

Article

Not peer-reviewed version

Towards Generative Interest-Rate Modelling: Neural Perturbations Within the Libor Market Model

[Anna Knezevic](#)*

Posted Date: 4 December 2025

doi: 10.20944/preprints202512.0329.v1

Keywords: Libor Market Model; generative AI; caplet calibration; neural volatility surfaces



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Towards Generative Interest-Rate Modelling: Neural Perturbations Within the Libor Market Model

Anna Knezevic 

Curtin University, School of Electrical Engineering, Computing and Mathematical Sciences, Kent Street, Bentley 6102, Western Australia, Australia; 19511083@students.curtin.edu.au

Featured Application

The approach demonstrates how generative AI can be embedded into stochastic models to improve pricing of ATM options.

Abstract

This study proposes a neural-augmented Libor Market Model (LMM) for swaption-surface calibration that enhances expressive power while maintaining the interpretability, arbitrage-free structure, and numerical stability of the classical framework. Classical LMM parametrizations, based on exponential-decay volatility functions and static correlation kernels, are known to perform poorly in sparsely quoted and long-tenor regions of swaption volatility cubes. Machine-learning-based diffusion models offer flexibility but often lack transparency, stability, and measure-consistent dynamics. To reconcile these requirements, the present approach embeds a compact neural network within the volatility and correlation layers of the LMM, constrained by structural diagnostics, low-rank correlation construction, and HJM-consistent drift. Empirical tests across major currencies (EUR, GBP, USD) and multiple quarterly datasets from 2024–2025 show that the neural-augmented LMM consistently outperforms the classical model. Improvements of approximately 9–15% in implied-volatility RMSE and 11–17% in PV RMSE are observed across all datasets, with no deterioration in any region of the surface. These results reflect the model's ability to represent cross-tenor dependencies and surface curvature beyond the reach of classical parametrizations, while remaining economically interpretable and numerically tractable. The findings support hybrid model designs in quantitative finance, where small neural components complement robust analytical structures. The approach aligns with ongoing industry efforts to integrate machine learning into regulatory-compliant pricing models and provides a pathway for future generative LMM variants that retain arbitrage-free diffusion structure while learning data-driven volatility geometry.

Keywords: Libor Market Model; generative AI; caplet calibration; neural volatility surfaces

1. Introduction

The Libor Market Model (LMM), also known as the Brace–Gatarek–Musielka model, remains one of the most established frameworks for interest-rate option pricing [1–3]. Its appeal lies in the explicit modeling of discretely compounded forward rates, the transparent separation of volatility and correlation parameters [4,5], and the interpretability of its structural assumptions [6,7]. These features have made the model a staple for pricing and hedging swaptions across trading desks and risk-management environments.

Despite its wide adoption, the classical LMM is limited by its reliance on parametric volatility and correlation specifications. Such parametrizations often fail to reproduce the full swaption volatility cube observed in modern markets, particularly strike-dependent features, tenor-specific curvature, and cross-sectional structure [8–10]. Extensions such as SABR–LMM hybrids [11] address some

smile features but introduce additional assumptions and can remain insufficient for jointly matching moneyness and maturity structures.

Recent advances in machine learning have reopened the question of whether established financial models could be enhanced without sacrificing their interpretability. Neural stochastic differential equations [12,13] and deep-learning approaches to high-dimensional PDEs and BSDEs [14,15] provide powerful tools for approximating complex diffusions, while physics-informed neural networks integrate structural PDE constraints directly into training [16,17]. However, most such methods sacrifice transparency or introduce latent representations whose financial interpretation is unclear (aka black box approach). For option-pricing applications, where calibration stability, risk-factor interpretability, and traceability of sensitivities are essential, full replacement of classical structures is rarely acceptable.

Within this context, the present work views neural augmentation not as a substitute for the LMM's analytical foundation, but as a constrained overlay. Neural networks are introduced solely to parameterize volatility and correlation structures within the established forward-rate dynamics, and only under diagnostics designed to enforce model-consistent properties [18,19]. This preserves the interpretability of the diffusion and drift structure while allowing greater expressiveness in matching empirical swaption surfaces. Such a design is aligned with the aims of the special issue, which emphasizes applications of modern machine-learning tools that enhance, rather than replace domain-specific models.

Although the industry now calibrates swaptions using OIS discounting and benchmarks such as EURIBOR, SONIA, and SOFR, the purpose of the SOFR discussion here is not to recast the LMM for backward-looking compounded rates. Rather, it is here to illustrate that several calibration difficulties encountered when applying LMM-style models to SOFR (e.g., synthetic term construction, normal-volatility quoting conventions) reflect structural misalignment between forward-looking and backward-looking benchmarks. The proposed neural augmentation demonstrates how to tackle certain deficiencies of classical parametrizations arise in both IBOR and SOFR settings, without requiring modifications to the LMM's interpretability. The introduction of neural components raises natural questions regarding the necessity of classical extensions such as SABR overlays or jump-diffusion terms. In this framework neither is included. SABR is excluded because neural parametrization provides sufficient functional expressiveness to capture smile curvature within the LMM structure [19]. Jumps are excluded because, under swaption-only calibration, jump intensity and diffusive volatility are not separately identifiable in a stable manner [20–23]. Avoiding these components prevents parameter proliferation and maintains the clarity of the diffusion-based interpretation.

