CIRJE-F-1262

Generating the Term Structure of Interest Rates with Diffusion Models

Yosuke Fukunishi
Graduate School of Economics
The University of Tokyo

Haorong Qiu
Graduate School of Economics
The University of Tokyo

Akihiko Takahashi
The University of Tokyo

Fan Ye
Graduate School of Economics
The University of Tokyo

December 2025

Generating the Term Structure of Interest Rates with Diffusion Models

Yosuke Fukunishi[1], Haorong Qiu[1], Akihiko Takahashi[1], and Fan Ye[1]

[1]Graduate School of Economics, University of Tokyo

First version: September 16, 2025 Revised version:November 30, 2025

Abstract

This study introduces a novel generative modeling framework for simulating the term structure of interest rates. In recent years, generative models have achieved significant progress in image generation and are increasingly being applied to finance. To the best of our knowledge, this is the first study to apply a generative model—specifically, a diffusion model—to the term structure of interest rates. Furthermore, we extend the framework to incorporate conditional generation mechanisms and v-parameterization. The training dataset consists of spot yield curves constructed from daily overnight index swap (OIS) rates using cubic Hermite splines. As base conditioning variables, we use short-term interest rates and changes in consumer price indexes (CPIs). Empirical analysis covering the period from 2015 to 2025 demonstrates that our model successfully reproduces the level and shape of yield curves corresponding to historical macroeconomic conditions and short-term interest rate environments. Additionally, when incorporating further conditioning variables related to quantitative easing policies, monetary base, current account balances, and nominal gross domestic product (GDP), we find that the inclusion of quantitative easing indicator notably enhances the model's output relative to the base conditioning case. This suggests improved robustness and representational capacity under expanded conditioning. In consideration of practical applications, we further analyze the generation outcomes derived from difference-based learning, confirming that the performance of out-of-sample generation is comparable to that of direct generation. Moreover, we also examine an alternative approach based on factor models commonly used in finance and macroeconomics to estimate the functional form of yield curves. Specifically, we consider the Nelson–Siegel–Svensson (NSS) model and investigate the indirect generation of synthetic yield curves by producing the NSS model's latent factors. Compared to direct generation, this factor-based indirect method enables faster generation while still achieving comparable reproducibility in terms of both the level and the shape of the yield curves.

Keywords: Machine Learning, Generative Models, Diffusion Models, Term Structure of Interest Rates, Yield Curve, Financial Time Series

1 Introduction

Generative models have experienced rapid growth in recent years, particularly within corporate environments over the past one to two years. However, their application in data-intensive domains —such as investment decision-making and risk management—remains in its early stages, though it holds considerable promise for future development. Against this backdrop, research into synthetic data generation in finance using generative models has begun to gain momentum. In this study, we focus on generating interest rate data—a topic that has received limited attention in prior research —and propose a method for constructing synthetic term structures of interest rates using diffusion models.

The foundation for the development of diffusion models was laid by Ho et al.[1], catalyzing rapid advancements in the field. For the implementation of synthetic data generation in this study, we incorporate two key technical innovations in diffusion modeling: conditional generation with cross-attention mechanisms and v-parameterization. The former was introduced by Rombach et al.[3], while the latter was proposed by Salimans and Ho[4]. These techniques have been widely adopted in the primary application areas of diffusion models, such as image and video generation.

Prior studies on synthetic data generation in finance using diffusion models include Chen et al. [5], which targets multivariate time series, and Tanaka et al.[6], which transforms time series into images for conditional training. However, these studies primarily focus on stock price time series. To the best of our knowledge, no prior work has addressed the term structure of interest rates in the context of diffusion-based synthetic data generation. Moreover, among the studies in finance reviewed in this section, none appear to have explicitly implemented v-parameterization. Therefore, our work is also distinctive in its use of v-parameterization for synthetic data generation.

The remainder of this paper is organized as follows. Section 2 provides an overview of diffusion models and discusses their technical advancements. Section 3 introduces our proposed methodology for preparing yield curves as training data and conditional generation of the term structure of interest rates using diffusion models. In Section 4, we present the results of conditional generation based on conditioning data from specific historical periods, demonstrating that the inclusion of conditioning variables enables the model to produce yield curves that closely resemble historical scenarios. In Section 5 we present the results of out-of-sample generation, showing a comparison between the yield curves generated by a model trained on data limited to the end of 2023 and the actual yield curves observed after 2024. This demonstrates that the model is capable of generating data that retains the characteristics of yield curves even under conditions different from the training period, and that it has expressive power to capture future data fluctuations over a horizon of approximately six months.In consideration of practical applications, we further analyze the generation outcomes derived from difference-based learning, confirming that the performance of out-of-sample generation is comparable to that of direct generation. In Section 6, we present a method of generating synthetic yield curves in combination with factor models that are already widely used in macroeconomic analysis, specifically the Nelson-Siegel-Svensson model. As a result, it is shown that the generated synthetic yield curves can still reproduce the shape of the actual yield curves, while enabling faster training and generation. Finally, Section 7 concludes the paper by examining practical implications and potential applications, and by outlining directions for future research.

2 Diffusion Models

Diffusion models have emerged as a powerful class of generative frameworks, particularly effective in modeling complex data distributions through iterative denoising processes and their application to finance has also been proposed. The field of diffusion models has experienced significant growth since the publication of the Denoising Diffusion Probabilistic Model(DDPM) introduced by Ho et al.[1]. The following provides an overview of DDPM.

Let $\mathbf{x}_0 \in \mathbb{R}^d$ be a sample drawn from an unknown data distribution characterized by the probability density function $p(\mathbf{x})$. In the context of generative modeling, we aim to approximate $p(\mathbf{x})$ using a model $p_\theta(\mathbf{x})$, parameterized by θ and then generate new samples by sampling from $p_\theta(\mathbf{x})$.

Since DDPM has developed from diffusion models, it is defined as a latent variable model. In this context, let $\mathbf{x}_0$ represent sample data and assume latent variables $\mathbf{x}_1, \mathbf{x}_2, \ldots, \mathbf{x}_T$ of the same dimension as the sample data.

An interesting aspect of DDPM lies in its formulation of the forward process, a Markov chain with Gaussian noise as follows:

$$\{\mathbf{x}_t\}_{t=1}^{T}, \quad \mathbf{x}_t = \sqrt{1-\beta_t}\mathbf{x}_{t-1} + \sqrt{\beta_t}\epsilon, \quad \epsilon \sim \mathcal{N}(0, I), \tag{1}$$

where $\mathcal{N}(\mu, \Sigma)$ denotes a normal distribution with mean μ and variance-covariance matrix Σ. In this explanation, the variance schedule$\beta_1 \ldots \beta_t$ are treated as a hyperparameter and assumed to be set to known values, satisfying $0 < \beta_1 < \beta_2 < \cdots < \beta_T < 1$.

Furthermore, by defining forward process as equation (1), when T is sufficiently large (e.g., $T = 1000$), $\mathbf{x}_T$ can be regarded as pure noise (i.e., $\mathbf{x}_T \sim \mathcal{N}(0, I)$).

According to Ho et al.[1], it is shown that the reverse process which traces the forward process in reverse order to progressively remove noises, is also a Markov chain with Gaussian transition probability and that the reverse process can be expressed as follows:

$$\{\mathbf{x}_t\}_{t=1}^{T}, \quad \mathbf{x}_{t-1} = \mu_\theta(\mathbf{x}_t, t) + \sqrt{\beta_t}\epsilon, \quad \epsilon \sim \mathcal{N}(0, I). \tag{2}$$

Then, given equation (2), we estimate $\mu_\theta(\mathbf{x}_t, t)$. This estimation can be performed by a neural network architecture such as U-Net, enabling the learning of the reverse process. As explained by Ho et al.[1], instead of estimating $\mu_\theta(\mathbf{x}_t, t)$,the reverse process can also be learned with a model that predicts the added noise $\epsilon_\theta(\mathbf{x}_t, t)$ given $\mathbf{x}_t$ and timestep t. This approach to learning the noise is referred to as *noise prediction.*

When performing sampling, we assume the initial value for the reverse process $\mathbf{x}_T \sim \mathcal{N}(0, I)$ so that $\mathbf{x}_T$ can be easily generated as a random vector. Then, by following the learned reverse process to progressively remove the noise, a sample can be obtained. This is the core idea behind DDPM.

Since the introduction of DDPM, theoretical advancements in diffusion models have included the unification with score-based models and the formulation of both the forward and reverse processes using continuous-time stochastic differential equations (SDEs),introduced by Song et al.[2]. Within this continuous-time SDE framework, the forward process in DDPM is represented as a stochastic process following the Ornstein-Uhlenbeck (OU) process, while the reverse process is expressed through an SDE derived from the relationship between the Kolmogorov forward and backward equations.

On the implementation side, one major advancement is conditional generation using cross-attention mechanisms, introduced by Rombach et al.[3], which enables image generation from words by connecting diffusion models with language models. Furthermore, to generate larger images more efficiently, techniques such as DDIM and *v-parameterization* proposed by Salimans and Ho[4]—both extensions of noise prediction—have been proposed and are now widely adopted in modern image generation frameworks, including Stable Diffusion v2.

3 Methodology

This paper aims to perform conditional generation of synthetic term structure of interest rate using a diffusion model. The procedure is outlined as follows:

Step1: Obtain historical swap rates and construct yield curves using methods commonly applied in finance. Based on the constructed yield curves, we construct datasets of zero-rates at quarterly (0.25-year) intervals,which is applied for training data.

Step2: Obtain the macro economic indicators as condition data to be applied during conditional learning.

Step3: Prepare a diffusion model that supports conditional generation based on both training data and conditional data.

Step4: Conduct training and perform conditional generation.

3.1 Data Preparation

3.1.1 Construction of Yield Curve and Interest Rate Dataset

In this study, we assume the simultaneous learning of interest rate data across multiple currencies and define the following requirements for the training interest rate dataset. The target currencies are Japanese Yen (JPY), US Dollar (USD), and British Pound (GBP). Furthermore, the interest rate targeted for training is chosen to be the zero rate used in the valuation of financial instruments, rather than the market-traded rate. Based on these requirements, the dataset is constructed through the following steps:

1. Time series data of OIS rates for each currency are obtained from Bloomberg. For details on the obtained OIS rates, see Appendix A. We collect data from 2010 to 2025. Due to differences in data availability across currencies, the starting dates of the datasets vary. However, in order to maximize the amount of data used for training, we do not align the starting dates and instead included all available data.

2. Since the obtained OIS rates correspond to instruments with annual interest payments, yearly tenor data is required to construct zero rates. However, due to missing market data for certain tenors, standard cubic spline interpolation is applied to generate interpolated OIS rate data for each currency.

3. Zero rates are constructed from the interpolated OIS rates using the bootstrap method for each currency. Additionally, following Healy[8], cubic Hermite spline interpolation is applied to the resulting zero rates. This yielded 120 data points per business day, with maturities ranging from 0.25 to 30 years for each currency.

4. The data for each currency are merged along the tenor axis to form a single training dataset. As a result, a dataset with the shape (11,185, 120) is created. where 11,185 represents the number of observations and 120 corresponds to the number of data points per observation.

3.1.2 Conditional Inputs

We use the following data as conditional inputs for yield curve generation:

Table 1: Condition Data

Data	Description	Update Frequency
Currency type	JPY = 1, USD = 2, GBP = 3	NA
Short-term interest rate	Overnight rate(JPY,GBP) or 1-week OIS rate (USD)	per Business day
Macroeconomic indicators	Inflation rates,and other indicators described in Appendix B	per Month

For details on the obtained condition data, see Appendix B. Although macroeconomic indicators are updated on a monthly basis, we use the indicators displayed on the Bloomberg or LSEG Datasteram for a given business day directly to construct the dataset. Given that the dimension of the conditional features is X, we construct data with a shape of (11,185, X). Both short-term interest rate and macroeconomic indicators are publicly available and easily accessible, and are selected as conditioning variables ,which seems closely related to formation of the term structure of interest rates. The short-term interest rate is considered to have a direct impact on the short end of the yield curve. Among the components of macroeconomic indicators, inflation rates are regarded as influencing the determination of medium- to long-term interest rate levels and the shape of the yield curve.

3.2 Diffusion Model Setup

After preparation of the training data set, we conduct training using a diffusion model. For the implementation of the model, we refer to the standard diffusion model program utilized in the evaluation of a diffusion factor model in Chen et al.[5][7], and extend it by incorporating additional code to support conditional generation and the v parameterization technique for the estimation object. Although Chen et al.[5] proposes a model to learn latent factors, deploying neural networks defined in (18) of their article, we understand that the program code corresponds to a standard diffusion model implementation, using equation (7) in their paper as the objective function, as stated in Appendix D of Chen et al.[5]. The neural network architecture employed is the 2D U-Net framework, which is widely adopted in diffusion models, its architecture is as follows.
We reshape the data from a one-dimensional vector of length 120 into a two-dimensional matrix of size (10, 12). The resulting (10, 12) matrices are then processed using a 2D U-Net architecture consisting of two downsampling layers, one intermediate layer, and two upsampling layers. To implement conditional learning and sampling, we incorporate cross-attention layers into the

intermediate stages of the U-Net framework according to Rombach et al. [3]. This adaptation enables the model to effectively utilize macro-level data—which strongly influences the shape of yield curves.

For the training hyperparameters, we use the cosine schedule for the beta schedule. The number of training steps is set to 1000, following the DDPM approach.

3.3 Model Training and Synthetic Yield Curve Generation

We train the model for 200 epochs using both interest rate data and conditional macroeconomic inputs. The training is carried out on a commercial GPU (Geforce 5070) and completed in approximately 30 minutes, demonstrating its practical feasibility. Synthetic yield curves are generated conditionally based on macroeconomic indicators using the trained model. The input conditional data must match the dimensionality of the features used during training. For example, if the model is trained with currency, short-term interest rates, and inflation rates as conditioning variables, then the scenarios should be provided as a array of shapes (1, 3). Each sample is generated in 1,000 time steps. The generation of 1,024 synthetic samples per condition takes approximately 90 seconds on the same GPU.

4 Results

This section discusses the analysis of the outputs generated by the trained model. To analyze the impact of conditioning data on generation, we train multiple models using different conditioning inputs and conduct a comparative analysis of their generated outputs. In this comparison, conditioning data from specific historical periods are used as inputs, and we also verify their consistency with the corresponding historical data.

4.1 Impact of Conditional Inputs

First, we choose a model conditioned on currency, short-term interest rates, and inflation rates as the baseline model for our experiment. To assess the impact of conditional input, we prepare three models that incrementally add each condition data. The description of each model is summarized in Table 2.

Figure 1 compares the yield curves generated by each model using the conditional data from April 2015 with the actual yield curve observed in April 2015. To align with the monthly historical data, 30 samples are generated. Each colored line represents the sampled yield curve. The results indicate that the model adapts effectively to different scenarios, producing curves that reflect the specified economic conditions.

Table 2: Definition of Models in figure 1

Model Name	Shape of Condition Inputs
CUR	(Currency)
ST Rate	(Currency , Short-Term interest Rates)
Baseline	(Currency , Short-Term interest Rates, Inflation Rates)

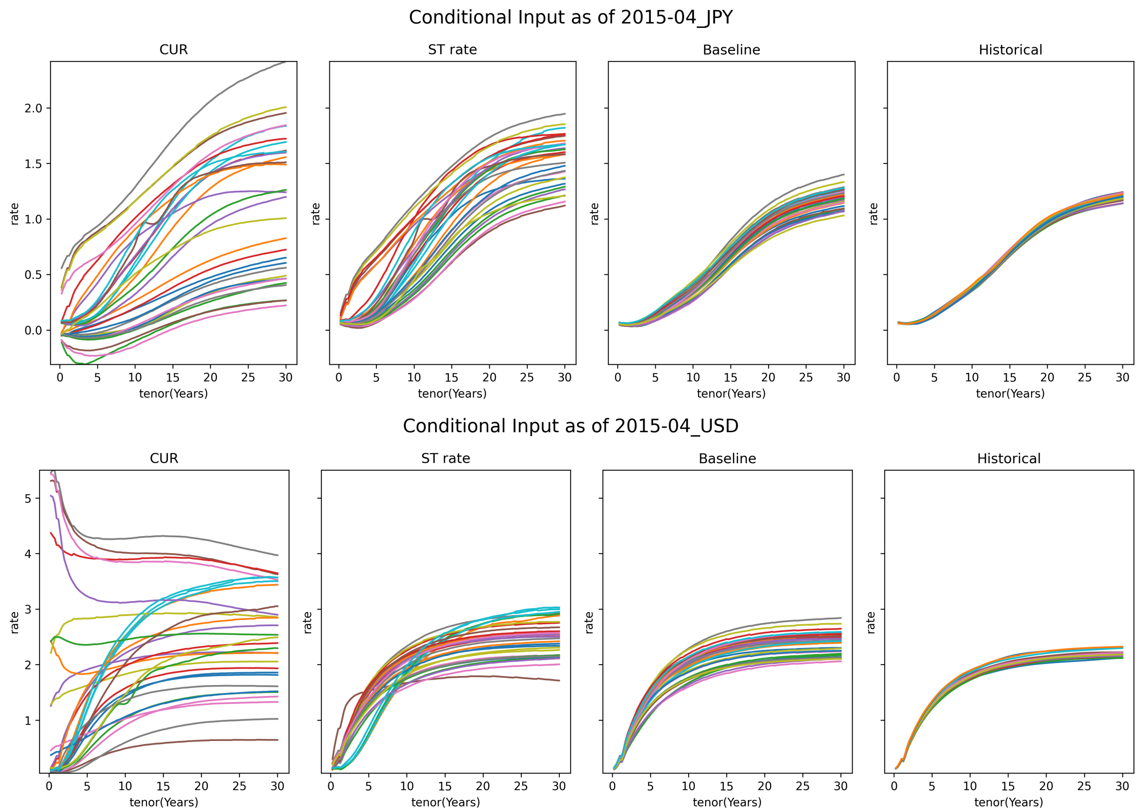

Figure 1: Comparison of generated and historical yield curves by adding conditions

4.2 Empirical Comparison with Historical Data

To evaluate the accuracy of our baseline model, we compare synthetic yield curves generated under the conditions of a short-term interest rate and an inflation rate for a given historical month with the actual yield curves from that same period. Since short-term interest rates are available as daily data, we use the monthly average as an input condition. The number of generated samples is 30 samples, which is determined based on the number of days in each month. Given that there are three currencies, three comparisons are conducted for each month.

