

Article

Not peer-reviewed version

A Controlled Comparison of Deep Learning Architectures for Multi-Horizon Financial Forecasting Evidence from 918 Experiments

[Nabeel Ahmad Saidd](#) *

Posted Date: 2 March 2026

doi: 10.20944/preprints202603.0066.v1

Keywords: financial time-series forecasting; deep learning; model benchmarking; multi-horizon forecasting; hyperparameter optimisation; reproducibility; seed robustness; directional accuracy



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

A Controlled Comparison of Deep Learning Architectures for Multi-Horizon Financial Forecasting Evidence from 918 Experiments

Nabeel Ahmad Saidd

Dr. A.P.J. Abdul Kalam Technical University (AKTU); nabeelahmadsaidd@gmail.com

Abstract

Multi-horizon price forecasting is central to portfolio allocation, risk management, and algorithmic trading, yet deep learning architectures have proliferated faster than rigorous comparisons on financial data can assess them. Existing benchmarks are weakened by uncontrolled hyperparameter budgets, single-seed evaluation, narrow asset-class coverage, and absent pairwise statistical correction. This study compares nine deep learning architectures—Autoformer, DLinear, iTransformer, LSTM, ModernTCN, N-HiTS, PatchTST, TimesNet, and TimeXer—from four families (Transformer, MLP, CNN, RNN) across three asset classes (cryptocurrency, forex, equity indices) and two forecasting horizons ($h \in \{4, 24\}$ hours). All 918 runs follow a five-stage protocol: fixed-seed Bayesian hyperparameter optimisation, configuration freezing per asset class, multi-seed final training, metric aggregation with uncertainty quantification, and statistical validation. ModernTCN achieves the best mean rank (1.333) with a 75% first-place rate across 24 evaluation points; PatchTST ranks second (2.000). The global rank leaderboard reveals a clear three-tier structure separating the top pair from a middle cluster and a bottom tier. Architecture explains 99.90% of raw RMSE variance versus 0.01% for seed randomness. Rankings remain stable across horizons despite 2–2.5 \times error amplification. Directional accuracy is indistinguishable from 50% across all 54 model–category–horizon combinations, indicating that MSE-trained architectures lack directional skill at hourly resolution. Four practical implications follow: (i) large-kernel temporal convolutions and patch-based Transformers consistently outperform alternatives; (ii) the complexity–performance relationship is non-monotonic, with architectural inductive bias mattering more than raw capacity; (iii) three-seed replication suffices given negligible seed variance; and (iv) directional forecasting requires explicit loss-function redesign. The full codebase, data, trained models, and evaluation outputs are released for independent replication.

Keywords: financial time-series forecasting; deep learning; model benchmarking; multi-horizon forecasting; hyperparameter optimisation; reproducibility; seed robustness; directional accuracy

JEL Classification: C45; C52; C53; G17

1. Introduction

1.1. Motivation

Multi-horizon price forecasting is central to modern finance: portfolio allocation relies on expected return estimates at multiple horizons, risk management depends on volatility projections, and algorithmic trading requires predictive signals whose quality degrades with the forecast window. Unlike physical systems governed by conservation laws, financial time series exhibit non-stationarity, heavy-tailed returns, volatility clustering, leverage effects, and abrupt regime transitions [1,2]—properties that make accurate multi-step prediction exceptionally difficult.

Deep learning architectures for temporal sequence modelling have proliferated rapidly over the past five years. Transformer-based models now span a wide range of inductive biases: auto-correlation in the frequency domain [3], patch-based tokenisation with channel independence [4],

inverted variate-wise attention [5], and exogenous-variable-aware cross-attention with learnable global tokens [6]. At the same time, decomposition-based linear mappings have been shown to match or surpass Transformer performance on standard benchmarks [7], while modern temporal convolutional architectures exploit large-kernel depthwise convolutions [8] and FFT-based 2D reshaping [9] to capture multi-scale temporal structure. Hierarchical MLP designs with multi-rate pooling [10] offer yet another approach to direct multi-step forecasting.

