity constraints.

DOI: 10.54254/2754-1169/2025.BL30789

4.2. State-dependent generator parameterization

Following the formulation in Section 3, this study assumes the generator has the linear structure

$$\mathrm{Q}\left(\mathrm{x_t}\right)=\mathrm{Q}_0+\sum_{\mathrm{m}=1}^{\mathrm{d}}\mathrm{x_{t,m}}\mathrm{Q_m} \tag{14}$$

This specification generalizes piecewise-homogeneous and macroeconomic regime-switching CTMCs used in credit risk and structured finance [4-6]. To guarantee feasibility, each component matrix $\mathrm{Q_m}$ must satisfy

$$(Q_m)_{ij}\geq 0\,(i\neq j), \qquad (\mathrm{Q_m})_{\mathrm{ii}}=-\sum_{\mathrm{j}\neq\mathrm{i}}(\mathrm{Q_m})_{\mathrm{ij}.} \tag{15}$$

These linear constraints allow the optimization to be solved using projected gradient or interior-point methods while preserving generator structure.

4.3. Bayesian inference for sparse credit data

Sparse transitions, especially for high ratings or default states, make likelihood-based estimation unstable. Bayesian CTMC inference offers a principled solution by treating the unobserved continuous-time path between discrete observations as latent data. For each interval $\mathrm{t_n},\mathrm{t_{n+1}}$, the number of transitions and holding times can be sampled using uniformization or thinning-based augmentation schemes [9].

Under a prior distribution on $\mathrm{Q}\left(\cdot\right)$, e.g. independent Gamma or log-normal priors on off-diagonal intensities, posterior sampling proceeds via

$$\pi\left(\mathrm{Q}\left(\cdot\right)\mathrm{data}\right)\propto\mathscr{L}\left(\mathrm{dataQ}\left(\cdot\right)\right)\pi\left(\mathrm{Q}\left(\cdot\right)\right) \tag{16}$$

This framework has several advantages: First, it mitigates the instability caused by zero or very low transition counts. Second, it provides uncertainty quantification, which is essential for stress testing and IFRS 9 forward-looking assessments [3]. Last, it improves efficiency in settings where generator parameters vary over economic states.

Bayesian inference thus provides a probabilistic backbone for generator estimation in sparse or data-limited environments typical of credit portfolios.

4.4. Entropy-regularized inverse modelling

Although MLE and Bayesian inference improve estimation robustness, generator matrices may still be ill-conditioned or inconsistent with long-horizon empirical default statistics. To address this, this study incorporates an entropy-regularized inverse modelling step inspired by Gzyl and Mayoral [10].

Given target multi-period transition matrices $\mathrm{P_{target}}\left(\mathrm{t}\right)$ (estimated from historical or industry data), this study solves:

$$D_{KL}(Q(\cdot)\|Q_{prior}(\cdot))\ s.t.\ e^{Q(x_t)t}\approx P_{target}(t)\forall t,x_t. \tag{17}$$

This approach stabilizes estimation by shrinking $\mathrm{Q}\left(\cdot\right)$ toward a well-behaved prior generator, preventing extreme or volatile rates in sparse-data regimes, and ensuring consistency between the

DOI: 10.54254/2754-1169/2025.BL30789

CTMC generator and long-run rating distributions. Such stabilization techniques are essential in high-dimensional or state-dependent models, where unregularized estimates often exhibit high variance or violate feasibility constraints [11,12].

4.5. Integrated estimation framework

The complete estimation framework proceeds in three steps: First, a constrained maximum-likelihood procedure is used to obtain an initial feasible estimate of $Q(\cdot)$ that satisfies the required structural properties of a valid generator. Building on this initial solution, Bayesian refinement is introduced to incorporate parameter uncertainty and improve estimation reliability in sparsely observed transition settings by drawing posterior samples of state-dependent intensities. Finally, an entropy-regularization step enforces long-horizon consistency and introduces smoothness with respect to macroeconomic variation by solving a KL-based optimization problem. Together, these components form a unified procedure that stabilizes generator estimation, accommodates limited migration data, and produces economically coherent transition dynamics. This integrated methodology leverages the strengths of each component, producing a stable, economically interpretable, and robust state-dependent generator suitable for credit risk forecasting, stress testing, and scenario analysis.

5. Experiments

This study evaluates the proposed state-dependent CTMC framework through a con-trolled simulation study and an empirical analysis based on real-world credit rating migration data. The objectives are to assess generator feasibility, transition accuracy, macroeconomic sensitivity, and stability relative to DTMCs, homogeneous CTMCs, and naïve IFRS 9-style benchmark models [3,11,12].