The comparison results for JPY interest rates are shown in Figure 2, those for USD interest rates are presented in Figure 3, and for GBP interest rates are in Figure 4. In Figures 2 through 4 , for better visual clarity, we calculate the mean and standard deviation of the generated yield curves under each condition, and visualize them using the mean $\pm$ 1 standard deviation bands. These are then compared with the corresponding mean and standard deviation of the historical yield curves.

This approach enabled us to clearly assess how closely the model could reproduce key characteristics of the yield curve, especially level, slope, and curvature, under specified macroeconomic conditions. While greater variance is observed during periods of extreme monetary policy (e.g., 2020), the generated curves generally capture the level, slope, and curvature of historical data. Focusing on the results for USD and GBP in 2025, it can be observed that the model is capable of reproducing not only normal yield curves but also inverted ones, confirming the model's ability to reproduce realistic yield curve shapes under standard macroeconomic conditions.

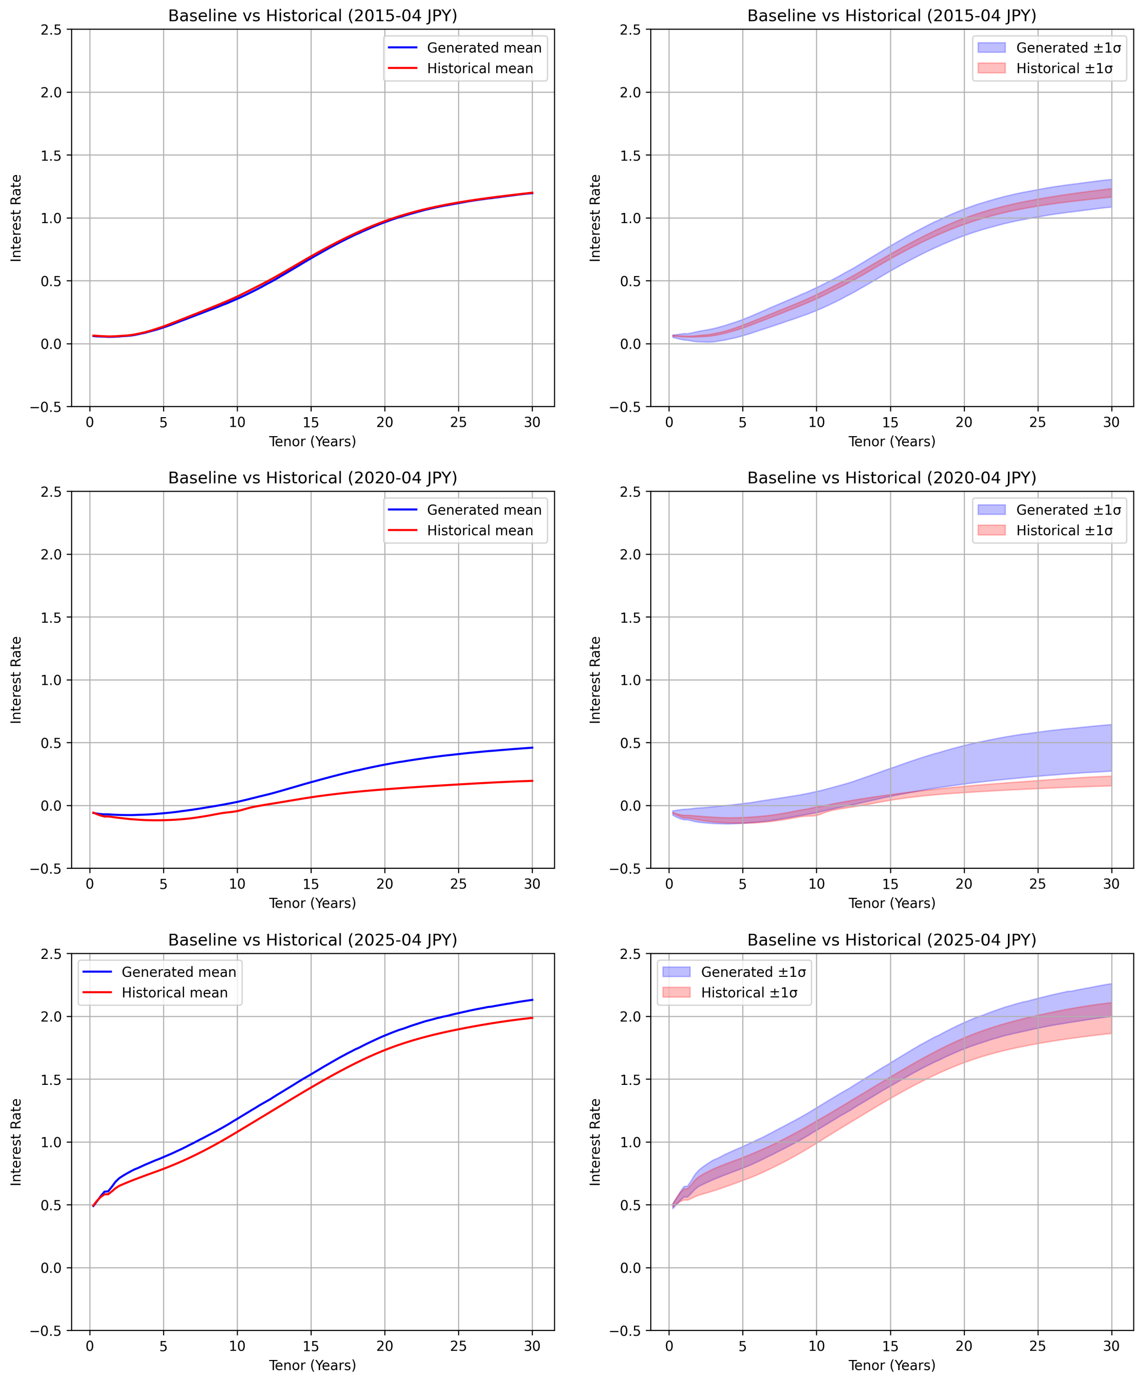

Figure 2: Comparison between generated sample and historical data: JPY

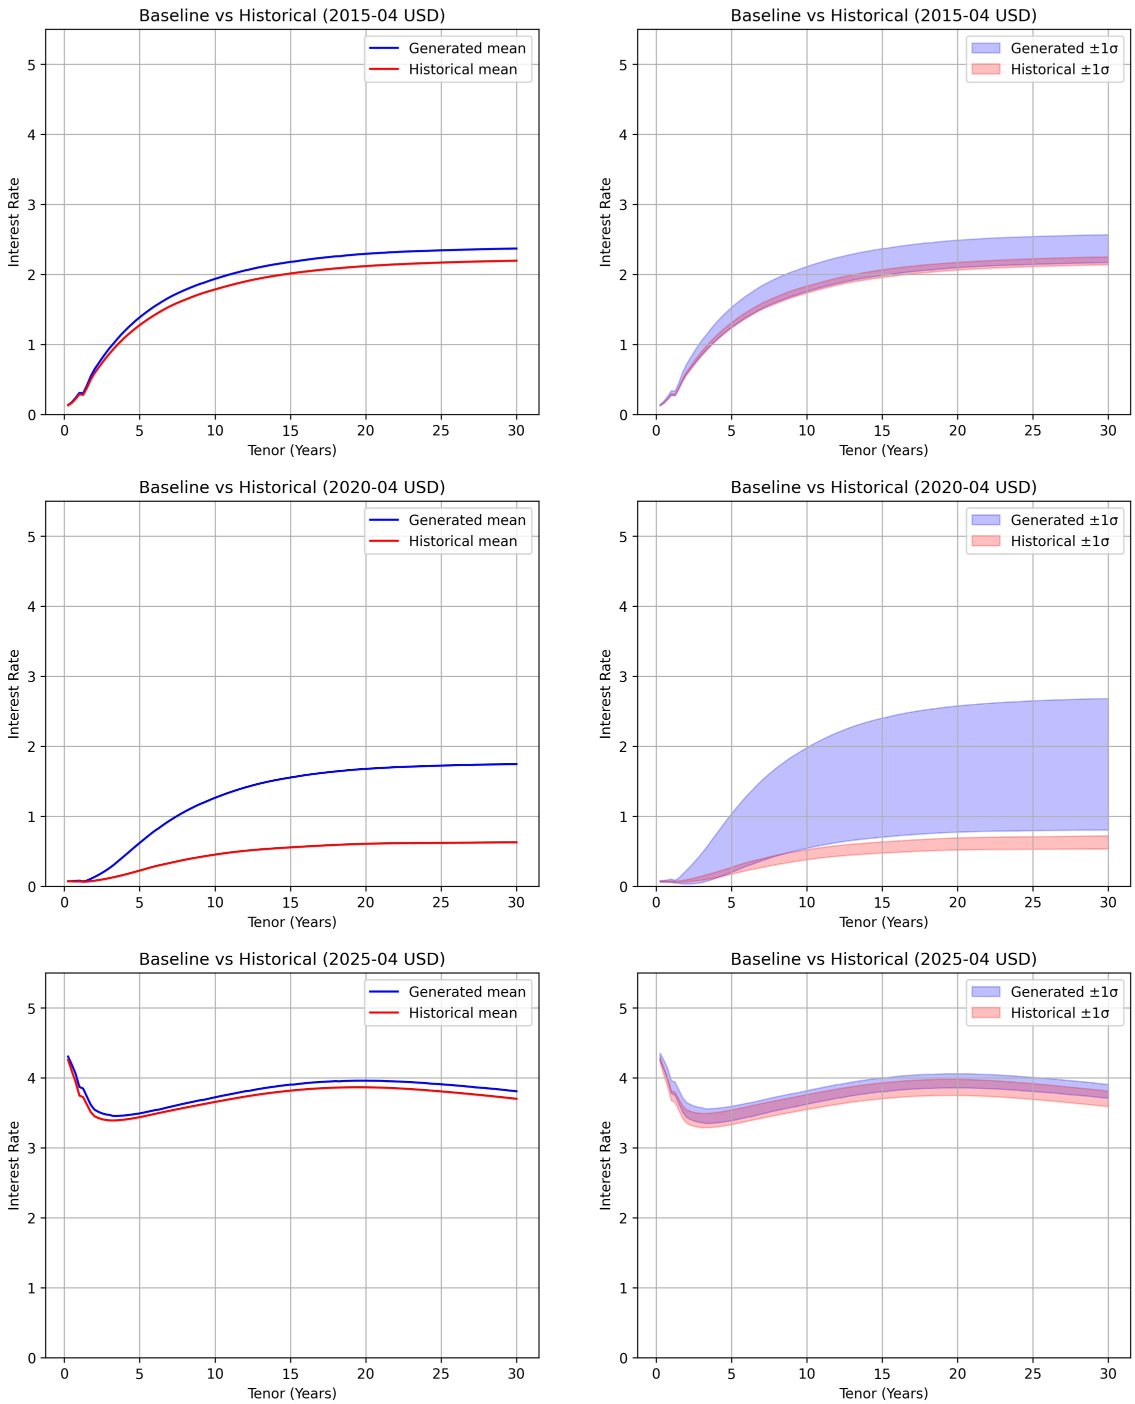

Figure 3: Comparison between generated sample and historical data: USD

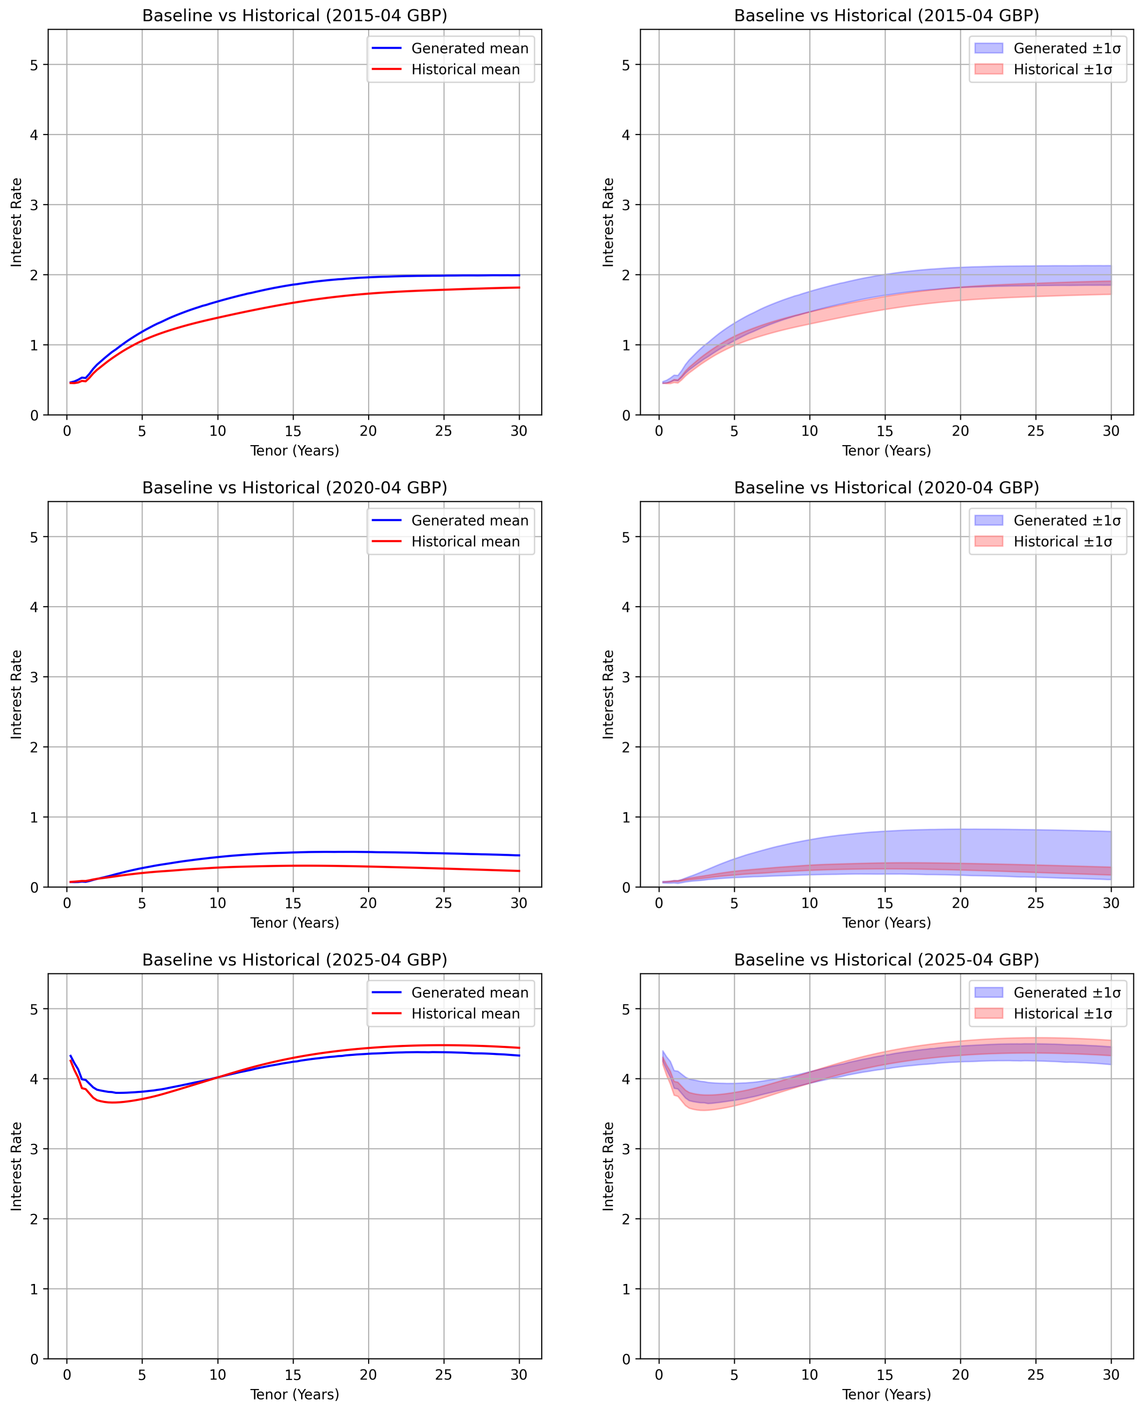

Figure 4: Comparison between generated sample and historical data: GBP

4.3 Further Analysis with Additional Macroeconomic Indicators

We further examine the incorporation of additional macroeconomic indicators into the baseline condition set(currency,short-term interest rates, and inflation rates) to capture broader structural features across different periods and market environments.

To better capture the extreme monetary policies in 2020, we introduce (i)the quantitative easing (QE) indicator specifically referring to the year-over-year (YoY) growth of the asset purchase amounts by each country's central bank, or proxy variables thereof. The exact definitions of the indicators are provided in Appendix B. For this dataset, we use LSEG Datastream and publicly available figures from each central bank as data sources.

In addition to the QE indicator, we consider indicators as follows: (ii) monetary base YoY growth, reflecting the intensity of monetary expansion, (iii) the current account balance to GDP ratio, and (iv) nominal GDP. Similarly to the QE indicator extension, these indicators are incorporated separately and simultaneously into the baseline condition set for model training and sample generation.

To assess the effect of these additional macro conditions, we evaluate the yield curves generated under the specified conditions by computing (i) the root mean square error (RMSE) between the generated and historical yield curves for each currency in representative months, and (ii) the maximum standard deviation (σ) across the samples generated, which reflects the dispersion and stability of the generation. The numerical results under different conditioning sets are summarized in Table 3 to Table 5. The results suggest that, for different time periods and currencies, augmenting the baseline conditioning set with an additional indicator that is relevant to the prevailing macro-financial environment can further improve the quality of the generated samples.

4.4 Enhancement with Additional Macro Economic Indicators as Conditions

As a result of RMSE comparison tests conducted with the addition of various macroeconomic indicators, it is found that the inclusion of the QE indicator and monetary base (QE+MB) for JPY and the inclusion of the QE indicator, the monetary base and the ratio of current account to GDP (QE + MB + CA% GDP) for USD and GBP is particularly effective in reproducing historical data. Based on this finding, an enhanced model is developed by adding these conditional data to the conditions of the baseline model and training it accordingly, followed by the generation of sample data.

Figures 5 to 7 present a graphical comparison between the generated samples and the historical data. The results generated by the enhanced model demonstrate superior reproducibility of historical data compared to those generated by other models. In particular, there is a significant improvement in reproducibility for the year 2020, which had been a challenge for the baseline model.

These findings indicate that the proposed method enhances the reproducibility of historical data by selecting and adding conditional models. In practical applications, depending on the intended use of synthetic data, different approaches may be considered: for cases requiring strict reproducibility of historical data, generation using the model with added conditional data is recommended; on the contrary, when a certain degree of flexibility in synthetic data is desirable, for example, in stress testing, generation using the baseline model may be more appropriate.