This architectural diversity raises a practical question: *which architecture should a practitioner deploy for a given financial forecasting task, at which horizon, and for which asset class?* A reliable answer requires controlled experimentation that isolates architectural merit from confounding factors—a requirement that existing benchmarks have not met, as Section 2.4 demonstrates.

1.2. Research Gap

Despite the growing body of forecasting studies, five persistent methodological shortcomings undermine published model comparisons:

- G1. Uncontrolled hyperparameter budgets.** Many benchmarks allocate different tuning effort to different models—or skip hyperparameter optimisation entirely—confounding tuning luck with architectural merit. Prior work has introduced architectures with custom-tuned configurations while evaluating competitors under default or unspecified settings [7,9], preventing fair attribution of performance differences.
- G2. Single-seed evaluation.** The vast majority of comparative studies report results from a single random initialisation. Seed-induced variance has been shown to exceed algorithmic differences in standard machine learning benchmarks [11], with analogous effects documented in reinforcement learning [12]. Without multi-seed replication, it is impossible to distinguish genuine architectural advantage from stochastic variation in weight initialisation.
- G3. Single-horizon analysis.** Most comparisons evaluate a single forecasting horizon, precluding investigation of how architectural inductive biases interact with prediction difficulty as the forecast window extends. Prior work has benchmarked recurrent networks at fixed horizons without characterising degradation behaviour, leaving cross-horizon generalisation unaddressed [13].
- G4. Absent pairwise statistical correction.** Even studies that report omnibus significance tests (e.g., Friedman) rarely apply post-hoc pairwise corrections with family-wise error control, leaving the significance of individual ranking differences unquantified. The M4 competition [14] provided aggregate rankings but did not report pairwise tests among deep learning participants.
- G5. Narrow asset-class coverage.** Benchmarks focusing on a single asset class—whether cryptocurrency [15], equities, or standard forecasting datasets (ETTh, Weather, Electricity)—cannot assess whether ranking conclusions generalise across the structurally distinct dynamics of different financial markets. Cross-class evaluation is necessary for reliable deployment guidance.

Collectively, these gaps prevent reliable inference about the relationship between architectural inductive biases and financial time-series structure, and deprive practitioners of evidence-based model selection guidance—despite the finding (Section 4.4) that architecture choice, not random initialisation, explains over 99% of total forecast variance.

1.3. Hypotheses and Objectives

Four testable hypotheses structure the empirical investigation. Each states a clear null and a designated statistical test:

- H1. Ranking Non-Uniformity.** The global performance ranking of the nine architectures is significantly non-uniform across evaluation points.
 H_0 : All models have equal expected rank.

Test: Rank-based leaderboard analysis across 24 evaluation points (12 assets \times 2 horizons); win-rate counts; mean rank gaps between tiers.

Evidence: Sections 4.1 and 4.6.

H2. Cross-Horizon Ranking Stability. Top-ranked architectures at $h = 4$ maintain their relative superiority at $h = 24$, despite absolute error amplification.

H_0 : Rankings at $h = 4$ and $h = 24$ are independent ($\rho_S = 0$).

Test: Spearman rank correlation between $h = 4$ and $h = 24$ model rankings per asset; percentage degradation analysis.

Evidence: Section 4.3.

H3. Variance Dominance. Architecture choice explains a significantly larger proportion of total forecast variance than random seed initialisation.

H_0 : Model and seed factors contribute equally to variance.

Test: Two-factor sum-of-squares variance decomposition; comparison of variance proportions across panels (raw, z-normalised all models, z-normalised modern only).

Evidence: Section 4.4.

H4. Non-Monotonic Complexity–Performance. The relationship between trainable parameter count and forecasting error is non-monotonic—architectural inductive bias matters more than raw model capacity.

H_0 : Monotonic negative correlation ($\rho_S = -1$) between parameter count and RMSE rank.

Test: Spearman rank correlation between parameter count and mean RMSE rank; visual inspection of complexity–performance scatter.

Evidence: Section 4.7.

The primary objective is to provide a controlled, statistically validated comparison of nine deep learning architectures across three asset classes at two forecasting horizons under identical experimental conditions, yielding evidence-based deployment guidance.