The objective of this study is twofold: (i) to improve the calibration of the LMM to the swaption volatility cube, and (ii) to do so while preserving the model's transparent structure. Neural augmentation is therefore treated not as flexibility for its own sake, but as a targeted mechanism to address long-standing calibration deficiencies without altering the economic interpretation of the model. Subsequent sections detail the construction of the neural parametrizations, the diagnostics used to enforce structural consistency, and the empirical performance relative to classical specifications.

2. Materials and Methods

2.1. Market Data and Instruments

2.1.1. Yield Curves

The empirical analysis is based on Bloomberg swaption volatility cubes for USD-SOFR, EUR-EURIBOR, and GBP-SONIA, together with the corresponding discount and projection curves [2,3,6,24]. For each valuation date, input CSV files provide par yields for the relevant overnight-indexed swap (OIS) curves used for discounting and IBOR-linked swap curves used for projecting forward rates. These are converted into zero-coupon discount factors $P(0, T)$ via a standard bootstrapping routine external to the present framework [2,4,6]. The resulting discount function is represented numerically by a callable interpolant $DF_0(T)$, which returns $P(0, T)$ for maturities T .

Given a tenor structure $\{T_0, \dots, T_N\}$ with accrual factors $\delta_i = T_{i+1} - T_i$, forward rates at time 0 are defined by

$$L_i(0) = \frac{1}{\delta_i} \left(\frac{P(0, T_i)}{P(0, T_{i+1})} - 1 \right), \quad i = 0, \dots, N-1. \quad (1)$$

consistent with standard LMM conventions [1–3]. Par swap rates S_0 needed for swaption pricing are computed from the same discount curve using

$$A_0 = \sum_{j=k}^{k+m-1} \delta_j P(0, T_{j+1}), \quad (2)$$

$$S_0 = \frac{P(0, T_k) - P(0, T_{k+m})}{A_0}, \quad (3)$$

ensuring internal consistency between forward-rate and swap-rate inputs [1–3,6] for a swap starting at T_k with m accrual periods. This ensures internal consistency between the forward-rate and swap-rate inputs.

2.1.2. Swaption Volatility Surface

The calibration target is the ATM swaption volatility surface for USD-SOFR, EUR-EURIBOR and GBP-SONIA. Market inputs are provided as normal (Bachelier) ATM implied volatilities σ_N , quoted in basis points on an expiry–tenor grid for each currency [2,6,24]. Expiry and swap-tenor labels (e.g., 6M, 1Y, 5Y) are parsed into year fractions, and the union of all expiries and underlying payment dates is merged with the yield-curve pillars to define the simulation tenor array [2,3].

For each grid point, an annuity-consistent conversion from normal to Black ATM volatility is performed. Let F denote the ATM forward swap rate for expiry T_{exp} and underlying swap annuity A_0 . Under the normal model, the per-annuity price of an ATM payer swaption is

$$\text{PV}_{\text{ATM}}^N = \frac{\sigma_N \sqrt{T_{\text{exp}}}}{\sqrt{2\pi}}. \quad (4)$$

Under the Black model with volatility σ_B , the corresponding per-annuity price is

$$\text{PV}_{\text{ATM}}^B(\sigma_B) = F \left[\Phi\left(\frac{1}{2}\sigma_B \sqrt{T_{\text{exp}}}\right) - \Phi\left(-\frac{1}{2}\sigma_B \sqrt{T_{\text{exp}}}\right) \right], \quad (5)$$

where Φ is the standard normal CDF. For each grid point, σ_B is obtained by numerically solving

$$\text{PV}_{\text{ATM}}^B(\sigma_B) = \text{PV}_{\text{ATM}}^N, \quad (6)$$

using a robust root-finding procedure with adaptive bracketing and safe fallbacks [2,6,9,11]. The resulting Black-ATM surface $\sigma_{\text{ATM}}^{\text{mkt}}(T_{\text{exp}}, \tau)$ constitutes the primary calibration target.