Table 3: RMSE and maximum standard deviation of generated yield curves under different conditioning sets (JPY). Values are reported as RMSE (basis points) with maximum standard deviation (percentage) in parentheses.

Conditions	2015-04	2020-04	2025-04
Historical	0.0 (0.03)	0.0 (0.04)	0.0 (0.12)
Baseline	1.1 (0.11)	16.1 (0.19)	10.9 (0.13)
Baseline+QE	0.7 (0.08)	9.2 (0.19)	1.8 (0.11)
Baseline+MB	3.2 (0.07)	2.7 (0.06)	2.9 (0.15)
Baseline+CA%GDP	0.5 (0.04)	2.5 (0.08)	10.8 (0.13)
Baseline+NOMGDP	3.1 (0.09)	10.4 (0.19)	11.4 (0.12)
Baseline+QE+MB	1.0 (0.07)	1.6 (0.05)	0.4 (0.10)
Baseline+QE+MB+CA%GDP	1.3 (0.06)	1.0 (0.05)	11.5 (0.13)
Baseline+QE+MB+NOMGDP	2.2 (0.08)	1.5 (0.06)	1.2 (0.11)
Baseline+QE+MB+CA%GDP+NOMGDP	2.5 (0.07)	0.5 (0.04)	4.6 (0.11)

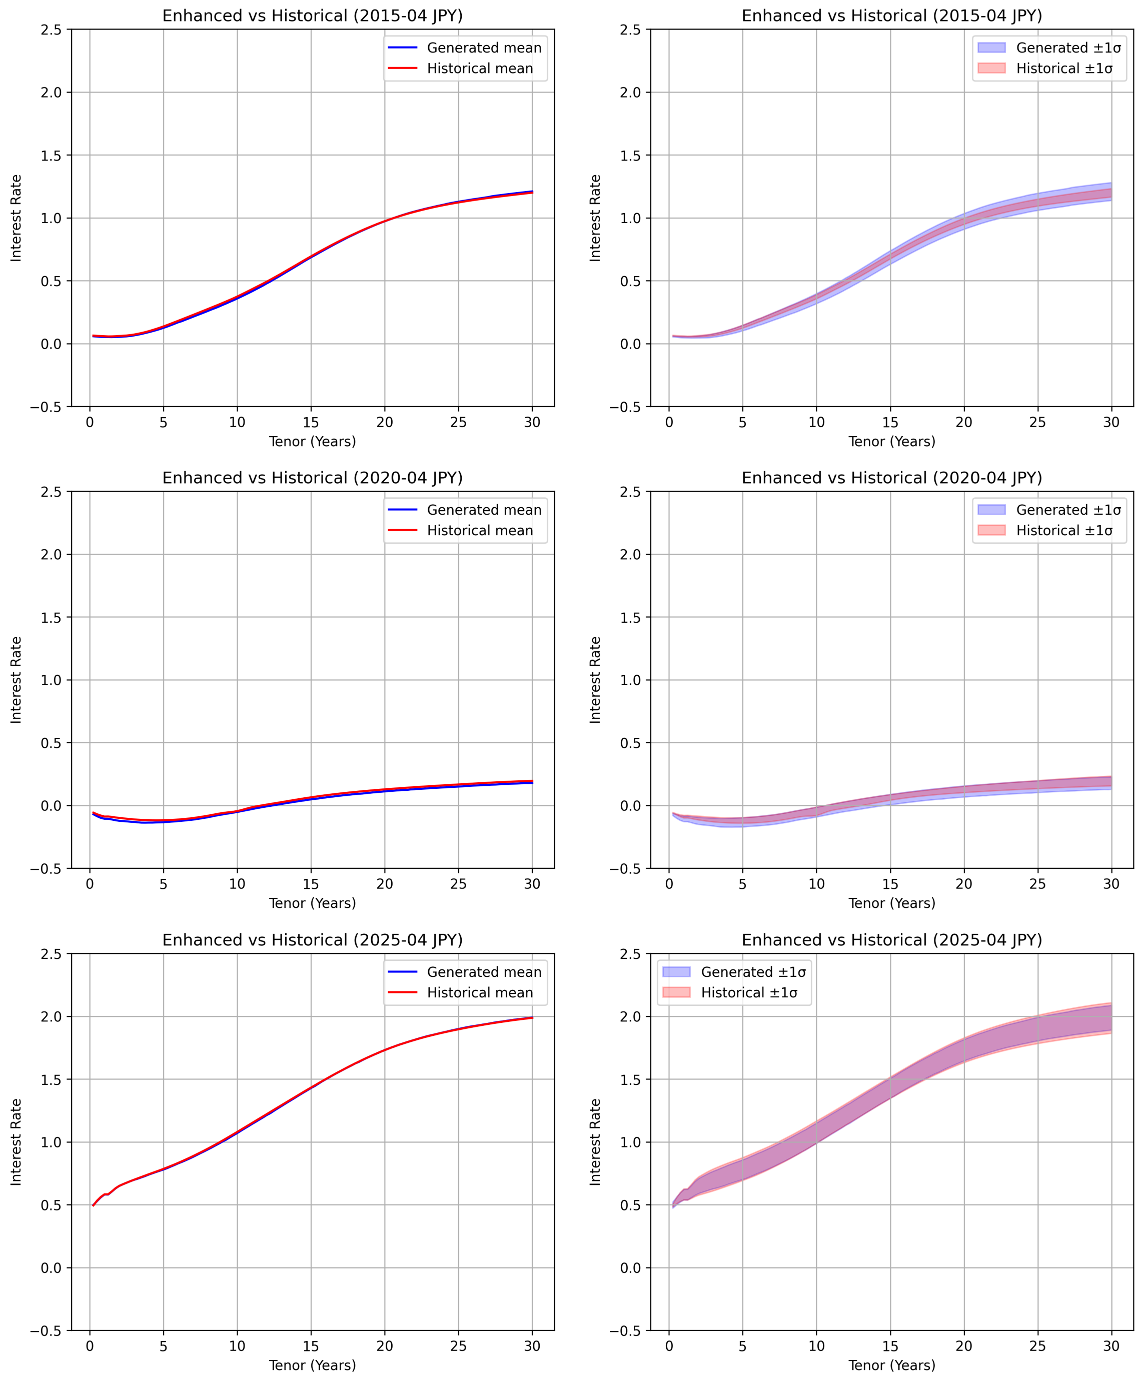

Figure 5: Enhancement with additional condition inputs: JPY

Table 4: RMSE and maximum standard deviation of generated yield curves under different conditioning sets (USD). Values are reported as RMSE (basis points) with maximum standard deviation (percentage) in parentheses.

Condition	2015-04	2020-04	2025-04
Historical	0.0 (0.06)	0.0 (0.09)	0.0 (0.11)
Baseline	15.2 (0.20)	89.6 (0.94)	8.7 (0.11)
Baseline+QE	15.3 (0.19)	1.9 (0.08)	4.9 (0.11)
Baseline+MB	9.7 (0.12)	33.1 (0.85)	6.0 (0.11)
Baseline+CA%GDP	13.4 (0.20)	207.7 (0.93)	9.1 (0.11)
Baseline+NOMGDP	14.3 (0.17)	45.8 (0.51)	7.0 (0.15)
Baseline+QE+MB	5.8 (0.11)	0.9 (0.07)	4.6 (0.10)
Baseline+QE+MB+CA%GDP	3.8 (0.10)	1.0 (0.08)	1.3 (0.10)
Baseline+QE+MB+NOMGDP	11.4 (0.14)	1.8 (0.09)	8.9 (0.13)
Baseline+QE+MB+CA%GDP+NOMGDP	9.1 (0.12)	1.0 (0.08)	8.5 (0.11)

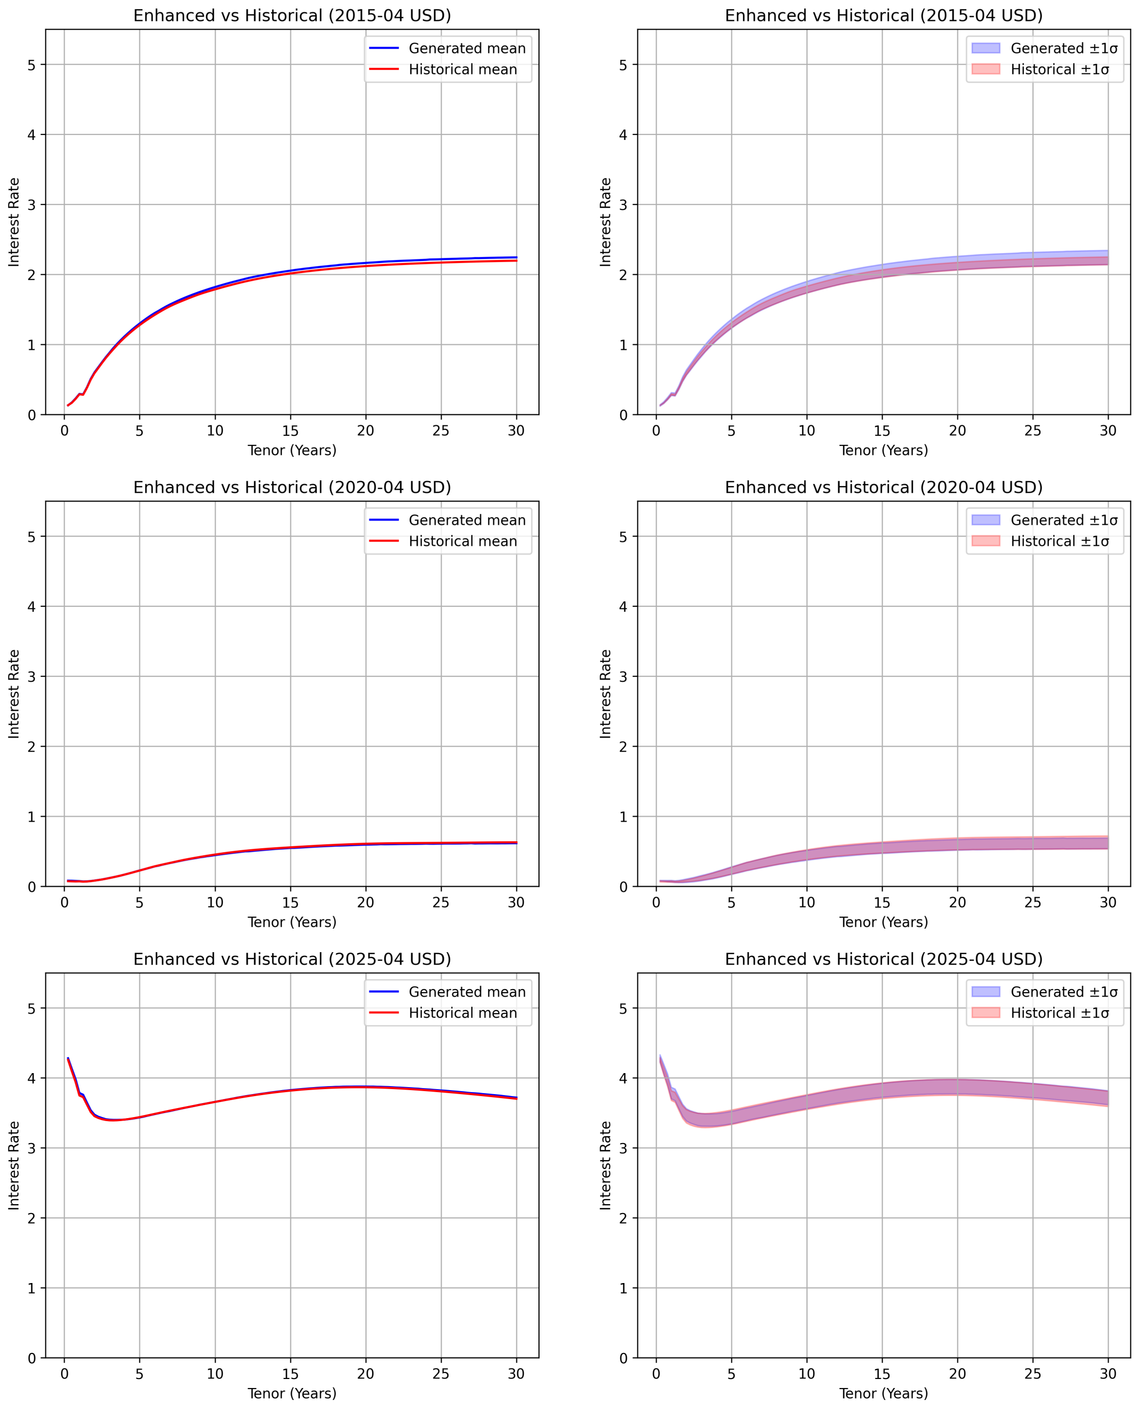

Figure 6: Enhancement with additional condition inputs: USD

Table 5: RMSE and maximum standard deviation of generated yield curves under different conditioning sets (GBP). Values are reported as RMSE (basis points) with maximum standard deviation (percentage) in parentheses.

Condition	2015-04	2020-04	2025-04
Historical	0.0 (0.10)	0.0 (0.06)	0.0 (0.11)
Baseline	20.5 (0.15)	17.3 (0.34)	8.8 (0.16)
Baseline+CA%GDP	22.8 (0.19)	1.5 (0.06)	7.0 (0.14)
Baseline+NOMGDP	18.7 (0.16)	23.5 (0.38)	13.2 (0.14)
Baseline+QE	20.4 (0.17)	0.4 (0.06)	2.3 (0.11)
Baseline+MB	11.3 (0.16)	32.4 (0.39)	11.4 (0.15)
Baseline+QE+MB	13.5 (0.16)	0.2 (0.06)	0.8 (0.10)
Baseline+QE+MB+CA%GDP	2.4 (0.15)	0.9 (0.05)	1.3 (0.10)
Baseline+QE+MB+NOMGDP	12.9 (0.16)	1.1 (0.05)	3.2 (0.10)
Baseline+QE+MB+CA%GDP+NOMGDP	4.4 (0.14)	1.4 (0.10)	2.6 (0.11)

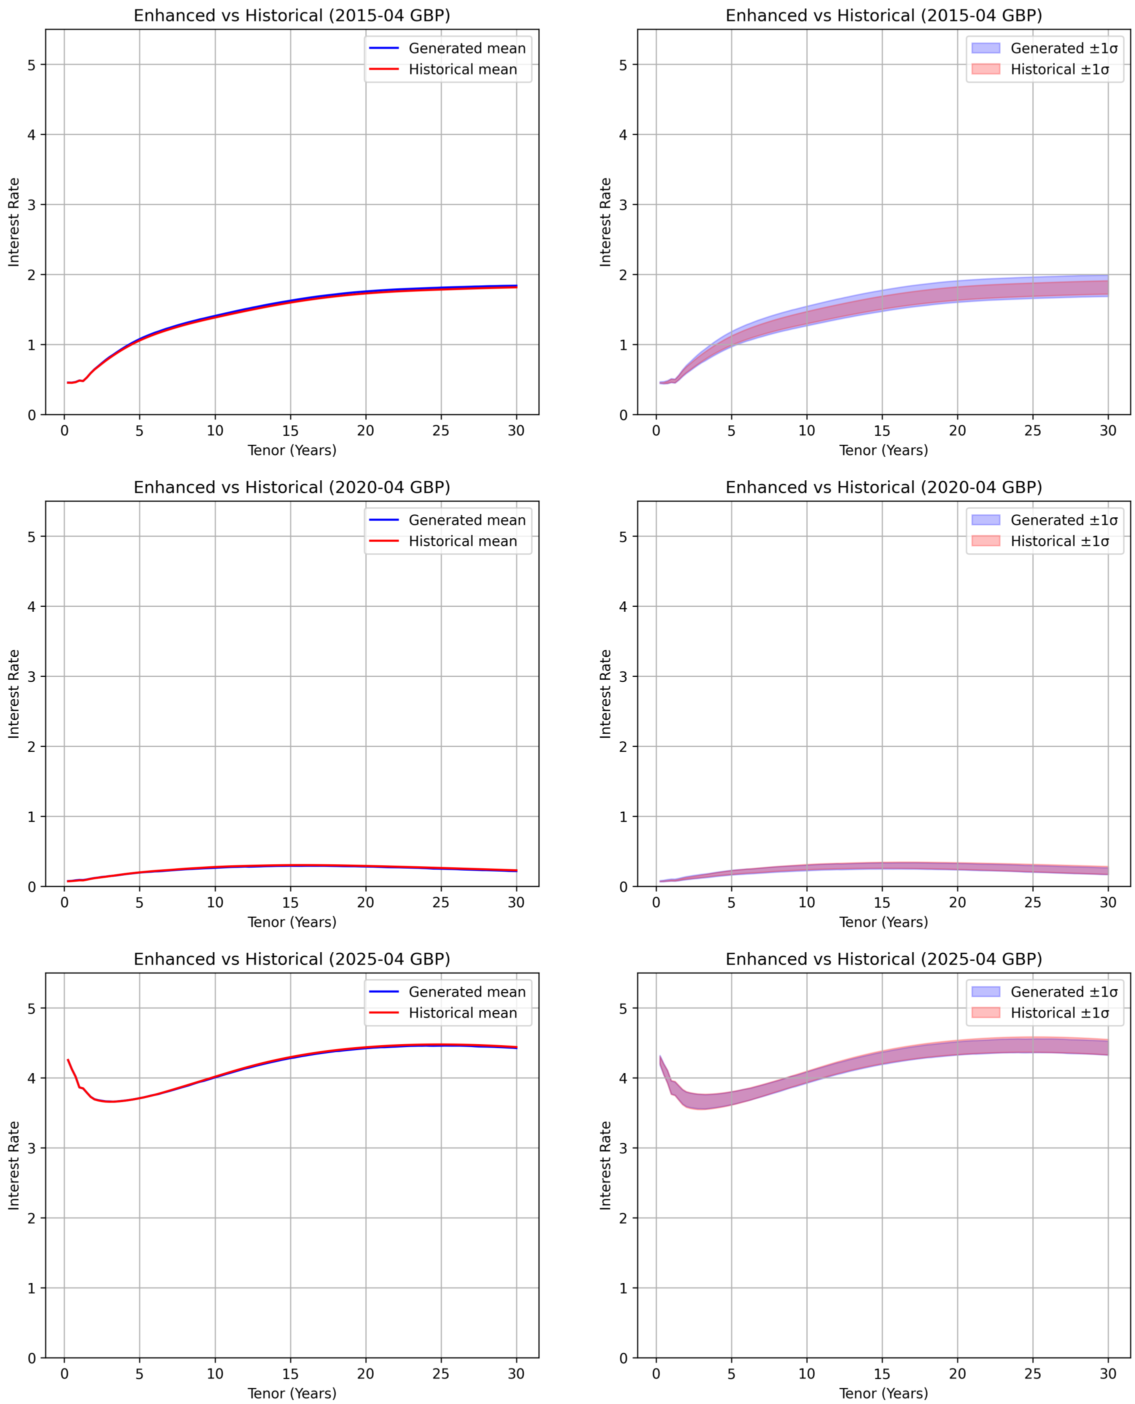

Figure 7: Enhancement with additional condition inputs: GBP

5 Results of Out-of-sample Generation

In the previous section, we compare the yield curves generated by the model using macroeconomic conditions within the training period with the historical yield curve data. On the other hand, we think that one of the important potential practical application of synthetic yield curves is scenario analysis in bond investment strategies. In such cases, it is crucial to determine whether yield curves can be generated by inputting hypothetical future macroeconomic conditions into a model trained on historical data up to that point. To evaluate this capability, we conducted the following simulation.