1.4. Contributions

Five specific contributions advance the state of knowledge:

- C1. Protocol-controlled fair comparison (addresses G1).** Nine architectures spanning four families (Transformer, MLP, CNN, RNN) are evaluated under a single five-stage protocol: fixed-seed Bayesian HPO (5 Optuna TPE trials, seed 42), configuration freezing per asset class, identical chronological 70/15/15 splits, a common OHLCV feature set, and rank-based performance evaluation across 12 instruments, three asset classes, and two horizons, totalling 918 runs (648 final training runs + 270 HPO trials).
- C2. Multi-seed robustness quantification (addresses G2).** Final training is replicated across three independent seeds (123, 456, 789). A two-factor variance decomposition shows that architecture choice explains 99.90% of total forecast variance while seed variation accounts for 0.01%, establishing model selection as the dominant lever for accuracy improvement and confirming that three-seed replication is sufficient.
- C3. Cross-horizon generalisation analysis (addresses G3).** Identical models are evaluated at $h = 4$ and $h = 24$ with matched protocol, characterising architecture-specific degradation (Table 10) and identifying which inductive biases scale with prediction difficulty.
- C4. Asset-class-specific deployment guidance (addresses G5).** Category-level analysis across cryptocurrency, forex, and equity indices (Table 9) shows that top-tier rankings (ModernTCN, PatchTST) hold across all three classes, while mid-tier orderings (DLinear, N-HiTS, TimeXer) are category-dependent. A per-asset best-model matrix (Table 8) further reveals niche advantages: N-HiTS achieves the lowest error on lower-capitalisation cryptocurrency assets, providing actionable asset-level deployment guidance.
- C5. Open, deterministic benchmarking framework (addresses G1–G5).** The complete pipeline—source code, configuration files, raw market data, processed datasets, all trained models, and

evaluation outputs—is released under an open licence. Accompanying documentation enables any researcher to reproduce every experiment via a unified command-line interface.

1.5. Paper Organisation

The remainder of this paper is organised as follows. Section 2 surveys related work and positions this study relative to existing benchmarks. Section 3 presents the unified experimental protocol. Section 4 reports the empirical findings structured by the four hypotheses. Section 5 interprets the results, provides economic context, and discusses limitations. Section 6 summarises contributions and offers deployment recommendations. Section 7 provides a reproducibility statement. Supplementary results are collected in the Appendix.

2. Related Work

This section positions the present study within the broader landscape of financial time-series forecasting. It traces the evolution from classical econometric models through four families of deep learning architectures, reviews multi-step forecasting strategies, and identifies the methodological gaps this benchmark addresses.

2.1. Classical and Statistical Approaches

The ARIMA family [16] remains a cornerstone of time-series analysis, capturing linear temporal dependencies through differencing and lagged-error terms. Exponential smoothing methods [17] offer computationally efficient trend–seasonality decomposition, while the ARCH [18] and GARCH [19] frameworks provide the standard toolkit for modelling conditional heteroscedasticity in financial returns.

These methods share three fundamental limitations. First, they assume linearity in the conditional mean or variance, yet financial returns exhibit nonlinear phenomena—leverage effects, long memory, and regime transitions [1]—that violate these assumptions. Second, classical models are inherently univariate (or require explicit cross-variable specification), limiting their ability to exploit joint OHLCV information. Third, multi-step forecasting under these frameworks typically proceeds recursively, compounding prediction errors at longer horizons [20]. These limitations motivate the use of deep learning architectures that learn nonlinear, multivariate mappings directly from data.

2.2. Deep Learning Architectures for Time-Series

Deep learning approaches to time-series forecasting have evolved along four architectural families, each encoding distinct assumptions about temporal structure. The specific models included in this benchmark are reviewed below, with emphasis on their temporal inductive biases.

Recurrent architectures.