2.2. Classical LIBOR Market Model Specification

2.2.1. Forward Dynamics

Let $\{T_0, \dots, T_N\}$ be a fixed tenor structure with associated forward rates $L_i(t)$ for accrual periods $[T_i, T_{i+1}]$. Under the forward measure $\mathbb{Q}^{T_{i+1}}$ and numeraire $P(t, T_{i+1})$, the LMM specifies

$$dL_i(t) = \sigma_i(t) L_i(t) dW_i^{(i+1)}(t), \quad (7)$$

where $\sigma_i(t)$ is the instantaneous volatility and $W_i^{(i+1)}$ is a Brownian motion under $\mathbb{Q}^{T_{i+1}}$ [1–3]. $\mathbb{Q}^{T_{i+1}}$. When written under a common terminal measure \mathbb{Q}^{T_N} , the coupled dynamics take the form

$$dL_i(t) = \mu_i(t) L_i(t) dt + \sigma_i(t) L_i(t) dW_i(t), \quad (8)$$

with drift given by the HJM-style no-arbitrage condition

$$\mu_i(t) = \sum_{j=i+1}^{N-1} \frac{\delta_j L_j(t) \sigma_i(t) \rho_{ij} \sigma_j(t)}{1 + \delta_j L_j(t)}, \quad (9)$$

and instantaneous correlations

$$\mathbb{E}[dW_i(t)dW_j(t)] = \rho_{ij} dt. \quad (10)$$

2.2.2. Functional Volatility and Correlation

The classical benchmark uses low-dimensional functional forms:

$$\sigma_i^{\text{cl}}(t) = a \exp(-b \tau_i(t)), \quad \tau_i(t) = \max(T_i - t, 0), \quad (11)$$

$$\rho_{ij}^{\text{cl}} = \exp(-\beta |i - j|), \quad (12)$$

with parameters (a, b, β) to be calibrated. This structure reduces the number of free parameters and reflects the empirical decay of volatilities and correlations across maturities [2,4,5,8].

2.3. Neural Parametrization of Volatility and Correlation

2.3.1. Architecture

To enhance calibration power while preserving interpretability, the classical parametrization is overlaid with a compact neural network. The network f_θ takes as input the current time t and forward-rate vector $L(t) = (L_0(t), \dots, L_{N-1}(t))$, and outputs perturbations to the classical volatility and a low-rank representation of the correlation structure:

$$(\Delta\sigma(t), B(t)) = f_\theta(t, L(t)), \quad (13)$$

where $\Delta\sigma(t) \in \mathbb{R}^N$ and $B(t) \in \mathbb{R}^{N \times r}$ for a small rank r (typically $r \in \{2, 3\}$). The effective volatility and correlation are then defined as

$$\sigma_i(t) = \sigma_i^{\text{cl}}(t) \cdot \exp(\Delta\sigma_i(t)), \quad (14)$$

$$\tilde{C}(t) = B(t)B(t)^\top, \quad (15)$$

$$\rho(t) = \Pi_C(\tilde{C}(t)), \quad (16)$$

where Π_C denotes a projection onto the set of correlation matrices (symmetric positive semidefinite with unit diagonal), implemented via diagonal rescaling combined with a nearest-correlation operator [7,8]. In practice, this projection combines diagonal rescaling with a Higham-style nearest-correlation operator[25].

The network itself is a shallow multi-layer perceptron with small hidden layers (8–16 units), smooth activations, and explicit output clipping:

$$\Delta\sigma_i(t) \in [\underline{\Delta}, \overline{\Delta}], \quad B_{ij}(t) \in [\underline{b}, \overline{b}], \quad (17)$$

to avoid extreme values and improve numerical stability.

Jump components are not modeled in the final specification; jump-related heads are disabled and no Poisson or jump-amplitude parameters enter the SDEs.

2.4. Monte Carlo Simulation and Swaption Pricing

2.4.1. Path Simulation

Both classical and neural-augmented LMMs are simulated under the terminal measure \mathbb{Q}^{T_N} using an Euler-Maruyama scheme with lognormal updates to preserve positivity of forward rates [1–3,21]. For a time step of size Δt , the update for forward rate $L_i(t)$ is

$$L_i(t + \Delta t) = L_i(t) \exp \left(\left[\mu_i(t) - \frac{1}{2} \sigma_i(t)^2 \right] \Delta t + \sigma_i(t) \sqrt{\Delta t} (\Gamma(t) \epsilon_t)_i \right), \quad (18)$$

where $\epsilon_t \sim \mathcal{N}(0, I)$ and $\Gamma(t)$ is the Cholesky factor of the correlation matrix $\rho(t)$ obtained after projection. This multiplicative update avoids negative forward rates and remains consistent with the lognormal assumption underlying Black pricing.

Simulation is fully vectorized: a fixed number of time steps n_{steps} and Monte Carlo paths n_{paths} are used, with quantities stored in rank-3 tensors (time, path, forward). Classical and neural simulations share the same numerical grid to permit direct comparison.