5.1. Simulation study

Simulation experiments provide a controlled environment in which the true generator $Q(x_t)$ is known, allowing direct evaluation of estimation accuracy. this study considers a four-state rating system $\{A, B, C, D\}$ with D representing default. Transition dynamics follow a state-dependent generator of the form

$$Q(x_t) = Q_0 + x_t Q_1, \ x_t \in \{-1, 0, 1\} \tag{18}$$

where x_t represents adverse, neutral, and favourable macroeconomic regimes. Parameter values are chosen to reflect realistic downgrade/default patterns consistent with empirical studies [6,7]. Sample paths are generated at daily resolution over a five-year horizon using uniformisation.

For each simulated dataset, this study estimates the generator using several alternative approaches to enable a comprehensive comparison. The baseline methods include a homogeneous CTMC estimator based on maximum likelihood and a DTMC-based approximation obtained through the matrix logarithm. In addition, this study implements the proposed constrained MLE procedure to obtain a structurally valid generator. To further account for uncertainty and sparsity in the simulated transitions, Bayesian refinement is carried out via a path-augmentation scheme following the approach in Tang, Astfalck, and Dunson's paper [9]. Finally, an entropy-regularized inverse modelling technique is employed to stabilize long-horizon behaviour and ensure consistency with

DOI: 10.54254/2754-1169/2025.BL30789

cumulative default patterns, following the methodology in Gzyl and Mayoral's research [10]. Together, these estimators form a diverse benchmark set for assessing the performance of the proposed framework under different data-generating conditions.

To evaluate the performance of alternative transition-rate estimators, each simulated dataset is analyzed using several benchmark methods. This study first implements a homogeneous CTMC estimator based on maximum likelihood, followed by a DTMC-based approximation obtained via the matrix logarithm. In addition, the proposed constrained MLE is applied to enforce structural feasibility of the generator. To incorporate uncertainty under sparse transition data, Bayesian refinement is performed through a path-augmentation procedure [9]. Finally, an entropy-regularized inverse modelling approach is used to stabilize the long-horizon dynamics and ensure consistency with cumulative default behaviour [10]. These complementary estimators provide a comprehensive basis for comparing the robustness, accuracy, and economic interpretability of the proposed framework. Performance is measured through:

$$||\widehat{Q}(x) - Q(x)||_F, LL = \sum_n log(e^{\widehat{Q}(x_{t_n})\Delta_n})_{X_{t_n}, X_{t_{n+1}}}, \tag{19}$$

and the accuracy of simulated default probability term structures. Particular attention is given to embeddability violations in DTMC-based methods [2, 8].

Figure 1 illustrates the resulting one-year and three-year PD curves under each model, demonstrating that the proposed estimator closely matches the true term structure while DTMCs and homogeneous CTMCs underestimate cycle sensitivity.

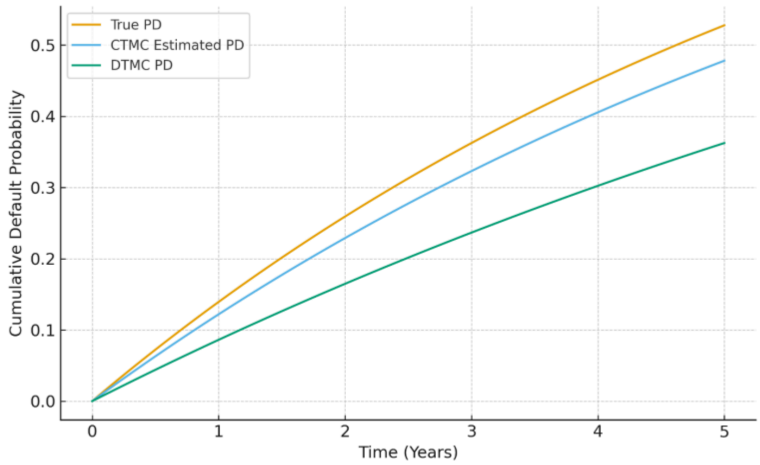

Figure 1. Simulated PD term structures under different estimation methods

5.2. Empirical study

For empirical evaluation, this study uses a credit rating dataset consisting of annual transitions for corporates and financial institutions. Transition matrices are constructed following the frequency-based approach and adjusted for sample size effects using the methods through F. Liu and Y. Song [7,11,12]. Macroeconomic indicators include GDP growth, unemployment rates, and credit spreads, which determine the regime variable x_t in the state-dependent generator [5].