5.1 Models Used for Out-of-sample Generation

The diffusion models used for out-of-sample generation are constructed in the same manner as those described in the previous section. We employed two models: the baseline model presented in Section 4.2, which uses currency, short-term interest rates, and CPI inflation as conditions, and the enhanced model introduced in Section 4.4, which incorporates additional macroeconomic indicators. The key difference from the models used in Section 4 lies in the training data. Both the interest rate and condition data are limited up to the end of December 2023. Notably, a characteristic of Japan's historical interest rate data is that, during the training period up to the end of 2023, the yield curve remained at low levels under continued monetary easing. However, interest rates began to rise in 2024 following the implementation of rate hikes. By restricting the training data to the end of 2023, we aimed to examine whether the model could generate yield curves that reflect the subsequent rise in interest rates.

5.2 Comparison of Out-of-sample Generation with Historical data

Using the model trained on data limited to the end of 2023, we generated synthetic yield curves by inputting the conditional data for each month after 2024, and compared them with the corresponding historical yield curves, following the same procedure as in Section 4. The results generated by the baseline model are shown in Figure 8 through Figure 10. Overall, even in out-of-sample generation, the shapes of the yield curves—such as upward-sloping or inverted curves—do not deviate significantly from the historical data. When comparing the average of the generated results with the monthly averages of the historical yield curves, we observe a maximum deviation of approximately 50 basis points for JPY and up to 1% for GBP. In contrast, for USD, the generated yield curves remain generally close to the historical curves for about half year following the end of the training period.

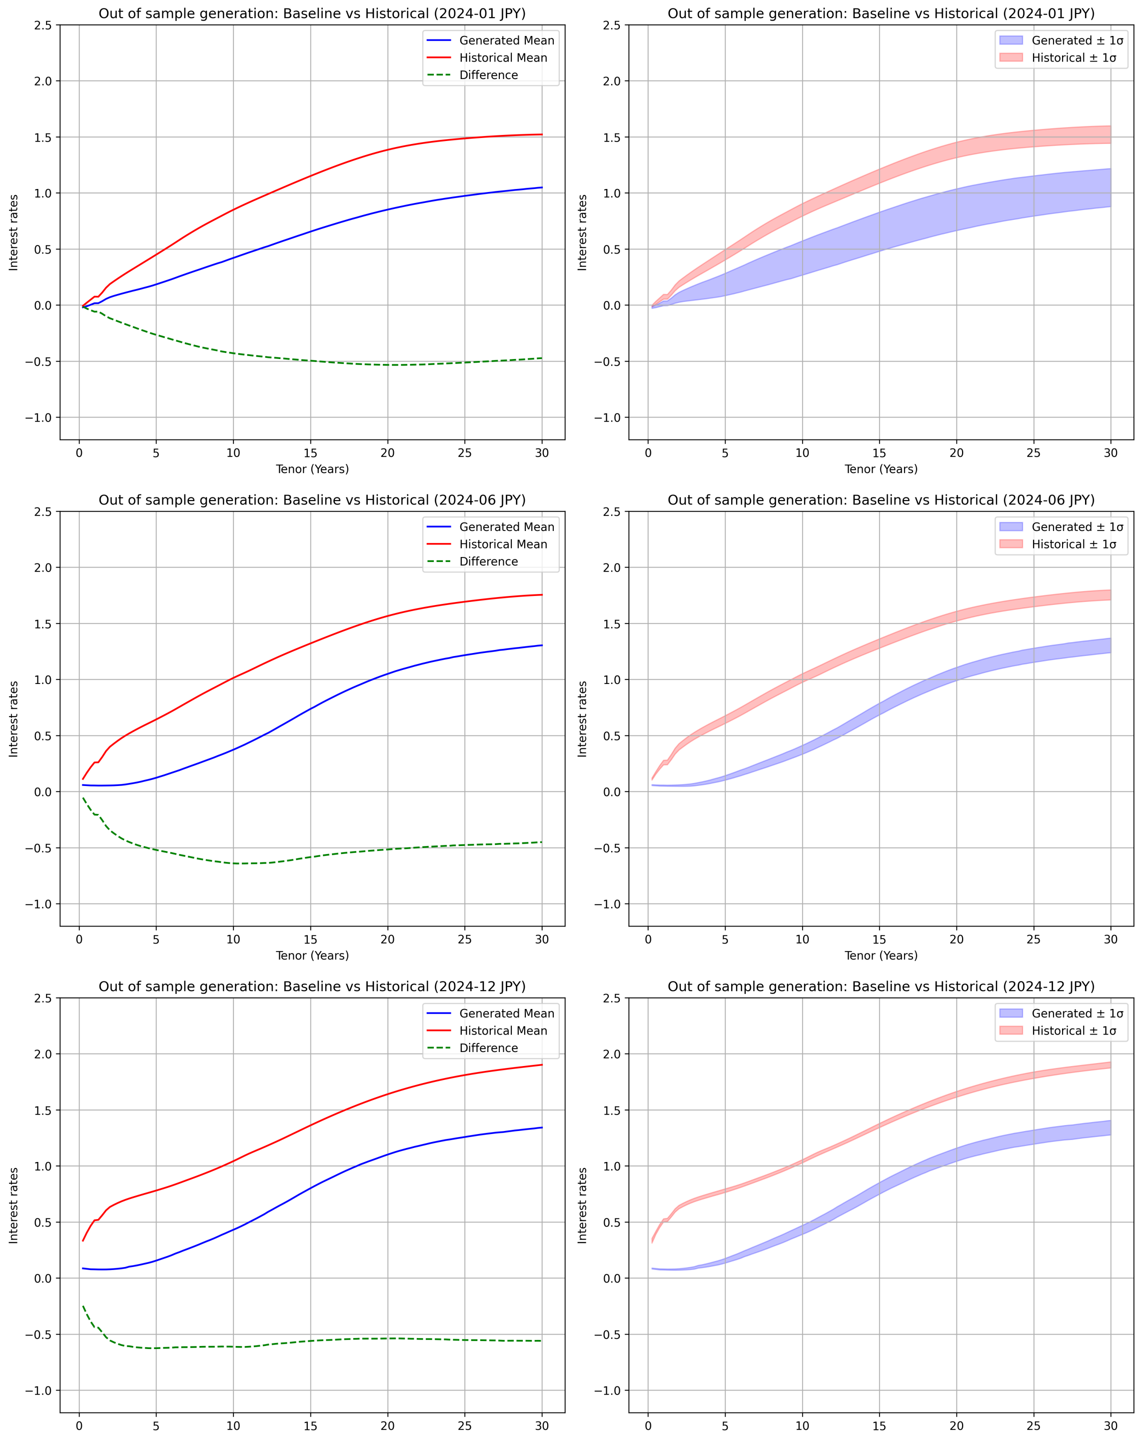

Figure 8: Out-of sample comparison of Baseline generation: JPY

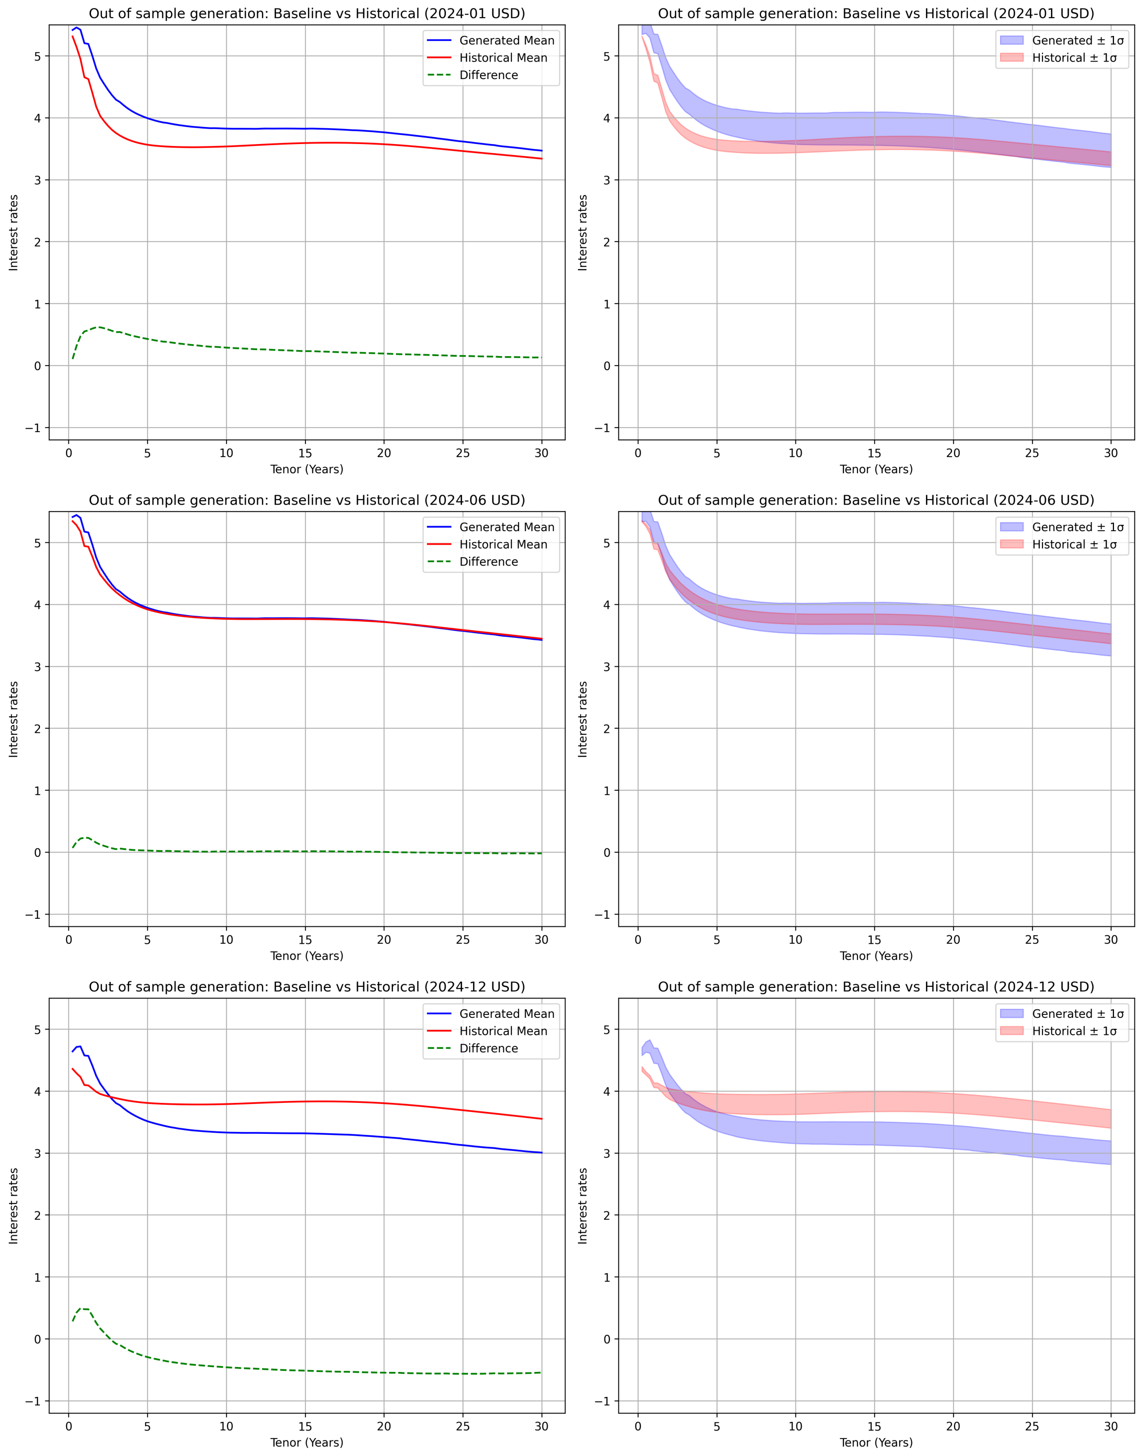

Figure 9: Out-of sample comparison of Baseline generation: USD

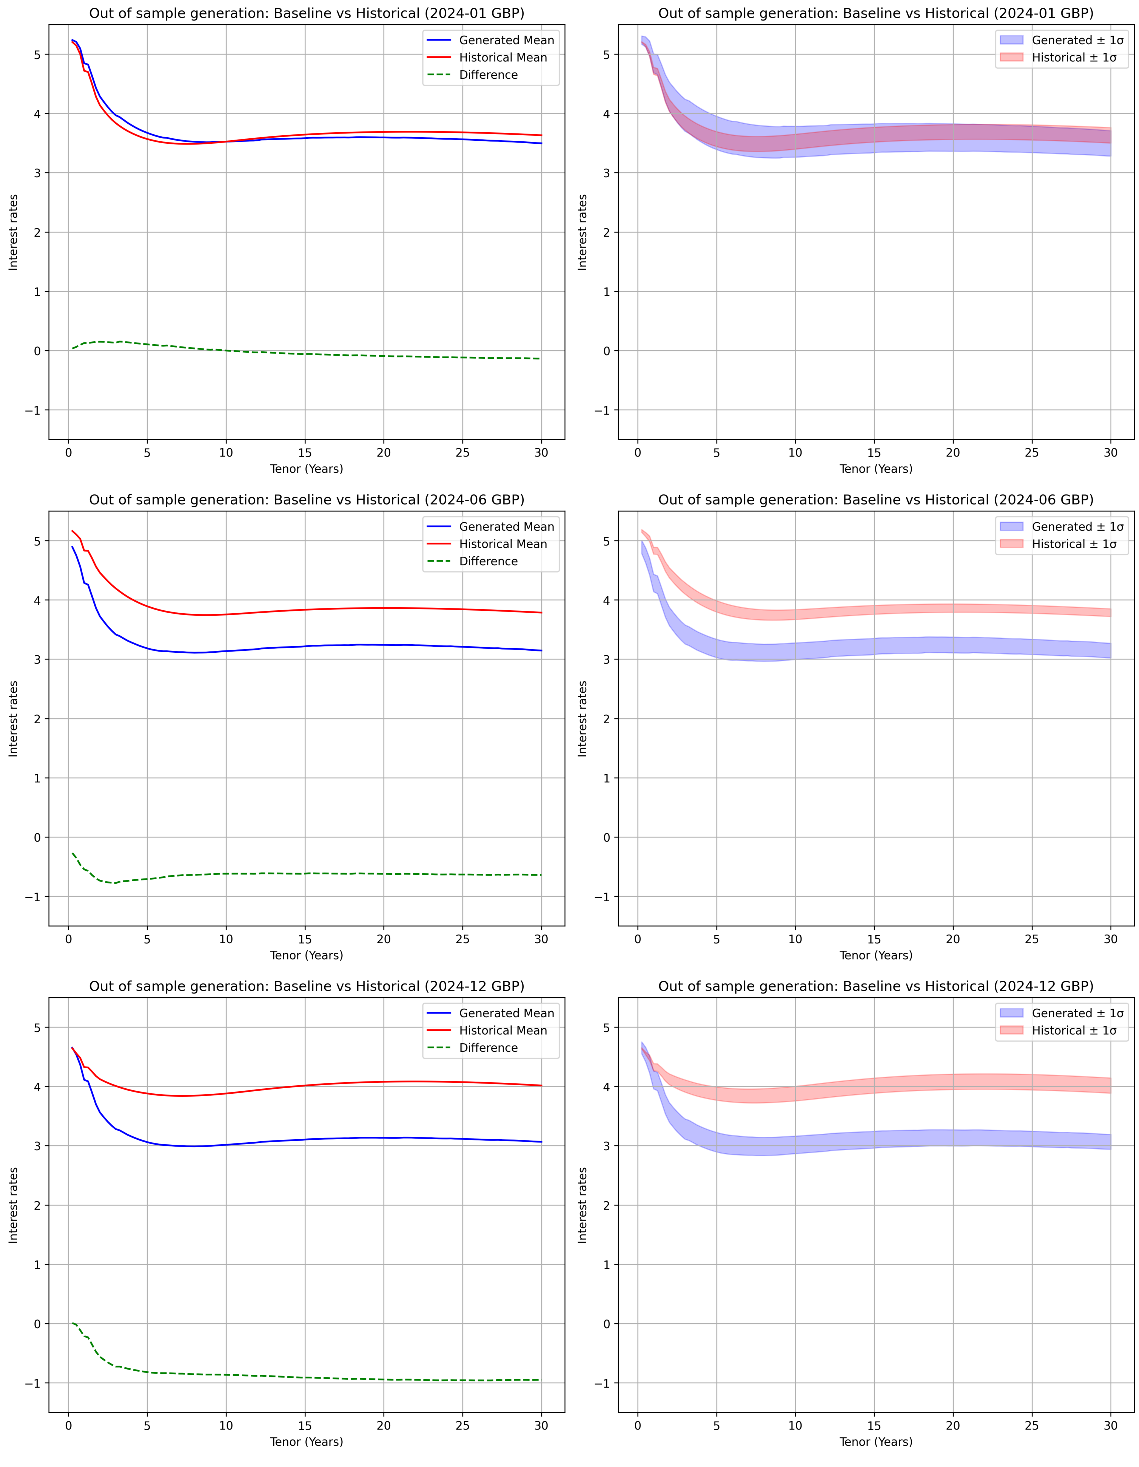

Figure 10: Out-of sample comparison of Baseline generation: GBP

Next, Figure11 through 13 presents a comparison of the generation results using the enhanced model, which incorporates additional conditioning variables beyond those used in the baseline model. For USD, the generated yield curves also remained generally close to the historical yield curves for approximately half year following the end of the training period. As for the Japanese yen, while the baseline model consistently produced lower interest rates than the actual yield curves, the enhanced model is able to generate higher interest rates—sometimes exceeding the actual rates —and for tenors of around 25 to 30 years, the generated rates are nearly at the same level as the historical yield curves. This suggests that the enhanced model successfully captures the rise in interest rates associated with the tapering of monetary easing, even though such conditions are not included in the training period.