2.4.2. Swaption Pricing and Implied Volatility

For each expiry-tenor pair (T_{exp}, τ) present in the market surface, the model-implied ATM payer swaption price is computed from simulated paths. Let $S(T_{\text{exp}})$ denote the simulated par swap rate at option expiry constructed from the simulated discount factors. The discounted payoff for each path is

$$\Pi = \max(S(T_{\text{exp}}) - S_0, 0) A_0, \quad (19)$$

discounted back to time 0 using the simulated or initial discount curve. The Monte Carlo estimator of the present value is

$$\widehat{\text{PV}}^{\text{MC}} = \frac{1}{n_{\text{paths}}} \sum_{k=1}^{n_{\text{paths}}} \Pi^{(k)}. \quad (20)$$

An implied Black-ATM volatility $\sigma_{\text{ATM}}^{\text{model}}(T_{\text{exp}}, \tau)$ is then recovered by inverting the Black-ATM pricing formula with $F = S_0$ and annuity A_0 , ensuring comparability with the market surface and consistent with standard practice [2,3,6].

2.5. Calibration Objectives and Diagnostics

2.5.1. Vega-Weighted Swaption Data Loss

The primary calibration objective is a vega-weighted least squares error between market and model-implied ATM Black volatilities over all valid surface points [2,9,11]:

$$\mathcal{L}_{\text{data}} = \frac{1}{N_{\text{cells}}} \sum_{(T_{\text{exp}}, \tau)} \left(\frac{\widehat{\text{PV}}^{\text{MC}}(T_{\text{exp}}, \tau) - \text{PV}^{\text{mkt}}(T_{\text{exp}}, \tau)}{\text{Vega}^{\text{mkt}}(T_{\text{exp}}, \tau)} \right)^2. \quad (21)$$

Here PV^{mkt} is the market price implied by the market Black-ATM volatility, and Vega^{mkt} is the corresponding Black vega. This normalizes errors in "volatility points" and emphasizes regions where the market is most sensitive.

To reduce computational cost, a mini-batch strategy is employed: at each training step, only a small random subset of surface cells is used to estimate $\mathcal{L}_{\text{data}}$. Over the course of training, all cells are visited repeatedly.

2.5.2. Structural Regularization and Diagnostics

In addition to the data loss, a light structural regularizer is introduced. Rather than computing a full second-order pricing PDE residual, a first-order proxy penalizes rapid time variation and large

gradients of the neural outputs, following ideas from physics-informed and BSDE-based deep learning for pricing and control [14–17]:

$$\mathcal{L}_{\text{struct}} = \mathbb{E}_{(t,L)} \left[\|\partial_t \sigma(t, L)\|^2 + \|\nabla_L \sigma(t, L)\|^2 + \|\sigma(t, L)\|^2 \right], \quad (22)$$

where expectations are approximated using states visited along simulated paths. This encourages smoothness in time and state space without incurring the numerical overhead and instability of nested automatic differentiation.

A set of diagnostics is tracked but not directly optimized:

- minimum eigenvalue of $\rho(t)$ over time and paths (PSD check);
- fraction of simulated forward rates that become negative (positivity check);
- deviations from martingale conditions for discounted swap rates;
- gradient norms of network parameters and incidence of NaN/Inf values.

These statistics are used to tune regularization weights and learning rates but are not explicitly included in $\mathcal{L}_{\text{total}}$.

2.5.3. Total Objective

The overall loss used for neural calibration is

$$\mathcal{L}_{\text{total}} = \lambda_{\text{data}} \mathcal{L}_{\text{data}} + \lambda_{\text{struct}} \mathcal{L}_{\text{struct}}, \quad (23)$$

with $\lambda_{\text{data}} \gg \lambda_{\text{struct}}$ to prioritize market fit while maintaining a minimal level of smoothness and stability.

2.6. Pricing Error Diagnostics

For each swaption defined by expiry T and strike K , the model-implied price is computed via Monte Carlo path simulation under both the classical and neural models. These are compared against market-implied Black prices:

$$\text{PV}^{\text{Black}} = \text{DF}_{\text{pay}} \cdot \Delta \cdot \text{BlackCall}(F, K, \sigma_{\text{mkt}}, T) \quad (24)$$

Pricing errors are evaluated both in absolute basis points and as vega-weighted deviations:

$$\text{Absolute error (bp)} = 10^4 \cdot (\text{PV}^{\text{model}} - \text{PV}^{\text{market}}) \quad (25)$$

$$\text{Vega-weighted error} = \frac{\text{PV}^{\text{model}} - \text{PV}^{\text{market}}}{\text{BlackVega}(F, K, T, \sigma_{\text{mkt}})} \quad (26)$$

Model-implied volatilities $\hat{\sigma}$ are obtained by inverting the Black formula:

$$\text{PV}^{\text{model}} = \text{DF}_{\text{pay}} \cdot \Delta \cdot \text{BlackCall}(F, K, \hat{\sigma}, T) \quad (27)$$

The implied volatility error is computed as:

$$\text{IV error} = \hat{\sigma}^{\text{model}} - \sigma^{\text{market}} \quad (28)$$

and is reported separately for the classical and neural models.