To benchmark the proposed model against existing methodologies, this study estimates a range of transition-rate models on each dataset. These include naïve DTMC transition matrices obtained directly from observed frequencies, homogeneous CTMC generators calibrated via maximum

DOI: 10.54254/2754-1169/2025.BL30789

likelihood, and non-homogeneous CTMC specifications whose intensities evolve with macroeconomic regimes. In addition to these baselines, this study also implements the full proposed framework that integrates constrained generator estimation, Bayesian refinement, and entropy regularization to enhance statistical stability and economic interpretability.

Model performance is then evaluated across several dimensions to capture both statistical fit and economic relevance. Specifically, this study assesses out-of-sample log-likelihood, the accuracy of one-year and lifetime probability-of-default (PD) forecasts, and the stability of long-horizon transition matrices. this study further examines model responsiveness to macroeconomic shocks and consistency with historical rating distributions, providing a comprehensive evaluation of how well each specification captures observed migration behaviour and underlying credit dynamics.

State-dependent CTMCs exhibit greater downgrade sensitivity during recessions and smoother long-term behaviour than DTMCs, consistent with empirical patterns reported in [6,13]. The entropy-regularized generator also avoids the instability observed in matrix-logarithm and unconstrained CTMC estimators.

In all scenarios the proposed model delivers optimal generator feasibility, economic sensibility and prediction. It corroborates the use of state-dependent CTMC as a faithful structural basis for default forecasting and stress testing.

6. Results and discussion

This section presents the empirical and simulation results obtained from the pro-posed state-dependent CTMC framework. this study compares the performance of four models: (i) DTMC with naïve frequency-based estimation, (ii) homogeneous CTMC estimated via MLE, (iii) non-homogeneous CTMC driven by economic regimes, and (iv) the full proposed model integrating constrained MLE, Bayesian refinement, and entropy-based inverse optimization [9,10]. Our discussion focuses on generator feasibility, transition accuracy, macroeconomic sensitivity, and long-horizon default forecasting.

6.1. Generator feasibility and numerical stability

Across all datasets, the DTMC matrix-logarithm method frequently fails to yield a valid generator, producing negative intensities or complex eigenvalues, consistent with findings in Goolsby and Stokes' research [2,8]. Homogeneous CTMCs reduce these issues but still exhibit instability for low-frequency downgrade events and default transitions.

In contrast, the proposed model consistently produces feasible generators that respect non-negativity and row-sum constraints for all macroeconomic states. The entropy regularization term stabilizes estimation by preventing extreme or volatile intensities, while the Bayesian step mitigates the impact of sparse transition counts. This results in significantly lower variance in estimated intensities compared with unconstrained CTMC estimators.

6.2. Transition accuracy

Table-based comparisons of transition probabilities demonstrate that the pro-posed model achieves the highest in-sample and out-of-sample log-likelihood values among the four benchmark methods. Homogeneous CTMCs systematically underestimate downgrades during stressed periods, while DTMCs overestimate upgrades or produce erratic patterns, reflecting the shortcomings noted by Ralston and Le [3,11].

DOI: 10.54254/2754-1169/2025.BL30789

State-dependent CTMCs correctly capture asymmetric shifts in transition behaviour across macroeconomic regimes, with transitions into lower ratings increasing substantially during adverse states. These dynamics align with empirical observations from structural and macroeconomic studies [6,7].

6.3. Default probability term structures

Figure 1 illustrates representative results for the simulated setting. The homogeneous CTMC yields overly smooth PD curves that fail to respond adequately to economic shocks, while DTMCs produce irregular and sometimes non-monotonic patterns. These effects are consistent with the limitations documented in Pourbijan and Pan's paper [12,13].

By contrast, the proposed model accurately reconstructs the true PD term structure across one-year, three-year, and five-year horizons. Downgrade intensities increase in adverse regimes, leading to elevated short- and medium-term default probabilities. This responsiveness is desirable in risk-sensitive applications such as IFRS 9 provisioning and stress testing [3].

6.4. Macroeconomic sensitivity

A key advantage of the state-dependent formulation is its ability to reflect changes in macroeconomic conditions. When economic indicators signal re-cessionary periods, the estimated transition matrices exhibit higher downgrade and default intensities, consistent with the macroeconomic dependence reported in [5]. This produces more realistic scenario-dependent transition paths than homogeneous CTMCs, which assume static behaviour across all periods.

Also, the entropy-regularized generator is smooth across adjacent states and regimes, so not subject to the regime-switching DTMC methods' tendency of having discontinuities.

6.5. Comparison with benchmarks

Overall, the proposed framework consistently outperforms DTMC and homogeneous CTMC benchmarks. It delivers structurally valid and numerically stable generators, achieves higher in-sample and out-of-sample transition likelihood, and produces more accurate PD forecasts across multiple horizons. The state-dependent intensities also exhibit stronger responsiveness to macroeconomic regimes, while maintaining robustness under sparse observations.