Based on these out-of-sample generation results, a practical application would be to generate a synthetic yield curve approximately six months ahead, grounded in the outlook for macroeconomic indicators over the same horizon. Such synthetic yield curves can then be utilized for investment strategy planning or risk management. Since the historical yield curve data and macroeconomic indicators used for training are continuously updated and can be obtained from information vendors such as Bloomberg, it is possible to periodically retrain and update the model.

As a validation designed for such usage, we train the enhanced model with data up to the end of June 2024, and generate synthetic yield curves using macroeconomic data as of December 2024 as conditional input. In Figure14, the generated results are compared with the actual yield curves observed in December 2024. The generation results based on the updated model are shown in orange, while the results from the model trained only on data up to the end of December 2023 are shown in blue. As a result, the discrepancies that had arisen prior to updating the training data—up to approximately 50 bps for JPY and around 1% for GBP—are largely corrected, and synthetic yield curves close to the historical yield curves are generated. In a similar vein, we train the enhanced model with data up to the end of December 2024, and generate synthetic yield curves using macroeconomic data as of June 2025 as conditional input. The results are shown in Figure15, and it presents that the outcomes are consistent with those in Figure14. Therefore, under the assumption that the conditional data evolves as expected, it is suggested that even in out-of-sample generation over a horizon of about six months, our enhanced model can produce synthetic yield curves that effectively represent the actual yield curves.