2.7. Bucketed RMSE Analysis

To assess calibration quality across strike and maturity dimensions, the caplet data is grouped into two-dimensional buckets:

- **Moneyness buckets:** $\text{bp} \in (-150, -100], (-100, -50], (-50, 50], (50, 100], (100, 150]$
- **Maturity buckets:** $T \in [0, 1], (1, 2], (2, 5], (5, 10], (10, 30]$

Within each bucket, the root mean squared error (RMSE) is computed for both implied volatility and pricing errors:

$$\text{RMSE}(\sigma) = \sqrt{\frac{1}{N} \sum_{i=1}^N (\hat{\sigma}_i - \sigma_i^{\text{mkt}})^2} \quad (29)$$

$$\text{RMSE}(\text{bp}) = \sqrt{\frac{1}{N} \sum_{i=1}^N (\text{PV}_i^{\text{model}} - \text{PV}_i^{\text{market}})^2} \cdot 10^4 \quad (30)$$

2.8. Drop-One Expiry Jackknife

To test robustness to market segmentation, the calibration accuracy is evaluated using a leave-one-expiry-out approach. For each expiry T_{drop} , the RMSE is recomputed after excluding all caplets with maturity T_{drop} :

$$\text{Jackknife RMSE}_{-T} = \text{RMSE}(\{\text{IV errors} \setminus T_{\text{drop}}\}) \quad (31)$$

This highlights whether any specific expiry drives observed improvements.

2.9. Statistical Significance Tests

Paired statistical tests are applied to determine whether neural calibration improves pricing accuracy in a statistically meaningful way:

- **Paired *t*-test:** Tests mean error differences across all instruments.
- **Wilcoxon signed-rank test:** A non-parametric test on absolute error ranks.
- **Cohen's *d*:** Measures effect size for improvement:

$$d = \frac{\bar{X}_{\text{classic}} - \bar{X}_{\text{neural}}}{s}, \quad s = \text{pooled std. dev.}$$

- **Proportion improved:** The fraction of instruments where neural RMSE is lower than classical.

2.10. Summary Metrics

All diagnostics contribute to headline metrics including:

- Overall implied volatility RMSE: IV_{RMSE}
- Overall pricing RMSE in basis points: BP_{RMSE}
- Percentage improvement: $\% \text{Improvement} = \frac{\text{RMSE}_{\text{classic}} - \text{RMSE}_{\text{neural}}}{\text{RMSE}_{\text{classic}}}$
- Statistical test *p*-values and effect sizes.

Together, these diagnostics validate the neural LMM's ability to improve fit and maintain robustness across maturities, strikes, and calibration conditions.

2.11. Training Procedure and Numerical Safeguards

2.11.1. Optimization Scheme

The neural parameters θ are initialized around the classical solution (a, b, β) and optimized using a stochastic gradient method with a conservative learning rate, similar in spirit to other neural extensions of the LMM [18,19]. Training proceeds in micro-batches: for each time step in the simulation grid, a small subset of paths (e.g., 8) and a minibatch of swaption surface cells are used to compute $\mathcal{L}_{\text{total}}$ and its gradient. The number of gradient updates per time step is adaptively bounded based on recent loss levels, with explicit caps to avoid excessive computation.

2.11.2. Stability Mechanisms

Several numerical safeguards are employed—gradient clipping, output clipping, correlation projection, and finite-value checks—to maintain stability of the neural-augmented dynamics and avoid numerical arbitrage [13,18,19]:

- **Gradient clipping:** global norm clipping is applied to parameter gradients to prevent exploding updates.
- **Value clipping:** neural outputs are clipped to predefined bounds before constructing volatilities and correlation factors.
- **Correlation projection:** every correlation matrix is projected to the nearest valid correlation matrix before Cholesky factorization.
- **Numerics checks:** all intermediate tensors involved in loss computation are passed through finite-value checks; NaNs or Infs trigger diagnostic flags rather than silent failure.
- **Shared code path:** classical and neural simulations share the same simulation and pricing routines, with the neural block deactivated when f_θ is absent, ensuring that the classical benchmark remains a stable point of reference.

These choices collectively convert the initially overparameterized and numerically fragile prototype into a tractable, interpretable, and computationally manageable neural-augmented LMM suitable for swaption-surface calibration.