Long Short-Term Memory networks [21,22] introduced gated recurrence to address vanishing gradients, maintaining a cell state that selectively retains or discards information. While effective for short-range dependencies, sequential processing limits parallelisation and hinders learning of very long-range patterns. Surveys confirm that LSTMs served as the default neural forecasting baseline during 2017–2021 [13,23]. In this benchmark, LSTM represents the recurrent family, providing a classical reference point. Its inductive bias is autoregressive: the hidden state compresses all past information into a fixed-dimensional vector, relying on recurrence to propagate long-range context.

Transformer-based architectures.

The self-attention mechanism [24] enables direct modelling of pairwise dependencies across arbitrary temporal lags, overcoming the sequential bottleneck of recurrence. Four Transformer variants are evaluated:

- Autoformer [3] replaces canonical attention with an auto-correlation mechanism operating in the frequency domain at $\mathcal{O}(L \log L)$ complexity, coupled with progressive trend–seasonal de-

composition. Its inductive bias assumes that dominant temporal patterns manifest as periodic auto-correlations detectable via spectral analysis.

- PatchTST [4] segments input series into patches, treats each as a token, and applies a Transformer encoder with channel-independent processing and RevIN normalisation [25] for distribution-shift mitigation. Its inductive bias prioritises local temporal coherence within patches while using attention for global inter-patch dependencies.
- iTransformer [5] inverts the attention paradigm: each variate's full temporal trajectory serves as a token, and attention operates across the variate dimension, directly capturing cross-variable interactions. Its inductive bias assumes that inter-variate relationships are the primary source of predictive information.
- TimeXer [6] separates target and exogenous variables, embedding the patched target alongside learnable global tokens and applying cross-attention to query inverted exogenous representations. Its inductive bias explicitly separates autoregressive dynamics from exogenous covariate influence.

MLP and linear architectures.

The necessity of attention for time-series forecasting has been challenged [7], demonstrating that DLinear—a decomposition-based model with dual independent linear layers mapping seasonal and trend components, without hidden layers, activations, or attention—can match or exceed Transformer performance on standard benchmarks. Its inductive bias assumes that temporal patterns are adequately captured by linear projections of decomposed components.

N-HiTS [10] employs a hierarchical stack of MLP blocks with multi-rate pooling: each block operates at a different temporal resolution, produces coefficients via a multi-layer perceptron, and interpolates them to the forecast horizon through basis-function expansion. Its inductive bias prioritises multi-scale temporal structure through hierarchical signal decomposition. Together, these results raise the question of whether attention contributes meaningfully to forecasting accuracy—a question this benchmark addresses.

Convolutional architectures.

Temporal convolutional networks (TCNs) apply causal, dilated convolutions to capture long-range dependencies through hierarchical receptive fields [26]. The two convolutional models in this benchmark—TimesNet and ModernTCN—each depart from the standard TCN template in distinct ways:

TimesNet [9] transforms the forecasting problem from 1D to 2D by identifying dominant FFT-based periods, reshaping the sequence into 2D tensors indexed by period length and intra-period position, and applying Inception-style 2D convolutions. Its inductive bias assumes that temporal dynamics decompose into inter-period and intra-period variations best captured through spatial convolutions.

ModernTCN [8] employs large-kernel depthwise convolutions with structural reparameterisation (dual branches merged at inference), multi-stage downsampling, and optional RevIN normalisation. Its inductive bias holds that local temporal patterns at multiple scales, captured through large receptive fields with efficient depthwise operations, suffice for accurate forecasting without frequency-domain or attention mechanisms.

A key question is whether these architectural differences yield *consistent* and *statistically significant* performance differences on financial data, or whether the experimental protocol dominates observed rankings. The present study disentangles these effects through the controlled protocol described in Section 3.

2.3. Multi-Step Forecasting Strategies

Multi-step-ahead prediction admits three principal strategies [20,27]. The *recursive* (iterated) strategy applies a one-step model iteratively, feeding predictions back as inputs; this approach is

straightforward but accumulates errors geometrically with the horizon length. The *direct* strategy trains independent output heads for each future step, avoiding error accumulation at the cost of ignoring inter-step temporal coherence. The *multi-input multi-output* (MIMO) strategy produces all horizon steps in a single forward pass, preserving inter-step dependencies without iterative error propagation.