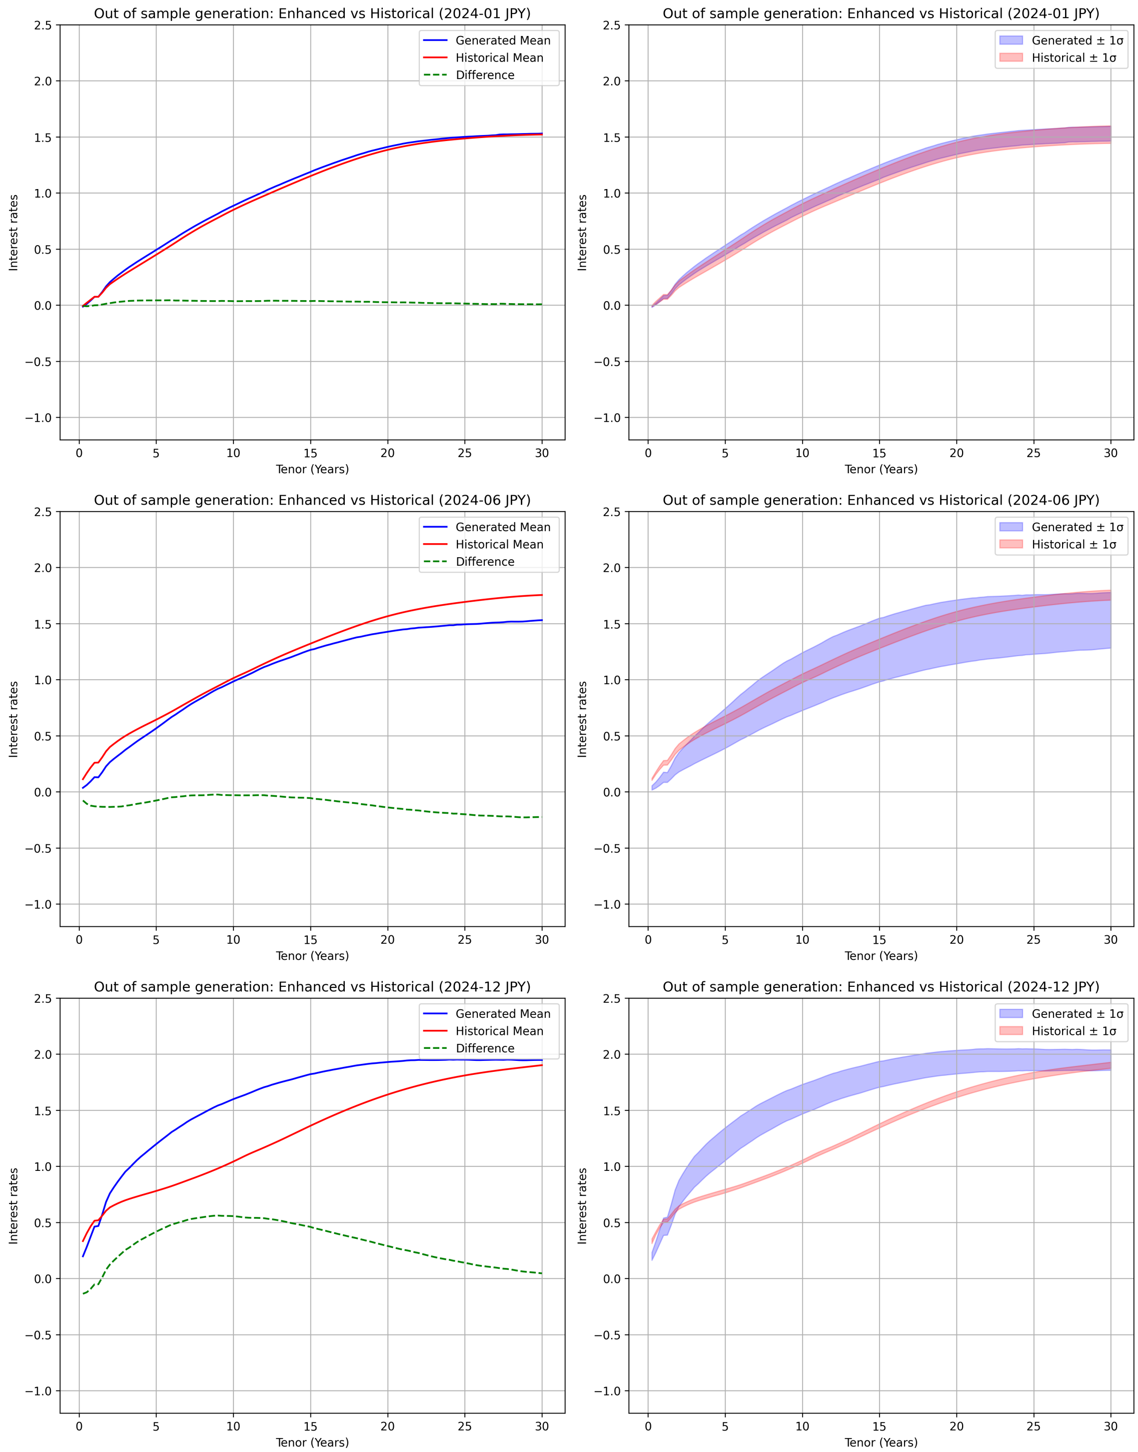

Figure 11: Out-of sample comparison of Enhanced generation: JPY

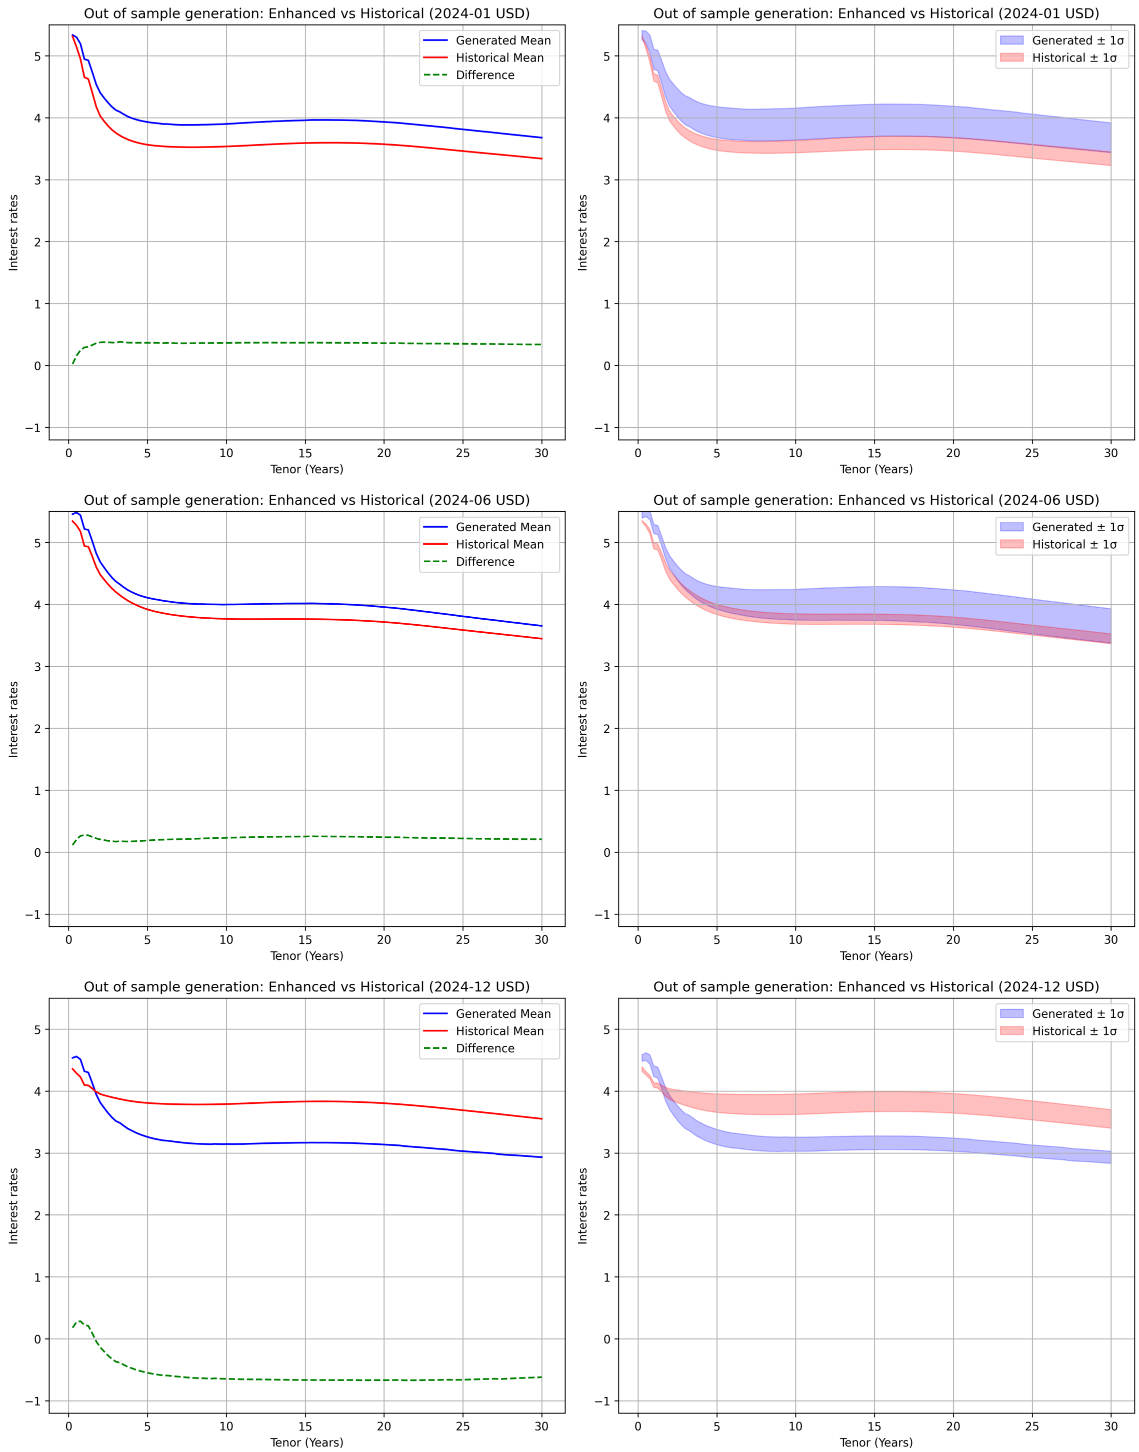

Figure 12: Out-of sample comparison of Enhanced generation: USD

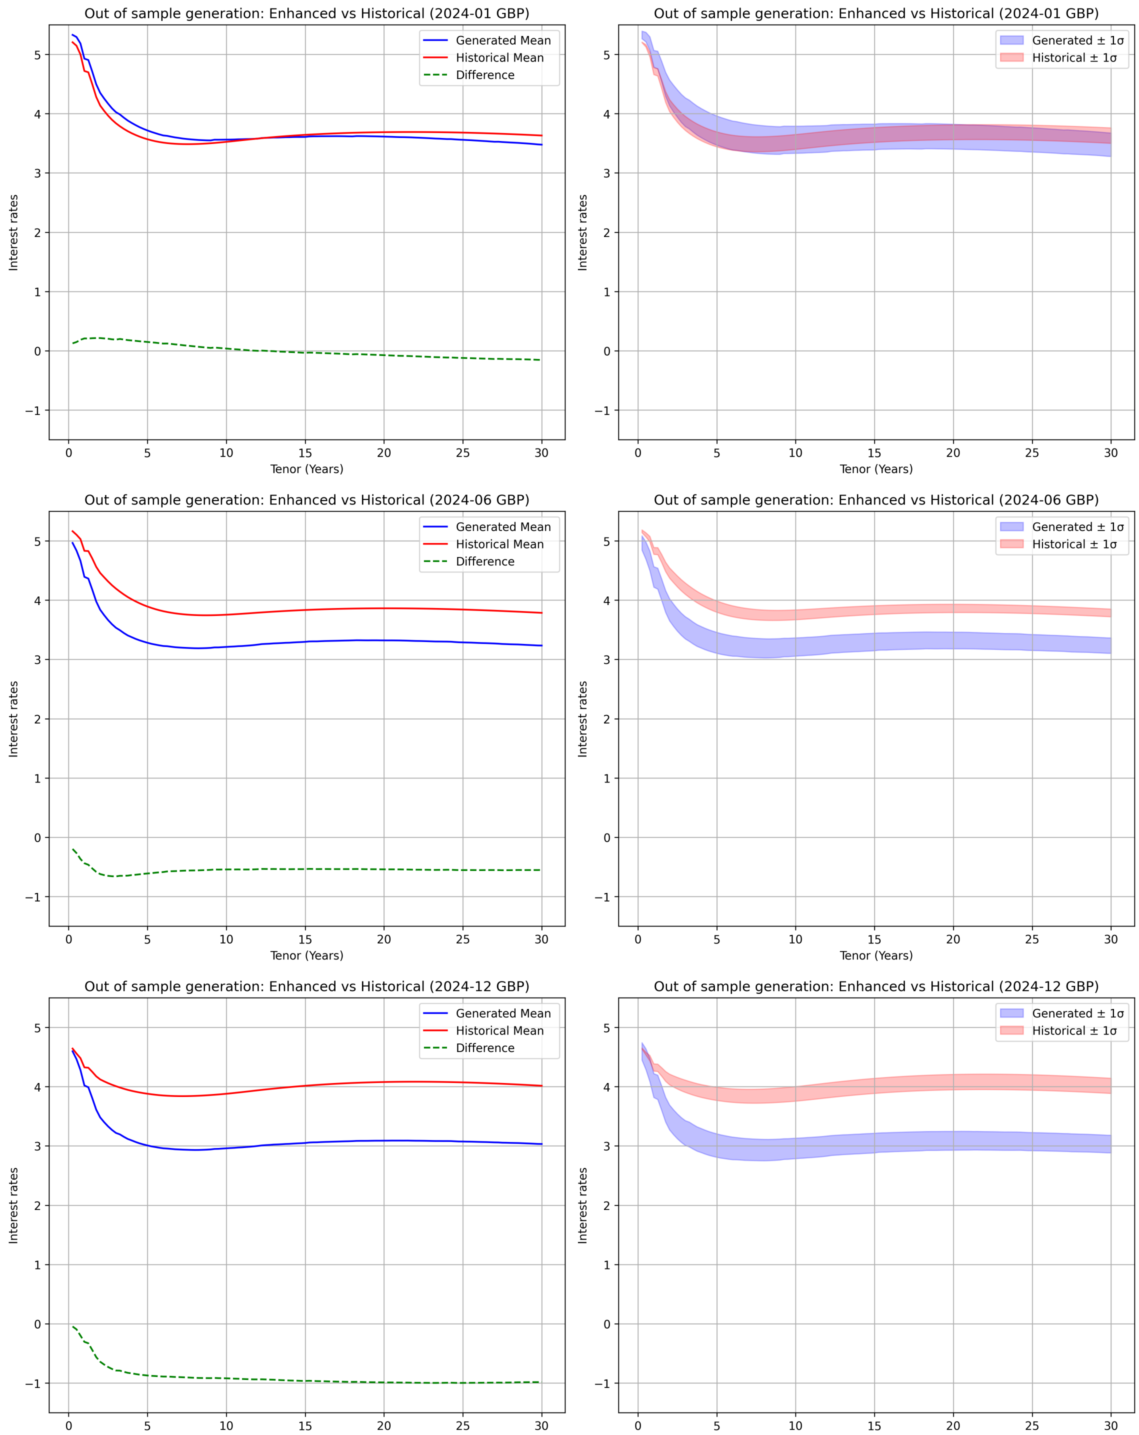

Figure 13: Out-of sample comparison of Enhanced generation: GBP

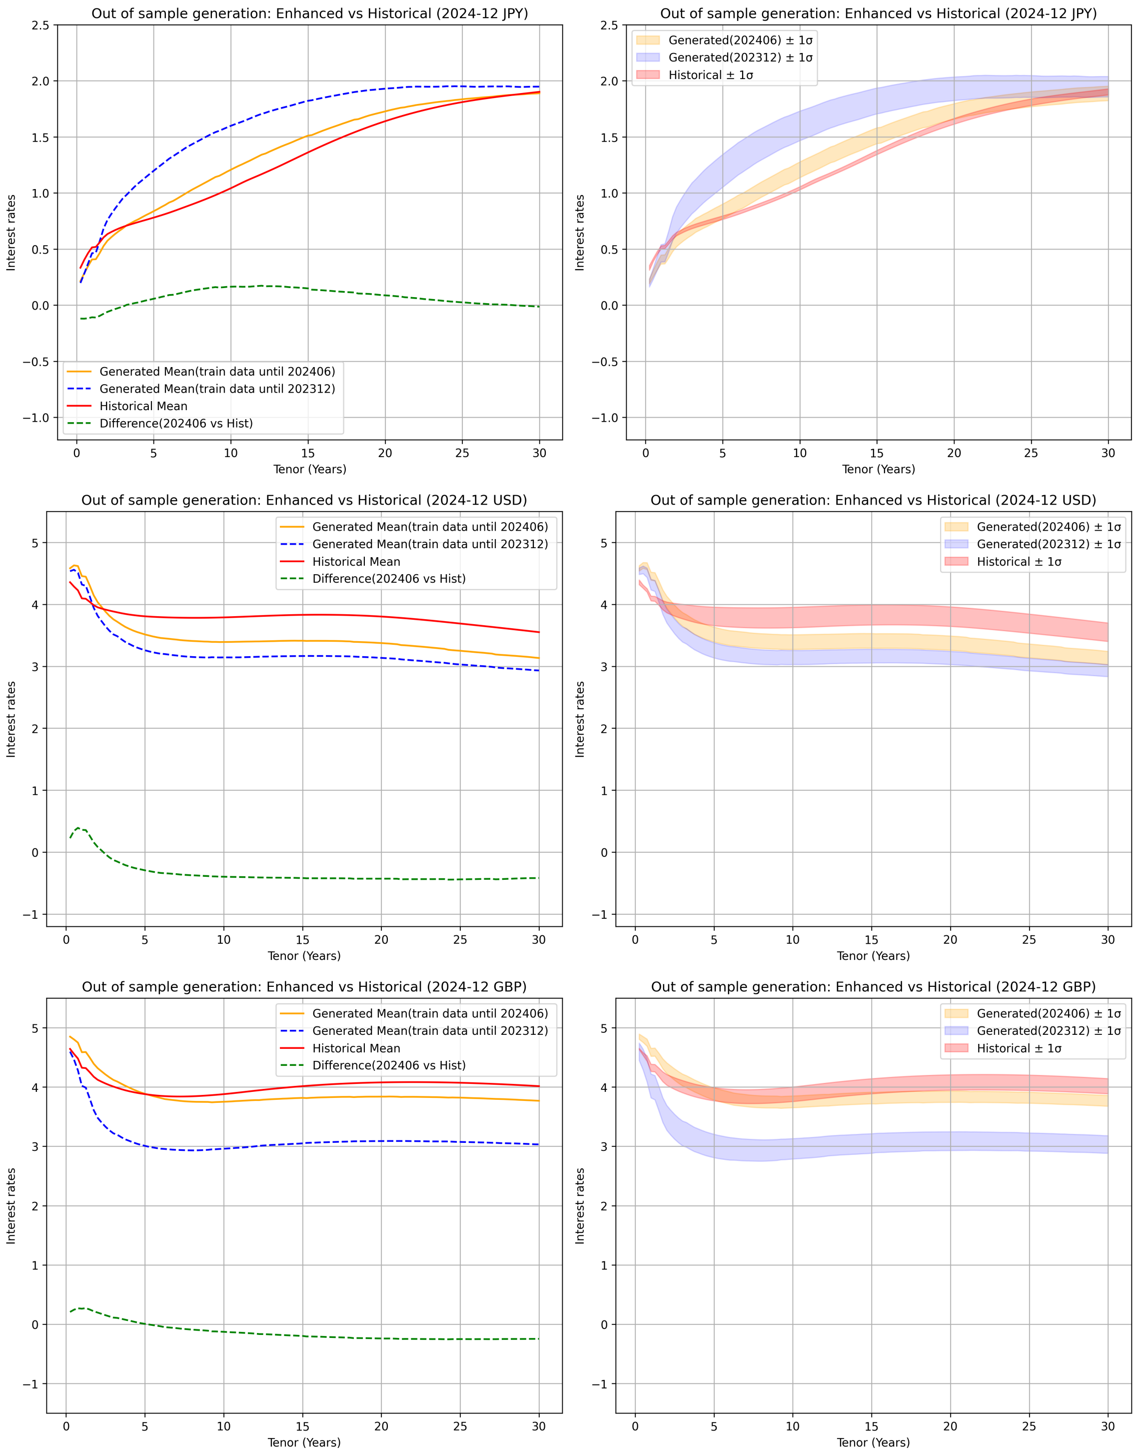

Figure 14: Six month ahead generation comparison: Enhanced model with data up to June 2024

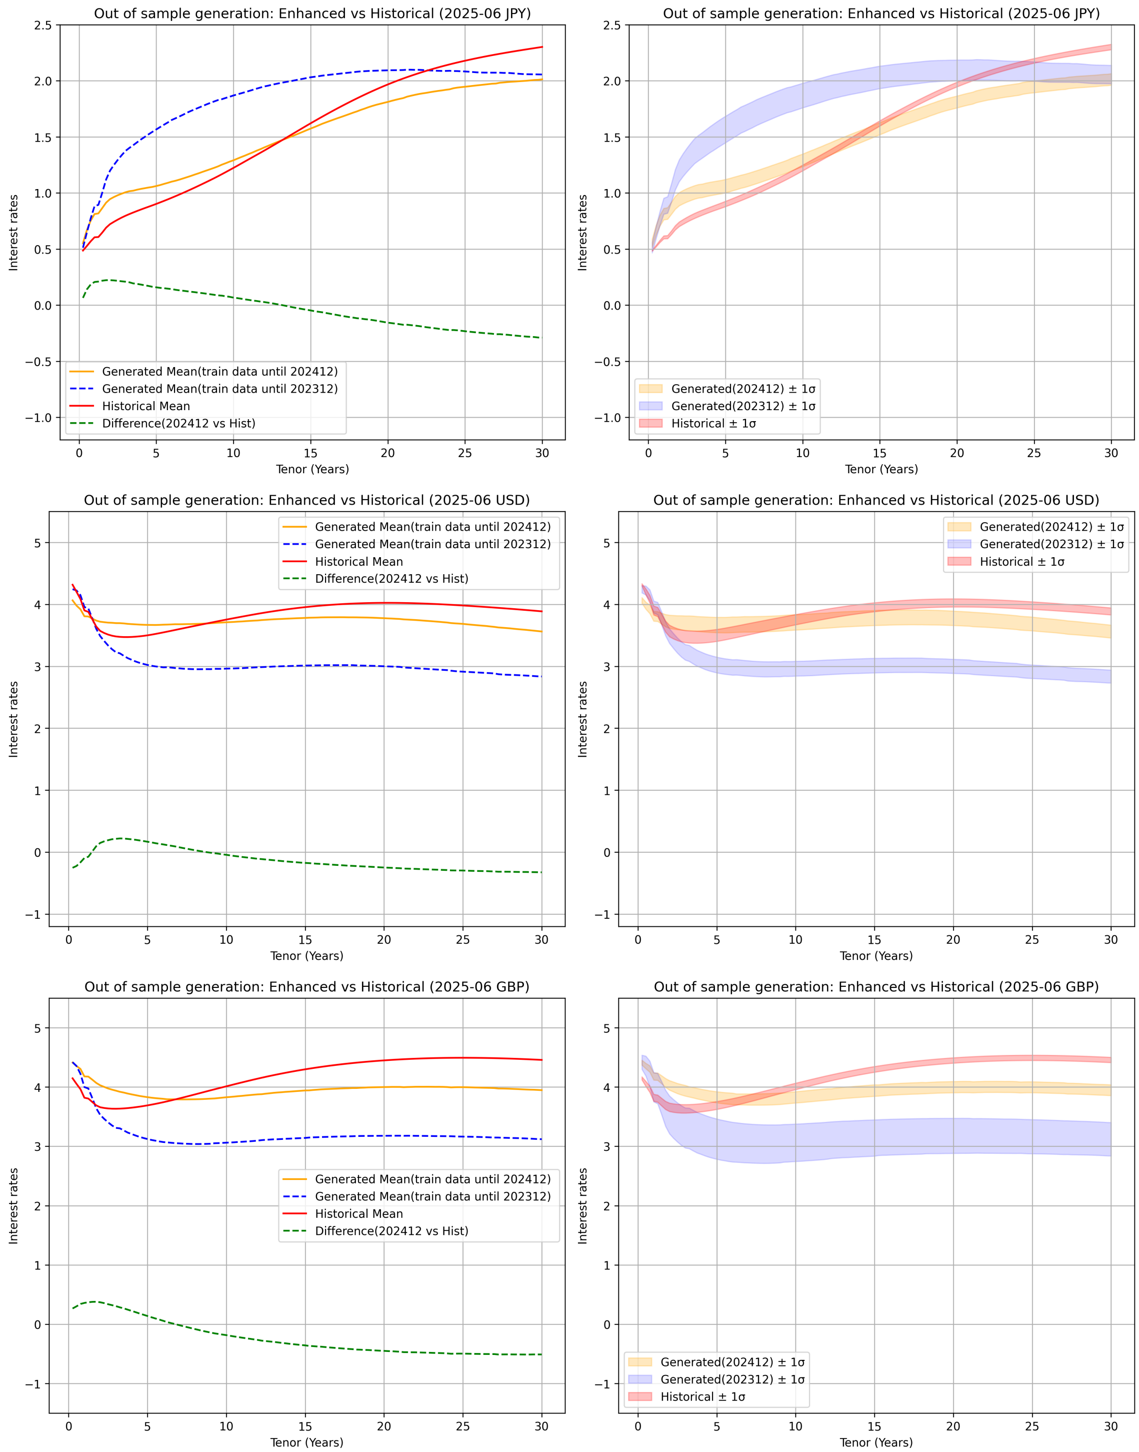

Figure 15: Six month ahead generation comparison: Enhanced model with data up to December 2024

5.3 Generation Results from a Model Trained on Yield Curve Variations Over a Specific Time Period

The practical application that has motivated the previous analysis of out-of-sample generation was to estimate the shape of the yield curve at a given time by using the future levels of macroeconomic indicators as conditional data. As another possible practical application, one could instead use the changes in macroeconomic indicators between the present and a future point in time as conditional data, in order to estimate the corresponding change in the yield curve from its current state. To examine the applicability of such an approach, we conduct generation using the differences in interest rate data over a certain period, derived from the interest data prepared in Section 3, as the training data, and likewise the differences in the conditional variables over the same period, derived from the conditional data prepared in Section 3 and Section 4, as the conditional input.

5.3.1 Data Preparation for Difference-Based Learning

The interest rate data and conditional data used for difference-based learning can be constructed by applying adjustments to the datasets prepared in Section 3. In this examination, the data are adjusted as follows:

- For each currency, spot rate curves at 0.25-year tenor intervals are created along with the corresponding conditional data, following the steps in Section 3. The reference date is set as the second business day of the month two months prior to the observation date. The differences in interest rates at each tenor between the reference date and the observation date are then calculated, and these differences are taken as the difference data for the respective observation dates. Similarly, differences are computed for the conditional variables except the label of currecy, and the resulting difference data for each currency are combined and used for learning with the diffusion model specified in Section 3.
- The macroeconomic indicators adopted as conditional data are, as in the previous generation, currency, short-term interest rates, and CPI growth rate for the Baseline model, and for the Enhanced model, the Baseline conditional data plus QE indicators and the CAGDP ratio (for USD and GBP only). Two models are thus constructed, each trained on its respective set of conditional data.

5.3.2 Result of Generation from Difference-Based Learning

Using the model constructed through difference-based learning, 1,024 samples are generated in the same manner as before, and their mean and one standard deviation are calculated. The mean and one standard deviation of the sampled difference data are then added to the yield curve at the reference date and compared with the historical yield curve two months after the reference date. In-sample generation, where the training data are not split, reproduced the historical yield curves in the same way as the previous analysis. For out-of-sample generation, in which the training data are limited to those up to the end of 2023, Figures 16, 17, and 18 compares the generation results from the prior section with those from the difference-based learning model. In each figure, the two rightmost columns show the results from the difference-based learning model: the orange dashed curve represents the yield curve at the reference date, the blue curve represents the yield curve at the reference date plus the generated results, and the green line represents the generated differences. The generation results from the difference-based learning model are broadly same as those obtained in the prior section. However, in some samples, it can be observed that the yellow curve lies between the red curve and the blue curve. This indicates that the actual direction of yield curve movements from the reference date differs from the direction suggested by the samples generated by the model. Identifying conditional data that contribute to improving the accuracy of direction remains an issue for future research. Moreover, while this examination focused on differences over a two-month period as the learning target, it is also conceivable to incorporate the interval of differences into the conditional data and construct a model that simultaneously learns interest rate difference data across varying intervals.

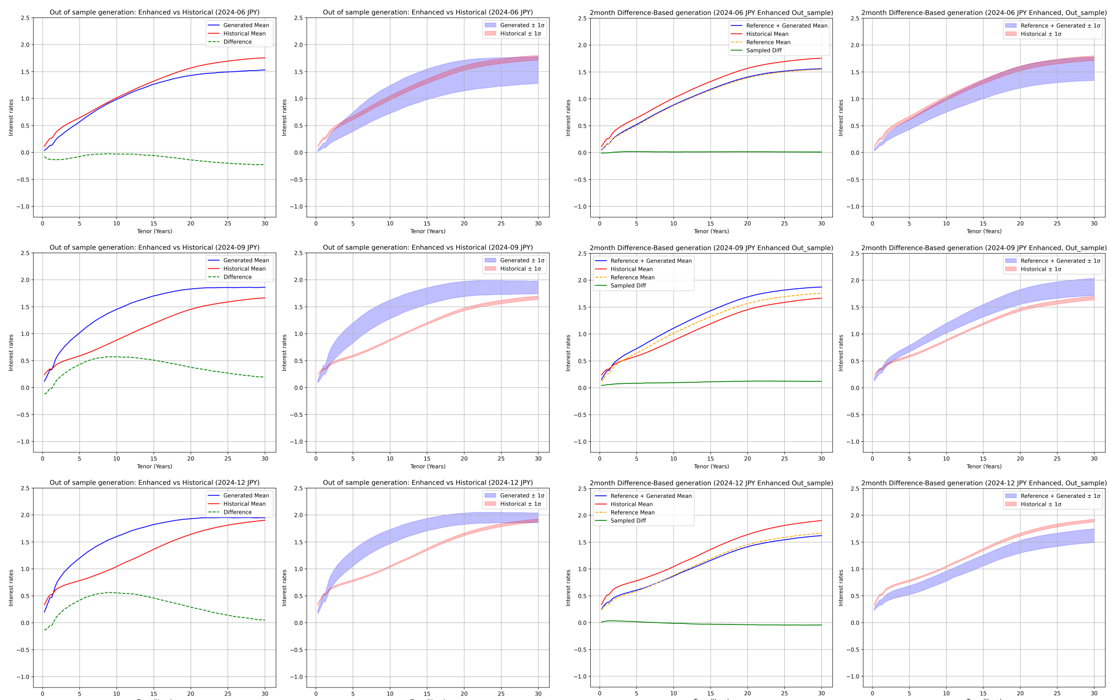

Figure 16: Out of sample generation comparison,right two columns present the results from Difference based model: JPY with Enhanced model

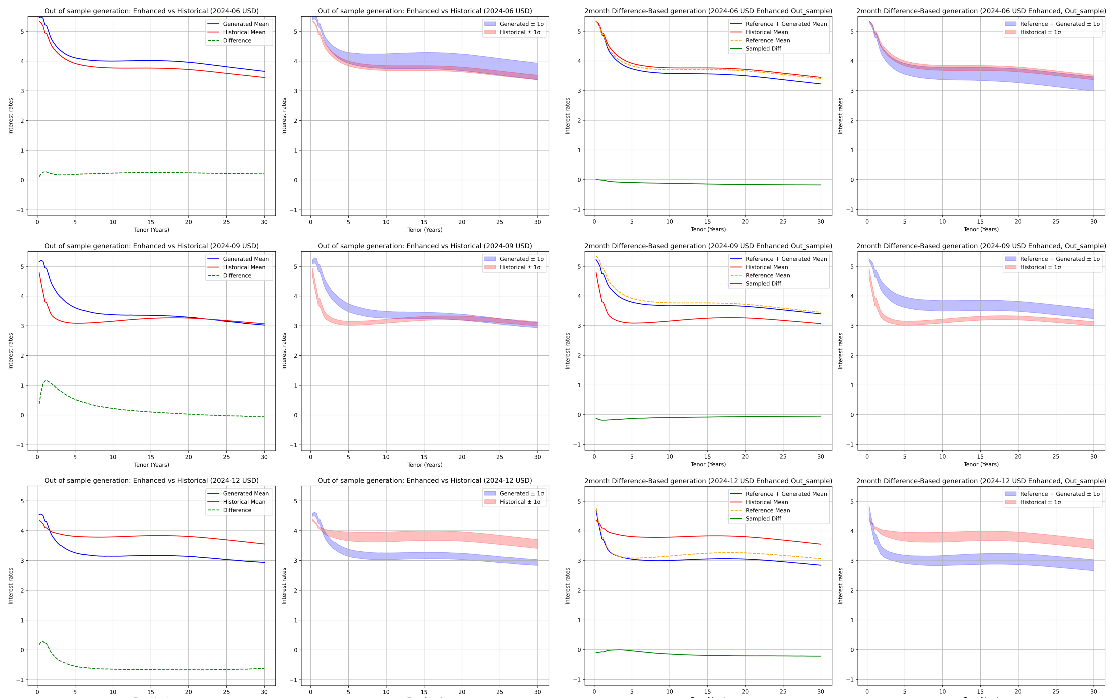

Figure 17: Out of sample generation comparison,right two columns present the results from Difference based model: USD with Enhanced model

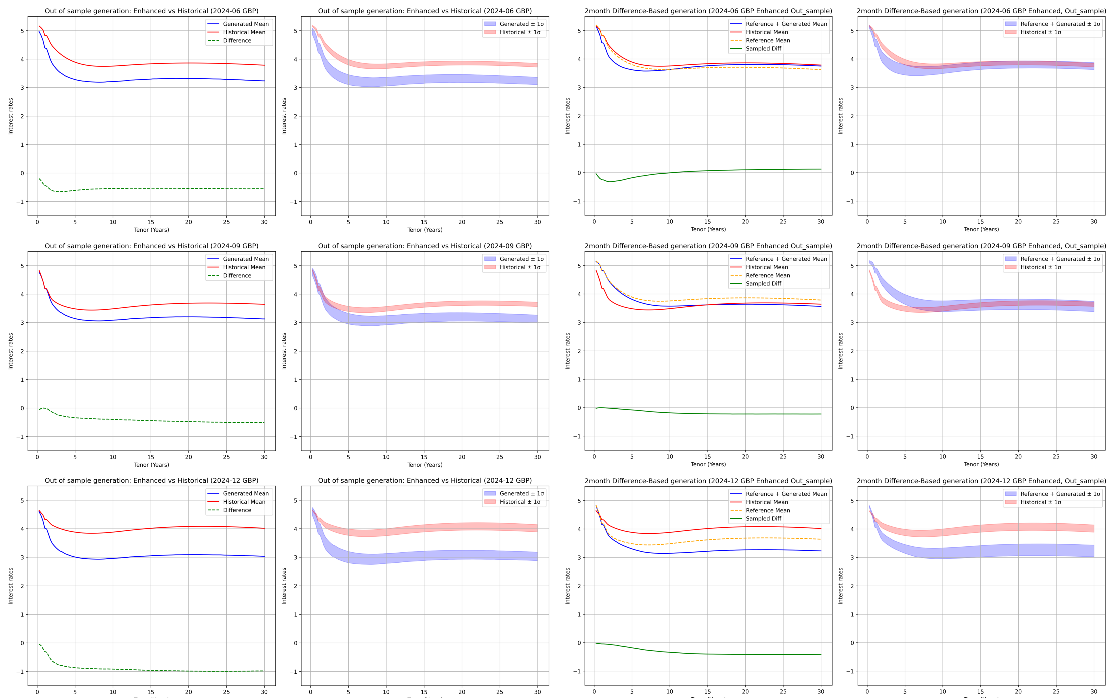

Figure 18: Out of sample generation comparison,right two columns present the results from Difference based model: GBP with Enhanced model

6 Results of Factor Model Based Generation

When considering applications in finance and macroeconomics, it is often desirable to analyze how the macroeconomic indicators—used as conditional data in previous generative models—affect changes in the shape of the yield curves, such as level, slope, and transitions between normal and inverted yield curves. With this application in mind, we conduct factor learning based on the Nelson-Siegel-Svensson model, which is a multi-factor model of the yield curve's functional form widely applied in finance and macroeconomics, to examine whether the generative model can learn how macroeconomic indicators are projected onto the characteristics of the yield curve.