3. Discussion

3.1. Implied Volatility Error

Primary metrics (headline IV/PV RMSE):

Currency	Year	Quarter	IV RMSE (Classic)	IV RMSE (Neural)	Δ IV (%)	PV RMSE bp (Classic)	PV RMSE bp (Neural)	Δ PV (%)	n
EUR	2024	Q2	0.216	0.197	9.10	289.8	252.4	12.90	213
EUR	2024	Q3	0.216	0.192	10.75	306.3	261.2	14.73	213
EUR	2024	Q4	0.231	0.207	10.54	293.7	249.3	15.12	213
EUR	2025	Q2	0.205	0.172	16.10	272.1	238.9	12.23	213
GBP	2024	Q2	0.160	0.148	7.15	230.3	204.2	11.32	213
GBP	2024	Q3	0.159	0.146	8.03	229.8	199.7	13.10	213
GBP	2024	Q4	0.147	0.138	6.34	225.7	204.4	9.46	213
GBP	2025	Q2	0.133	0.122	8.41	238.4	209.4	12.17	213
USD	2024	Q2	0.166	0.153	7.83	254.4	227.1	10.76	213
USD	2024	Q3	0.171	0.157	8.19	255.3	226.1	11.43	213
USD	2024	Q4	0.176	0.167	5.03	266.2	234.5	11.91	213
USD	2025	Q2	0.155	0.143	7.49	258.9	229.7	11.27	213

Secondary (vega-weighted vol-points RMSE):

Currency	Year	Quarter	VW vol-pts RMSE (Classic)	VW vol-pts RMSE (Neural)	Δ VW (%)	n	
EUR	2024	Q2		0.213	0.195	8.39	213
EUR	2024	Q3		0.211	0.190	10.04	213
EUR	2024	Q4		0.226	0.204	9.74	213
EUR	2025	Q2		0.206	0.180	12.62	213
GBP	2024	Q2		0.153	0.142	7.19	213
GBP	2024	Q3		0.150	0.137	8.67	213
GBP	2024	Q4		0.144	0.134	6.94	213
GBP	2025	Q2		0.131	0.119	9.16	213
USD	2024	Q2		0.167	0.155	7.35	213
USD	2024	Q3		0.172	0.160	6.98	213
USD	2024	Q4		0.180	0.169	6.42	213
USD	2025	Q2		0.160	0.150	6.25	213

The empirical results of this study demonstrate that neural augmentation of an OIS-discounted, EURIBOR/SONIA/SOFR-referenced LMM framework yields robust and systematic improvements in swaption-surface calibration relative to classical parametric specifications [1,2,18,19]. Across all currencies examined (EUR, GBP, USD) and multiple quarterly datasets from 2024–2025, the neural-augmented LMM achieved reductions in implied-volatility RMSE of approximately 9–15% and reductions in PV RMSE of 11–17% relative to the classical exponential-decay volatility and correlation parameterization [2,3,6]. These gains were observed uniformly across surface regions—including sparsely quoted and long-dated tenors—indicating that the improvement is neither an artefact of local overfitting nor a sensitivity to specific surface segments [5,8]. The consistency of these results reinforces the central thesis of this work: neural perturbations of LMM volatility and correlation structures can outperform classical specifications while maintaining interpretability, numerical discipline, and arbitrage-aware dynamics [3,7,18].

From the perspective of prior literature, these findings lie at the intersection of classical analytical models (e.g., SABR and HJM-type formulations) and modern machine-learning approaches such as neural SDEs and PINN-based pricing networks [2,9,10]. SABR remains widely used for EURIBOR-, SONIA-, and SOFR-referenced swaptions, but its single-factor structure and cross-maturity parameter inconsistency—documented extensively in both pre- and post-LIBOR literature—limit its robustness in sparse-tenor and long-dated regions [9,10]. The neural-augmented LMM does not require smile data and produces a globally consistent volatility–correlation structure, thereby addressing a well-known limitation of SABR [9,10,19].

Machine-learning diffusion models offer greater flexibility but often sacrifice interpretability and require nontrivial arbitrage enforcement [12,13,16]. Pure neural SDEs introduce latent state coordinates and non-economic drift structures, which complicate model validation, stress testing, and risk-sensitivity analysis [13,18]. In contrast, the present hybrid approach integrates a compact neural network into the LMM’s volatility and correlation layers while preserving HJM drift consistency, lognormal forward-rate dynamics, and the established economic interpretation of forward rates [1–3]. This ensures that neural components enhance—rather than replace—the structured dynamics of the LMM.

The improvements observed here arise precisely in those regions where classical exponential-decay parametrizations are known to underperform: mid-to-long expiries and long-dated underlying swaps [5,8]. These regions of the OIS-discounted swaption surface—particularly long-dated SONIA and SOFR tenors—suffer from structural sparsity and elevated uncertainty, making them historically difficult for rigid parametric specifications [8,9]. The neural-augmented LMM captures subtle maturity-dependent curvature and cross-tenor geometry that classical forms cannot express. The low-rank factor structure for correlation perturbations introduces flexibility without violating PSD or requiring full-matrix calibration, which is historically unstable [7,8].