This benchmark adopts *direct multi-step forecasting*: each model outputs all h forecast steps simultaneously ($h \in \{4, 24\}$) in a single forward pass, matching the native output design of all nine architectures. No model feeds predictions back as inputs. Two separate experiments per horizon use distinct lookback windows ($w = 24$ for $h = 4$; $w = 96$ for $h = 24$), enabling isolation of horizon-dependent degradation from architecture-dependent effects. Three considerations motivate this choice: (i) it avoids the error-accumulation confound of recursive strategies; (ii) it matches every architecture's native mode, preventing protocol mismatch; and (iii) it enables clean cross-horizon comparison (H2) by ensuring that $h = 4$ and $h = 24$ results differ only in task difficulty.

2.4. Benchmarking Practices and Identified Gaps

Several prior studies have compared deep learning architectures for time-series forecasting, but persistent methodological limitations constrain the conclusions that can be drawn. The most relevant benchmarks are reviewed below, mapped to the five gaps from Section 1.2.

Large-scale forecasting competitions have advanced standardised evaluation methodologies. The M4 competition [14] introduced a common evaluation protocol across 100,000 series but focused on macroeconomic and demographic data, did not include modern deep learning architectures (released after 2020), and did not apply pairwise statistical corrections (G1, G4, G5). The M5 competition [28] extended to retail sales forecasting with hierarchical structure but again excluded recent Transformer, TCN, and MLP architectures (G5).

Within the deep learning literature, a comprehensive benchmark of recurrent networks [13] excluded Transformer- and CNN-based alternatives and evaluated a single horizon (G3, G5). The competitiveness of linear models against Transformers was demonstrated [7], but with default hyperparameters for competitors and a single seed (G1, G2). TimesNet was benchmarked against multiple baselines [9] but with uncontrolled HPO budgets and single-seed evaluation (G1, G2). PatchTST was introduced with strong benchmark results [4] but on non-financial datasets and without multi-seed replication (G2, G5). iTransformer was evaluated across standard time-series benchmarks [5] but without pairwise statistical tests or multi-horizon degradation analysis (G3, G4). Within the financial forecasting literature specifically, a comprehensive survey [15] noted the absence of controlled experimental comparisons without providing one.

Table 1 summarises these prior studies and their gap coverage. The central finding is that *no prior study simultaneously addresses G1–G5*: every existing benchmark leaves at least two gaps uncontrolled. The present study fills this compound gap by providing the first controlled, multi-seed, multi-horizon, multi-asset-class comparison with full statistical validation for financial time-series forecasting.

Table 1. Summary of prior comparative studies in time-series forecasting. Columns indicate the number of models evaluated, number of datasets or asset classes, horizons tested, whether multi-seed evaluation was performed, and whether post-hoc pairwise statistical tests were applied.

Study	Models	Datasets	Horizons	Multi-Seed	Pairwise Tests	Open Code
[14] (M4)	Many	100K series	Multiple	No	No	Partial
[13]	RNN only	6	Multiple	No	No	Partial
[7]	6	9	4	No	No	Yes
[9]	8	8	4	No	No	Yes
[4]	7	8	4	No	No	Yes
[5]	8	7	4	No	No	Yes
Present study	9	12 (3 classes)	2	Yes (3)	Yes	Yes

3. Experimental Design

This section presents the complete experimental protocol as a unified, replicable specification. Every design choice is stated with its rationale. A reader equipped with the accompanying repository can reproduce every reported number by following this section sequentially.

3.1. Formal Problem Definition

Let $\mathbf{X}_t \in \mathbb{R}^{w \times d}$ denote a multivariate input window of w consecutive hourly observations with $d = 5$ OHLCV features (Open, High, Low, Close, Volume), ending at time index t . The forecasting task is to learn a parametric mapping

$$f_{\theta} : \mathbb{R}^{w \times d} \rightarrow \mathbb{R}^h, \quad \hat{\mathbf{y}}_{t+1:t+h} = f_{\theta}(\mathbf{X}_t), \quad h \in \{4, 24\}, \quad (1)$$

where $\hat{\mathbf{y}}_{t+1:t+h}$ is the predicted vector of future Close prices and θ collects all learnable parameters. Two horizon configurations are evaluated as *completely separate experiments*, with no shared weights or intermediate results:

- **Short-term** ($h = 4$): lookback window $w = 24$ hours, predicting 4 hours ahead.
- **Long-term** ($h = 24$): lookback window $w = 96$ hours, predicting 24 hours ahead.