6.1 Overview of Nelson-Siegel-Svensson Model

The Nelson-Siegel-Svensson (NSS) model, originally introduced by Nelson and Siegel[9] and later extended by Svensson[10] is a parsimonious function widely used to capture the term structure of interest rates. It augments the original Nelson-Siegel model with a second curvature component, thereby enhancing its flexibility to fit complex yield curve shapes, including those exhibiting secondary humps.

The model defines the zero-coupon yield $Y(\tau)$ for a given time-to-maturity τ as:

$$\begin{aligned} Y(\tau) = \beta_0 + \beta_1 \frac{1 - e^{-\tau/\lambda_1}}{\tau/\lambda_1} + \beta_2 \left(\frac{1 - e^{-\tau/\lambda_1}}{\tau/\lambda_1} - e^{-\tau/\lambda_1} \right) \\ + \beta_3 \left(\frac{1 - e^{-\tau/\lambda_2}}{\tau/\lambda_2} - e^{-\tau/\lambda_2} \right). \end{aligned} \tag{3}$$

The parameters have clear economic interpretations: β_0 represents the long-term level (the asymptote as $\tau \to \infty$); $\beta_0 + \beta_1$ represents the short-term level; β_2 and β_3 represent the magnitudes of the two curvature components. The parameters λ_1 and λ_2 dictate the maturity location of these curvature components and the overall decay speed of the term structure.

6.2 Estimation of Factors

To generate the training data for the factor-based model, we estimate the daily time-series for all six NSS parameters $(\beta_0, \beta_1, \beta_2, \beta_3, \lambda_1, \lambda_2)$, following the methodology of Sekine[11]. Our procedure involves a three-stage estimation process: (1) determining global fixed decay parameters (λ_1, λ_2) using the Differential Evolution (DE) algorithm on monthly subsets; (2) calculating the β coefficients via Ordinary Least Squares (OLS); and (3) refining all parameters using Non-linear Least Squares (NLS). The detailed procedure is described as the SV-NL model in the Appendix of Sekine[11].

Tables 6 to 8 summarize the estimated β factors (β_0 through β_3) and λ parameters derived from both the historical datasets.

Table 6: NSS factors and average RMSE across the periods (USD).All the factors represented are the mean of the monthly observations. Average RMSE represents the monthly average of daily RMSE.

Period	2015-04	2020-04	2025-04
β_0	0.0177	0.0026	0.0168
β_1	-0.0164	-0.0015	0.0277
β_2	-0.0193	-0.0073	0.0000
β_3	0.0187	0.0146	0.0675
λ_1	1.0629	1.8678	1.1536
λ_2	15.3437	12.3190	11.6298
average RMSE	0.0088	0.0050	0.0159

Table 7: NSS factors and average RMSE across the periods (JPY).

Period	2015-04	2020-04	2025-04
β_0	0.0702	0.0422	0.0778
β_1	-0.0690	-0.0426	-0.0716
β_2	-0.0581	-0.0323	-0.0551
β_3	-0.1500	-0.1069	-0.1500
λ_1	6.7174	6.5987	6.7425
λ_2	49.6361	42.8390	49.7002
average RMSE	0.0189	0.0051	0.0256

Table 8: NSS factors and average RMSE across the periods (GBP).

Period	2015-04	2020-04	2025-04
β_0	0.0040	-0.0016	0.0178
β_1	-0.0007	0.0022	0.0255
β_2	0.0081	-0.0254	0.0000
β_3	0.0441	0.0386	0.0849
λ_1	4.6771	4.0888	1.5449
λ_2	21.1370	8.2729	14.8194
average RMSE	0.0137	0.0023	0.0120

6.3 Results of Factor-Based Generation

To evaluate the performance of the factor-based model, we conduct the in-sample generation experiment outlined in previous sections. The model is conditioned on the macroeconomic test set to generate a synthetic time series of the six NSS parameters. Subsequently, we reconstruct the full synthetic yield curves by substituting these generated parameters back into the NSS equation. This procedure enables a direct comparison between the curves implied by the factor-based method and the realized historical curves.

Figures 19 and 20 present the generated samples. It is evident that even the baseline model, despite utilizing minimal conditional data, tracks the historical yields closely. The enhanced model, which incorporates the comprehensive conditional dataset described in Section 4, exhibits superior performance. This is particularly notable in the US market during April 2020, where the model successfully captures the extreme market volatility triggered by the COVID-19 shock.

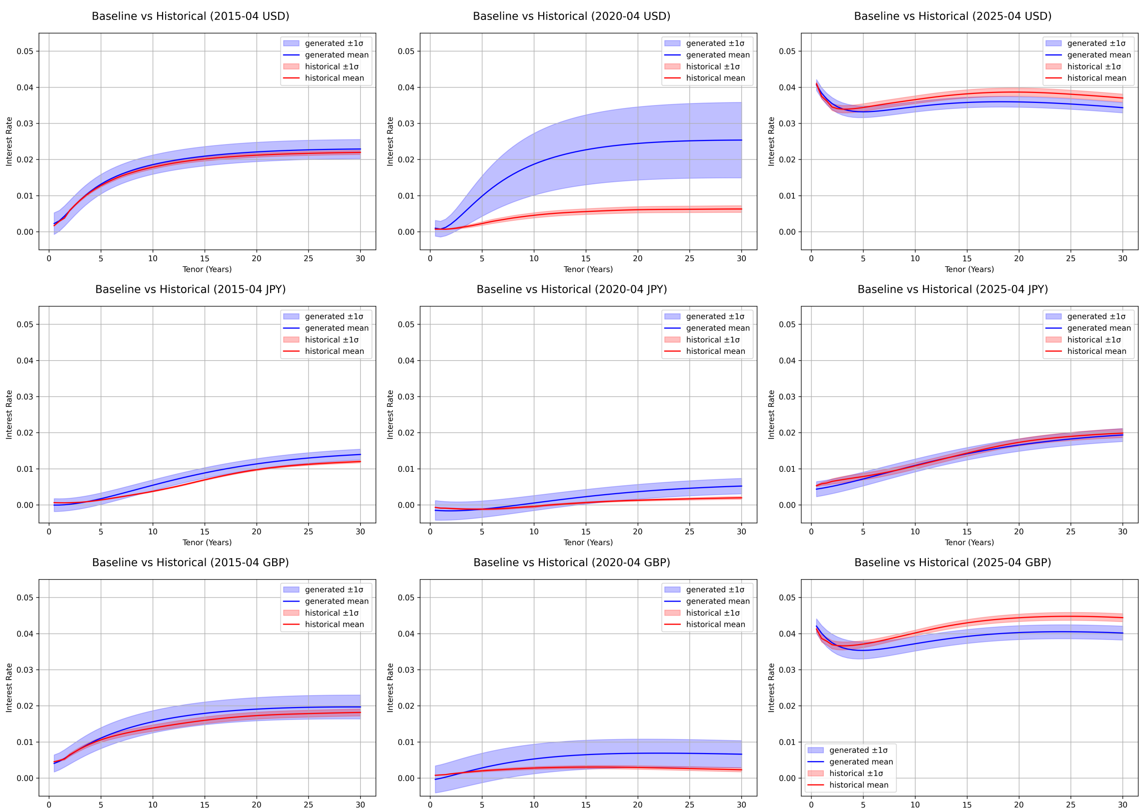

Figure 19: In-sample performance of NSS-baseline generation

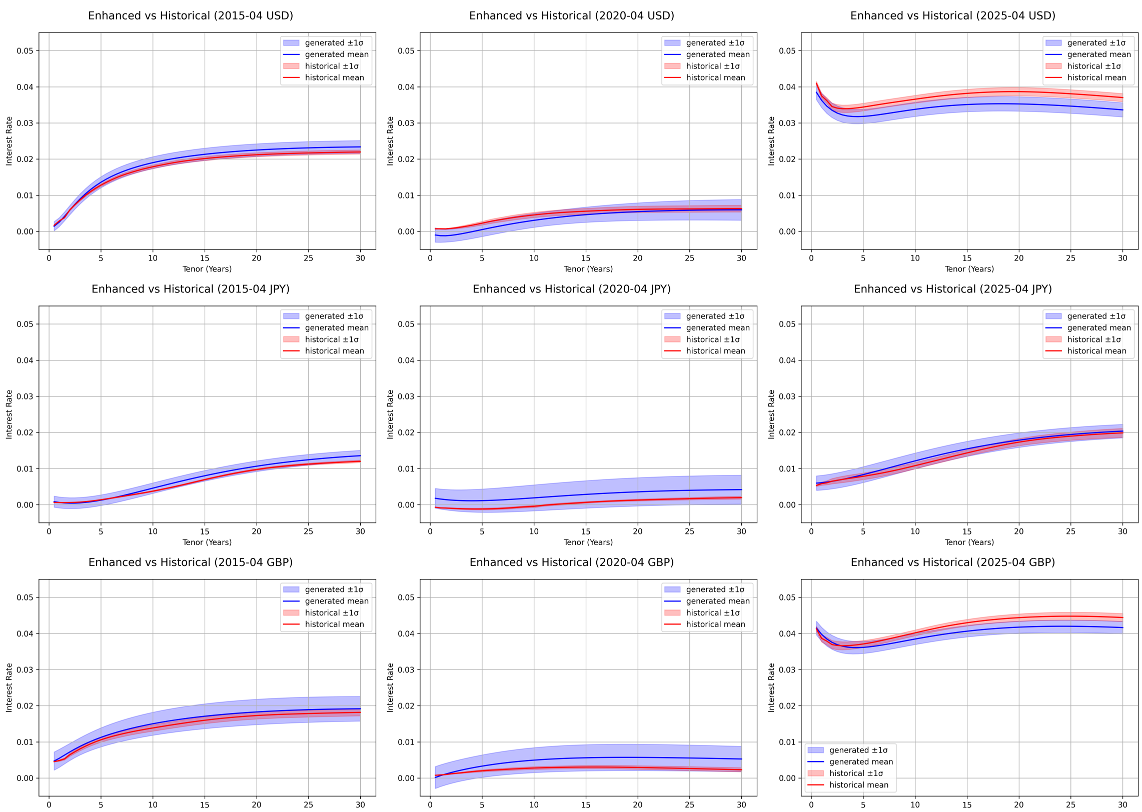

Figure 20: In-sample performance of NSS-Enhanced generation

6.4 Comparison of Performance with Other Models

We further conduct an out-of-sample experiment to benchmark the factor-based model against the direct generation model presented in Section 5. By generating NSS parameters and reconstructing the curves, we establish a framework for a direct comparative analysis between the two approaches.

Figures 21, 22, and 23 summarize the goodness-of-fit metrics of generation from both models, which are trained only on data up to the end of December 2023 against historical data. Figures 24 and 25 summarize the results of six-month-ahead generation in the same way as Section 5. The left two columns present the results of direct generation, while the right two columns show the results generated by the factor-based approach using the NSS model. The comparison indicates that the out-of-sample generation by the factor-based model is generally comparable to that of the direct model. Notably, for USD interest rates, under June 2024 condition the factor-based model even reproduced historical data more accurately than the direct model. On the other hand, for JPY interest rates, the generation results based on the conditional data of June 2024 and December 2024—scenarios assuming an upward interest rate environment—showed larger deviations than the direct model, with differences exceeding 1% relative to historical data. Consistent with Section 5, we also update the training dataset to construct synthetic yield curves six months ahead based on interest rate data available up to the present. The results are similar to those obtained with the direct model, confirming that both approaches can reproduce yield curves derived from historical data when generating approximately six months into the future.

Beyond the quality of the generated samples, the factor-based approach demonstrates a significant advantage in computational efficiency. By compressing the high-dimensional yield curve data into a low-dimensional latent space (the six NSS parameters), the model significantly reduces the computational burden. We observe that the training time required for the factor-based model is approximately 18 minutes, reduced by 40% compared to the direct generation model on the same GPU. The generation of 1,024 synthetic samples per condition with the factor-based model takes approximately 50 seconds, also reduced by 44% compared to the direct generation model on the same GPU. This efficiency gain makes the factor-based approach particularly scalable and practical for scenarios requiring frequent model retraining or rapid scenario generation.

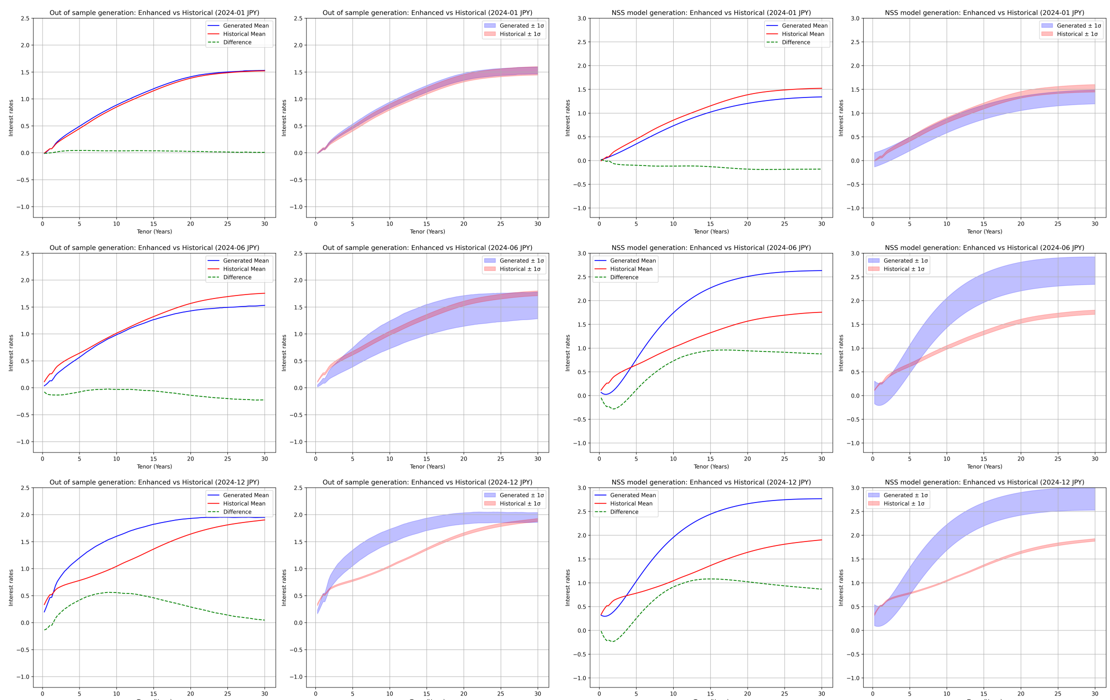

Figure 21: Out-of sample comparison of Enhanced generation, right two columns present the results from NSS factor-based approach: JPY

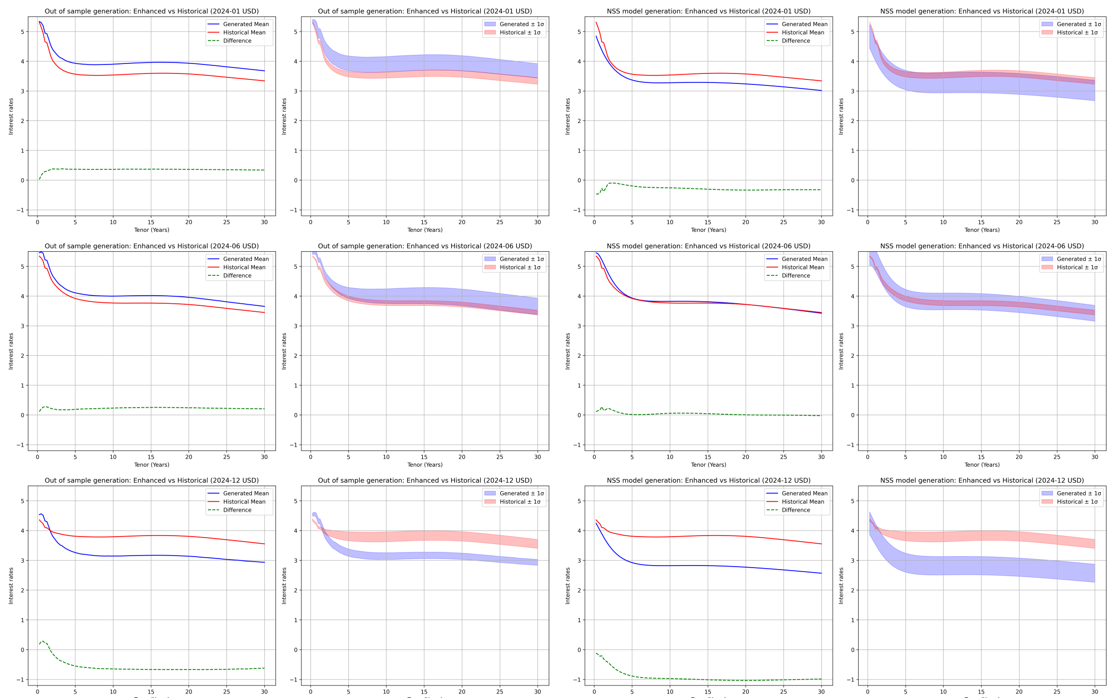

Figure 22: Out-of sample comparison of Enhanced generation, right two columns present the results from NSS factor-based approach: USD