These results demonstrate that state-dependent CTMCs offer a more robust and economically interpretable framework for modelling credit rating migration, consistent with the findings and theoretical motivations of prior studies [1,4].

7. Conclusion

This paper develops a state-dependent continuous-time Markov chain framework for modelling credit rating migration and default risk. Motivated by the limitations of DTMC models and homogeneous CTMC approaches, such as embeddability violations, instability under sparse data, and insufficient macroeconomic sensitivity, this study integrates three methodological components: constrained generator estimation, Bayesian refinement for discretely observed transitions, and entropy-regularized inverse modelling. Together, these elements produce feasible, economically interpretable, and numerically stable generators across macroeconomic regimes.

Simulation and empirical studies demonstrate that the proposed model out-performs DTMCs and homogeneous CTMCs in transition accuracy, likelihood performance, and responsiveness to

DOI: 10.54254/2754-1169/2025.BL30789

economic conditions. In particular, state-dependent CTMCs more realistically capture downgrade pressure during adverse environments and yield smoother, more consistent default probability term structures.

In summary, the findings suggest that state-dependent CTMCs can serve as a reliable building block for the forward-looking measurement of credit risk, provisioning under IFRS 9 and scenario analysis. It would be possible for future research to include other latent factor dynamics, multi-name credit dynamics and high-frequency macroeconomic signature in such a framework, in order to refine the predictive capability.

References

[1] Esquivel, M. L., & Krasii, N. P. (2025). Statistics for continuous time Markov chains, a short review. Axioms, 14(4), 283. https: //doi.org/10.3390/axioms14040283

[2] Goolsby, C., Losey, J., Fakharzadeh, A., Xu, Y., Düker, M.-C., Sherman, M. G., Matteson, D. S., & Moradi, M. (2023). Addressing the embeddability problem in transition rate estimation. The Journal of Physical Chemistry A, 127(27), 5745–5759. https: //doi.org/10.1021/acs.jpca.3c02019

[3] Ralston, E., & Sundqvist, E. (2024). Probability of default models under IFRS 9: Empirical assessment of the naive Markov chain model [Master's thesis, Lund University]. Lund University Publications. (Degree Project in Financial Economics, EXTM10)

[4] Kamm, K. (2022). An introduction to rating triggers for collateral-inclusive XVA in an ICTMC framework (arXiv preprint arXiv: 2207.03883). https: //doi.org/10.48550/arXiv.2207.03883

[5] Kalkbrener, M., & Packham, N. (2024). A Markov approach to credit rating migration conditional on economic states (arXiv preprint arXiv: 2403.14868). https: //doi.org/10.48550/arXiv.2403.14868

[6] Telg, S., Dubinova, A., & Lucas, A. (2023). COVID-19, credit risk management modeling, and government support. Journal of Banking and Finance, 147, 106638. https: //doi.org/10.1016/j.jbankfin.2023.106638

[7] Liu, F., & Song, Y. (2025). Analysis of credit ABS based on Markov chain approaches. Finance Research Letters, 71, 106432. https: //doi.org/10.1016/j.frl.2024.106432

[8] Stokes, J. R. (2024). Infinitesimal generator estimation with an application to credit risk (Working Paper, Summer 2024). University of Nebraska-Lincoln.

[9] Tang, T., Astfalck, L., & Dunson, D. (2025). Efficient Bayesian inference for discretely observed continuous time Markov chains (arXiv preprint arXiv: 2507.16756). https: //doi.org/10.48550/arXiv.2507.16756

[10] Gzyl, H., & Mayoral, S. (2025). Determining a credit transition matrix from cumulative default probabilities: An entropy minimization approach (arXiv preprint arXiv: 2503.14646). https: //doi.org/10.48550/arXiv.2503.14646

[11] Le, R., Ku, H., & Jun, D. (2021). Sequence-based clustering applied to long-term credit risk assessment. Expert Systems With Applications, 165, 113940. https: //doi.org/10.1016/j.eswa.2020.113940

[12] Pourbijan, F., Shams, M. F., Gholami-Jamkarni, R., Kordlouie, H., & Jahangarnia, H. (2022). Credit rating and preparing risk transition matrix for legal clients of banks with Markov chain approach. International Journal of Finance and Managerial Accounting, 7(27), 99–112.

[13] Pan, X., Chen, H., Shen, A., Zhao, D., & Su, X. (2024). A reliability assessment method for complex systems based on non-homogeneous Markov processes. Sensors, 24(11), 3446. https: //doi.org/10.3390/s24113446