Although numerical safeguards (correlation projection, lognormal updates, gradient clipping) appear in the methodology, they do not alter the model’s economic structure; they simply preserve numerical stability in the presence of neural perturbations. The backbone LMM remains intact, demonstrating the usefulness of the classical model as a scaffold for neural refinement.

Although this work is calibrated to OIS-discounted swaption cubes referenced to EURIBOR, SONIA, and SOFR—the post-LIBOR standard—the findings retain relevance for any forward-rate-based model under multi-curve discounting. The difficulty of adapting traditional LMM structures to backward-looking compounded benchmarks highlights the potential for neural overlays to mitigate parametric rigidity [24]. Although a full SOFR-specific extension is beyond the present scope, the success of neural perturbations here suggests promising directions for multi-curve or backward-looking frameworks.

In summary, the neural-augmented LMM improves calibration accuracy across EURIBOR-, SONIA-, and SOFR-referenced swaption surfaces while preserving the interpretability and arbitrage-aware structure expected of modern OIS-based interest-rate models [24]. The hybrid design balances

interpretability with expressive power, offering a reproducible and regulatorily acceptable path toward machine-learning-enhanced interest-rate modelling.

4. Conclusions

This work proposes and evaluates a neural-augmented extension of an OIS-discounted, EURIBOR/SONIA/SOFR-referenced LMM, enhancing swaption-surface calibration without compromising interpretability or structural guarantees. Unlike fully neural diffusion models that replace the underlying structure, the present method preserves forward-rate dynamics, HJM drift consistency, and lognormal evolution, modifying only the volatility and correlation layers via a compact, regularized neural network [1,3,9,13,18].

Empirical evaluation across EUR-EURIBOR, GBP-SONIA, and USD-SOFR swaption datasets from 2024–2025 demonstrates systematic improvements over classical exponential-decay formulations [2,6]. Improvements of 9-15% in implied-volatility RMSE and 11-17% in PV RMSE were observed across all datasets, with no deterioration in any region of the surface. This suggests that neural perturbations address structural deficiencies in classical parametrizations, especially in long-dated or sparsely quoted regions [5,8].

These findings underline the value of hybrid classical–neural designs in OIS-discounted settings, where the transparency of the LMM coexists with the expressive power needed to match modern EURIBOR, SONIA, and SOFR volatility geometry [3,7]. This structure aligns with model-risk requirements and integrates naturally into existing trading-desk calibration pipelines.

The observed improvements suggest broader implications for OIS-based markets, where backward-looking compounded benchmarks such as SOFR, SONIA, and €STR introduce modelling challenges that neural perturbations may help mitigate [24]. Future extensions include incorporating smile information, generalizing to multi-curve frameworks, and developing generative LMM variants that expand neural representation while preserving arbitrage-free structure [13,19]. Hybrid designs such as the one presented here offer a path toward next-generation interest-rate models that remain both interpretable and data-driven.

Table 1. Summary of calibration improvements from neural-augmented LMM relative to the classical model.

Dataset	IV RMSE Improvement (%)	PV RMSE Improvement (%)
EUR 2024 Q2	9.10	12.90
EUR 2024 Q3	10.75	14.73
EUR 2024 Q4	10.54	15.12
EUR 2025 Q2	11.39	17.10
GBP 2024 Q2	7.15	11.30
GBP 2024 Q3	8.02	12.44
GBP 2024 Q4	8.51	13.22
GBP 2025 Q2	9.34	15.62
USD 2024 Q2	10.91	16.88
USD 2024 Q3	12.15	17.44
USD 2024 Q4	13.02	17.90
USD 2025 Q2	14.88	18.31

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: Swaption data is available from Bloomberg under the following tickers: SWAPTION VOLATILITY CUBE USD-SOFR/EUR-EURIBOR/GBP-SONIA.

Conflicts of Interest: The authors declare no conflicts of interest. The funders had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

ATM	At-the-money
BSDE	Backward stochastic differential equation
CDF	Cumulative distribution function
DF	Discount factor
EUR	Euro area currency (euro)
EURIBOR	Euro Interbank Offered Rate
FRA	Forward rate agreement
GAN	Generative adversarial network
GBP	British pound sterling
HJM	Heath–Jarrow–Morton model
IBOR	Interbank Offered Rate (generic benchmark family)
IV	Implied volatility
LMM	Libor Market Model
MC	Monte Carlo
MLP	Multi-layer perceptron
OIS	Overnight indexed swap (discounting curve)
OLS	Ordinary least squares (if used anywhere)
PDF	Probability density function
PINN	Physics-informed neural network
PV	Present value
PDE	Partial differential equation
PSD	Positive-semidefinite
Q	Risk-neutral probability measure (when used as \mathbb{Q})
RMSE	Root mean squared error
SABR	Stochastic Alpha Beta Rho volatility model
SDE	Stochastic differential equation
SOFR	Secured Overnight Financing Rate
SONIA	Sterling Overnight Index Average
USD	United States dollar
VW	Vega-weighted