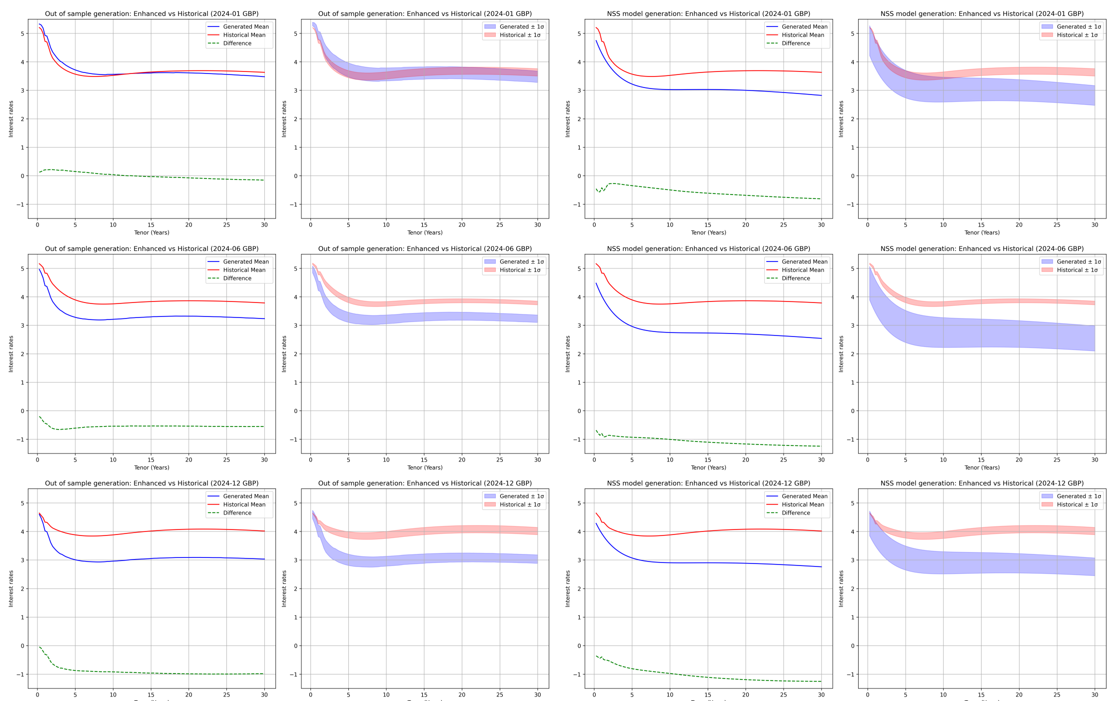

Figure 23: Out-of sample comparison of Enhanced generation,right two columns present the results from NSS factor-based approach: GBP

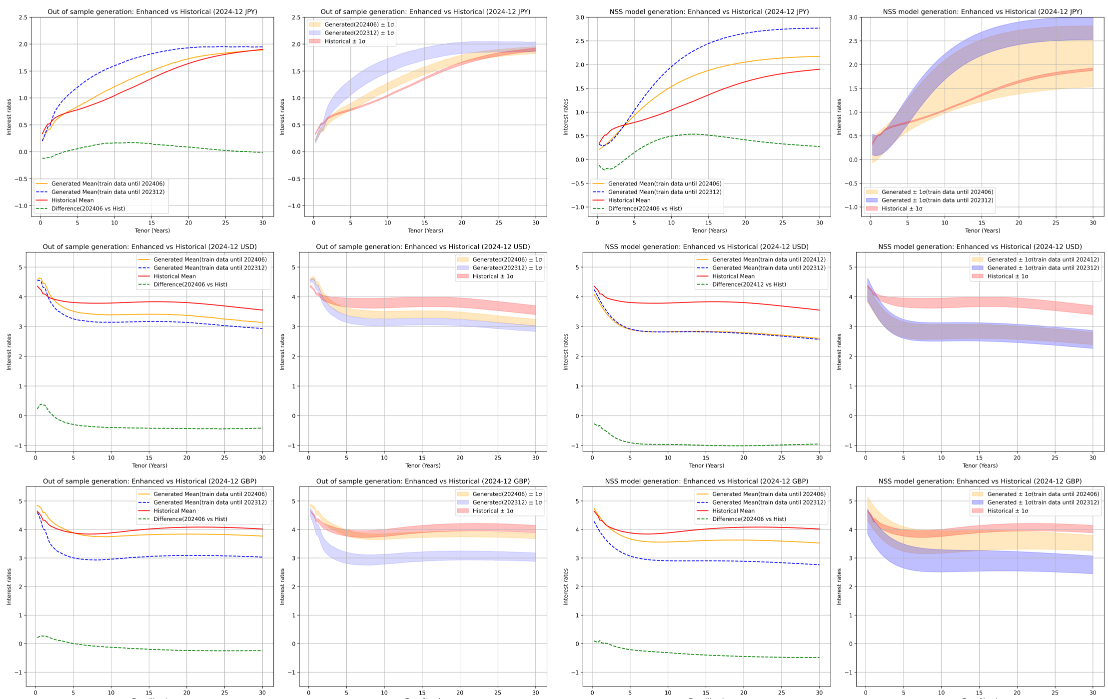

Figure 24: Six month ahead generation comparison,right two columns present the results from NSS factor-based approach: Enhanced model with data up to June 2024

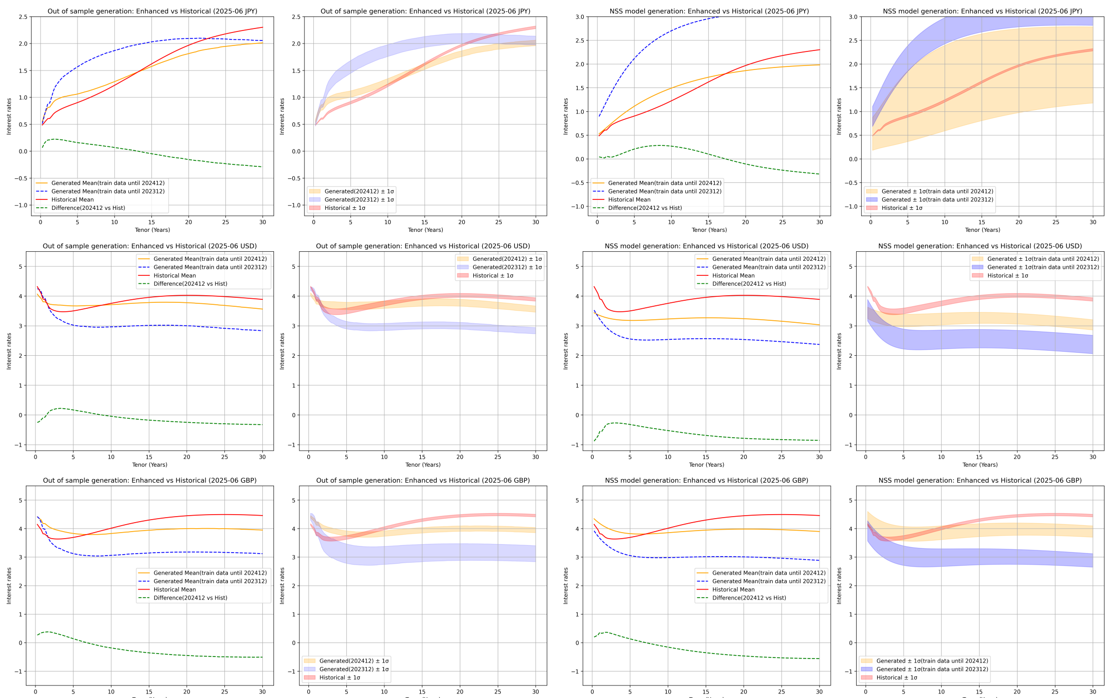

Figure 25: Six month ahead generation comparison,right two columns present the results from NSS factor-based approach: Enhanced model with data up to December 2024

7 Conclusion

In this study, we have proposed a conditional generative framework for synthesizing yield curves using diffusion models. To the best of our knowledge, this is the first application of diffusion models to the term structure of interest rates. By conditioning on macroeconomic indicators such as short-term interest rates and inflation, both the direct generation model and the difference-based learning model is able to generate realistic and historically consistent yield curves that reflect underlying economic environments. In addition, by assuming the Nelson-Siegel-Svensson (NSS) model as the functional form of the yield curve, we have demonstrated that realistic and historically consistent yield curves reflecting underlying economic environments can also be generated indirectly through the synthesis of NSS model factors.

This framework offers promising avenues for practical application. First, it can be utilized for scenario analysis: by specifying future economic scenarios and conditioning on the corresponding macroeconomic variables—predicted separately—the model can generate yield curves that support the formulation of bond investment strategies. Second, it can be applied to risk management: by conditioning on macroeconomic variables associated with stress scenarios, the model can simulate term structures to estimate potential losses and derive risk metrics. In both cases, the use of observable macroeconomic indicators and historical data as inputs enables intuitive and transparent analysis, which we consider a key advantage of the proposed approach. Third, as discussed in Section 4, when policy indicators such as quantitative easing metrics are used as conditioning variables, the model can be employed to assess the potential impact and effectiveness of policy measures prior to their implementation in actual markets. This opens up new possibilities for policy evaluation and macro-financial experimentation.

Future research will focus on expanding the set of conditioning variables to capture more nuanced economic scenarios, enhancing model robustness under extreme market conditions, and exploring practical applications in areas such as risk management, asset allocation and development of new interest rate-related financial products.

References

[1] Jonathan Ho, Ajay Jain, and Pieter Abbeel. *Denoising Diffusion Probabilistic Models.* arXiv preprint arXiv:2006.11239, 2020.

[2] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. *Score-Based Generative Modeling through Stochastic Differential Equations.* arXiv preprint arXiv:2011.13456, 2021.

[3] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. *High-Resolution Image Synthesis with Latent Diffusion Models.* arXiv preprint arXiv:2112.10752, 2022.

[4] Tim Salimans and Jonathan Ho. *Progressive Distillation for Fast Sampling of Diffusion Models.* arXiv preprint arXiv:2202.00512, 2022.

[5] Minshuo Chen, Renyuan Xu, Yumin Xu, and Ruixun Zhang. *Diffusion Factor Models: Generating High-Dimensional Returns with Factor Structure.* arXiv preprint arXiv:2504.06566, 2025.

[6] Yuki Tanaka, Ryuji Hashimoto, Takehiro Takayanagi, Zhe Piao, Yuri Murayama, and Kiyoshi Izumi. *CoFinDiff: Controllable Financial Diffusion Model for Time Series Generation.* arXiv preprint arXiv:2503.04164, 2025.

[7] `https://github.com/xymmmm00/diffusion_factor_model` Accessed: September 7, 2025.

[8] Jherek Healy. *Equivalence between Forward Rate Interpolations and Discount Factor Interpolations for Yield Curve Construction.* arXiv preprint arXiv:2005.13890, 2020.

[9] Nelson, R.C. and A.F. Siegel. *Parsimonious Modeling of Yield Curves.* Journal of Business, 60(4), 473-489.

[10] Svensson, L.E.O. *Estimating and Interpreting Forward Interest Rates: Sweden 1992-94.* CEPR Discussion Paper 1051.

[11] Atsushi Sekine. *Estimation of the Nelson–Siegel Model Using Daily Yield Data.* Public Policy Review, Policy Research Institute, Ministry of Finance, Japan, Vol.17, No.3, June 2021. (in Japanese).

Appendix A: List of obtained interest rate data

Table 9: Overview of Training Data:JPY

Tenor	Ticker	Start Date	End Date
3mo	JYSOC Curncy	2010/01/01	2025/07/02
6mo	JYSOF Curncy	2010/01/01	2025/07/02
9mo	JYSOI Curncy	2010/01/01	2025/07/02
1YR	JYSO1 Curncy	2010/01/01	2025/07/02
2YR	JYSO2 Curncy	2010/01/01	2025/07/02
3YR	JYSO3 Curncy	2010/01/01	2025/07/02
4YR	JYSO4 Curncy	2010/01/01	2025/07/02
5YR	JYSO5 Curncy	2010/01/01	2025/07/02
6YR	JYSO6 Curncy	2010/01/01	2025/07/02
7YR	JYSO7 Curncy	2010/01/01	2025/07/02
8YR	JYSO8 Curncy	2010/01/01	2025/07/02
9YR	JYSO9 Curncy	2010/01/01	2025/07/02
10YR	JYSO10 Curncy	2010/01/01	2025/07/02
11YR	JYSO11 Curncy	2011/02/01	2025/07/02
12YR	JYSO12 Curncy	2010/01/01	2025/07/02
15YR	JYSO15 Curncy	2010/01/01	2025/07/02
20YR	JYSO20 Curncy	2010/01/01	2025/07/02
25YR	JYSO25 Curncy	2010/01/01	2025/07/02
30YR	JYSO30 Curncy	2010/01/01	2025/07/02

Table 10: Overview of Training Data:USD

Tenor	Ticker	Start Date	End Date
3mo	USSOC Curncy	2012/10/01	2025/07/02
6mo	USSOF Curncy	2012/10/01	2025/07/02
9mo	USSOI Curncy	2012/10/01	2025/07/02
1YR	USSO1 Curncy	2012/10/01	2025/07/02
2YR	USSO2 Curncy	2012/10/01	2025/07/02
3YR	USSO3 Curncy	2012/10/01	2025/07/02
4YR	USSO4 Curncy	2012/10/01	2025/07/02
5YR	USSO5 Curncy	2012/10/01	2025/07/02
6YR	USSO6 Curncy	2012/10/01	2025/07/02
7YR	USSO7 Curncy	2012/10/01	2025/07/02
8YR	USSO8 Curncy	2012/10/01	2025/07/02
9YR	USSO9 Curncy	2012/10/01	2025/07/02
10YR	USSO10 Curncy	2012/10/01	2025/07/02
11YR	USSO11 Curncy	2024/07/22	2025/07/02

Tenor	Ticker	Start Date	End Date
12YR	USSO12 Curncy	2012/10/01	2025/07/02
15YR	USSO15 Curncy	2012/10/01	2025/07/02
20YR	USSO20 Curncy	2012/10/01	2025/07/02
25YR	USSO25 Curncy	2012/10/01	2025/07/02
30YR	USSO30 Curncy	2012/10/01	2025/07/02

Table 11: Overview of Training Data:GBP

Tenor	Ticker	Start Date	End Date
3mo	BPSWSC Curncy	2010/11/29	2025/07/09
6mo	BPSWSF Curncy	2010/11/29	2025/07/09
9mo	BPSWSI Curncy	2010/11/29	2025/07/09
1YR	BPSWS1 Curncy	2010/11/29	2025/07/09
2YR	BPSWS2 Curncy	2010/11/29	2025/07/09
3YR	BPSWS3 Curncy	2010/11/29	2025/07/09
4YR	BPSWS4 Curncy	2010/11/29	2025/07/09
5YR	BPSWS5 Curncy	2010/11/29	2025/07/09
6YR	BPSWS6 Curncy	2010/11/29	2025/07/09
7YR	BPSWS7 Curncy	2010/11/29	2025/07/09
8YR	BPSWS8 Curncy	2010/11/29	2025/07/09
9YR	BPSWS9 Curncy	2010/11/29	2025/07/09
10YR	BPSWS10 Curncy	2010/11/29	2025/07/09
12YR	BPSWS12 Curncy	2010/11/29	2025/07/09
15YR	BPSWS15 Curncy	2010/11/29	2025/07/09
20YR	BPSWS20 Curncy	2010/11/29	2025/07/09
25YR	BPSWS25 Curncy	2010/11/29	2025/07/09
30YR	BPSWS30 Curncy	2010/11/29	2025/07/09

Appendix B: List of obtained conditon data

Table 12: Overview of Training Data:USD

Indicator	Ticker	Detail	Transformation process
ST-rate	JYMUON Curncy	Overnight interest rate of JPY	None
ST-rate	USSO1Z Curncy	1week OIS rate of USD	None
ST-rate	SONIO/N Index	Overnight interest rate of GBP	None
Inflation rate	JNCPIYOY Index	CPI Index of Japan (yoy %)	None
Inflation rate	CPI YOY Index	CPI Index of U.S. (yoy %)	None
Inflation rate	UKRPCJYR Index	CPI Index of U.K. (yoy %)	None
QE Indicator	aJPBOJAGOVB[1]	Japan BOJ accounts, assets, Japan government bonds, JPY	Adjusted to yoy % terms
QE Indicator	aUSRHTBA[1 2]	U.S. Factors Affecting Reserve Balances of Depository Institutions, Reserve bank credit, Securities held	Adjusted to yoy % terms
QE Indicator	aGBBLTEA[1]	U.K. Holdings of gilts by the BOE asset purchase facility, GBP	Adjusted to yoy % terms
Monetary Base	JNMBYOY Index	Monetary Base of Japan(yoy %)	None
Monetary Base	ARDIMOYY Index	Monetary Base of U.S.(yoy %)	None
Monetary Base	UKMSM41Y Index	M4 Money Supply of U.K. (yoy %, sa)[3]	None
CA to GDP	EHCAJP Index	Current Account % GDP of Japan	None
CA to GDP	EHCAUS Index	Current Account % GDP of U.S.	None
CA to GDP	EHCAGB Index	Current Account % GDP of U.K.	None
Nominal GDP	ECOXJPS Index	Nominal GDP of Japan (USD bn)	Scaling(1/1,000)
Nominal GDP	GDP CUR$ Index	Nominal GDP of U.S. (USD bn)	Scaling(1/1,000)
Nominal GDP	ECOXUKS Indexy	Nominal GDP of U.K. (USD bn)	Scaling(1/1,000)

[1]Data from LSEG Datastream
[2]Missing values in the data sources are supplemented using publicly available data from central banks.
[3]Due to data availability issues, alternative data is used.