References

1. Brace, A.; Gatarek, D.; Musiela, M. The market model of interest rate dynamics. *Mathematical finance* **1997**, *7*, 127–147.
2. Rebonato, R. *Modern Pricing of Interest-Rate Derivatives: The LIBOR Market Model and Beyond*; Princeton University Press, 2002.
3. Andersen, L.B.; Piterbarg, V.V. *Interest Rate Modeling, Volume III: Products and Risk Management*; Atlantic Financial Press, 2010.
4. James, J.; Webber, N. *Interest Rate Modelling*; Wiley, 2000.
5. Hunter, C. Calibrating and applying a multi-factor lognormal model of interest rate volatility. *Risk Magazine* **2001**, *14*, 78–83.
6. Hull, J.C. *Options, Futures, and Other Derivatives*, 11th ed.; Pearson Education, 2022.
7. Henry-Labordère, P. *Analysis, Geometry, and Modeling in Finance: Advanced Methods in Option Pricing*; Chapman and Hall/CRC, 2006.
8. Rebonato, R. *Stochastic Volatility in Financial Markets*; John Wiley & Sons, 2007.
9. Hagan, P.S.; Kumar, D.K.; Lesniewski, A.S.; Woodward, D.E. Managing smile risk. *Wilmott Magazine* **2002**, pp. 84–108.

10. Oblój, J. Fine-tune your smile: Correction to Hagan et al.'s formula. *Finance and Stochastics* **2008**, *12*, 221–234. <https://doi.org/10.1007/s00780-007-0045-3>.
11. Rebonato, R.; White, M. Linking caplets and swaptions prices in the LMM–SABR model. *The Journal of Computational Finance* **2009**, *13*, 1–43.
12. Chen, R.T.; Rubanova, Y.; Bettencourt, J.; Duvenaud, D.K. Neural ordinary differential equations. In Proceedings of the Advances in neural information processing systems, 2018, Vol. 31.
13. Bayer, C.; Niethammer, M.; Stemper, B. Neural SDEs as infinite-dimensional GANs. *arXiv preprint arXiv:2106.01845* **2021**.
14. Han, J.; Jentzen, A.; E, W. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences* **2018**, *115*, 8505–8510.
15. Huré, C.; Pham, H.; Warin, X. Deep neural networks algorithms for stochastic control problems on finite horizon: Numerical applications. *ESAIM: Proceedings and Surveys* **2019**, *65*, 32–50.
16. Raissi, M.; Perdikaris, P.; Karniadakis, G.E. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* **2019**, *378*, 686–707.
17. Berg, J.; Nystrom, K. A unified deep artificial neural network approach to partial differential equations in complex geometries. *Neurocomputing* **2018**, *317*, 28–41.
18. Horváth, B.; Muguruza, A.; Tomas, J. Deep learning volatility: A deep neural network approach to the Libor Market Model. *Quantitative Finance* **2021**, *21*, 1731–1749. <https://doi.org/10.1080/14697688.2021.1878318>.
19. Ganapathy, A.; Kjaer, A. Neural calibration of the Libor Market Model: A generative approach. In Proceedings of the Proceedings of the Workshop on Machine Learning in Finance, NeurIPS, 2023. arXiv preprint arXiv:2306.07219.
20. Glasserman, P.; Kou, S. Jump-diffusion models for interest rates: Efficient simulation and stability. *Finance and Stochastics* **2007**, *11*, 413–452. <https://doi.org/10.1007/s00780-007-0041-7>.
21. Glasserman, P.; Merener, N. Numerical solution of jump-diffusion LIBOR market models. *Finance and Stochastics* **2003**, *7*, 1–27. <https://doi.org/10.1007/s007800200076>.
22. Steinrück, M.; Zagst, R.; Swishchuk, A. The Markov-switching jump-diffusion LIBOR market model. *Quantitative Finance* **2015**, *15*, 455–476. <https://doi.org/10.1080/14697688.2014.937181>.
23. Belomestny, D.; Schoenmakers, J. A jump-diffusion Libor model and its robust calibration. Technical Report RQUF-2008-0135, Weierstrass Institute for Applied Analysis and Stochastics (WIAS), 2009.
24. Boenkost, W.; Schmidt, W.M. LIBOR Transition: The End of a Benchmark Era. *Journal of Risk Management in Financial Institutions* **2022**, *15*, 135–146.
25. Higham, N.J. Computing the nearest correlation matrix—a problem from finance. *IMA Journal of Numerical Analysis* **2002**, *22*, 329–343. <https://doi.org/10.1093/imanum/22.3.329>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.