All models employ *direct multi-step forecasting*: the entire horizon vector $\hat{\mathbf{y}} \in \mathbb{R}^h$ is produced in a single forward pass. No model feeds predictions back as inputs, avoiding the error accumulation of recursive strategies [20,27]. The training objective is mean squared error:

$$\mathcal{L}(\theta) = \frac{1}{n} \sum_{i=1}^n \|\mathbf{y}_i - \hat{\mathbf{y}}_i\|^2, \quad (2)$$

and the model selection criterion is minimum validation MSE. This common loss function and selection criterion apply identically to all architectures, eliminating confounds from differential training objectives.

3.2. Data

3.2.1. Asset Universe

The benchmark spans 12 financial instruments across three asset classes, each with structurally distinct market microstructure:

- **Cryptocurrency** (4 assets): BTC/USDT, ETH/USDT, BNB/USDT, ADA/USDT. These instruments trade continuously (24/7), exhibit high volatility, and are subject to rapid regime changes driven by speculative activity and regulatory events.
- **Forex** (4 assets): EUR/USD, USD/JPY, GBP/USD, AUD/USD. Major currency pairs are characterised by high liquidity, short-term mean-reverting tendencies, and sensitivity to macroeconomic announcements and central-bank policy.
- **Equity indices** (4 assets): Dow Jones, S&P 500, NASDAQ 100, DAX. These indices track broad equity markets, exhibiting trending behaviour, lower intra-day volatility relative to cryptocurrency, and session-based trading hours.

All instruments are sampled at H1 (1-hour) frequency, providing uniform temporal resolution across asset classes. For hyperparameter optimisation, one representative asset per class is designated: BTC/USDT (cryptocurrency), EUR/USD (forex), and Dow Jones (equity indices). Optimised configurations are frozen and applied to all assets within the corresponding class, preventing asset-level overfitting while preserving category-level calibration (Section 3.4.2).

3.2.2. Feature Specification

All models receive identical input tensors comprising five raw market features: Open, High, Low, Close, and Volume (OHLCV). This deliberate restriction to unprocessed market data isolates the contribution of architectural design from feature-engineering confounds. No technical indicators,

lagged returns, calendar variables, or external covariates are introduced. The sole forecast target is the Close price at each horizon step, ensuring that performance differences reflect architecture rather than feature availability.

3.2.3. Preprocessing

The preprocessing pipeline transforms raw market data into normalised, windowed tensors through four stages:

1. **Loading.** Raw hourly OHLCV records are ingested from CSV files containing datetime, open, high, low, close, and volume columns.
2. **Truncation.** To ensure comparable dataset sizes across instruments with different histories, the most recent 30,000 time steps are retained for every asset prior to windowing.
3. **Normalisation.** Standard z-score normalisation (zero mean, unit variance) is fitted *exclusively on the training partition* and applied unchanged to validation and test partitions. This prevents leakage of future distributional information into the scaling statistics. After inference, predictions are inverse-scaled to the original price domain before metric computation.
4. **Windowing.** Rolling windows of length $w + h$ are constructed. Two configurations are employed: $(w, h) = (24, 4)$ for short-term and $(w, h) = (96, 24)$ for long-term forecasting. Each window yields an input matrix $\mathbf{X} \in \mathbb{R}^{w \times 5}$ and a target vector $\mathbf{y} \in \mathbb{R}^h$ (Close prices only).

3.2.4. Chronological Splits

All partitions are strictly chronological to prevent future-data leakage:

- **Training:** first 70% of samples (approximately 21,000 windows per asset per horizon).
- **Validation:** next 15% of samples (approximately 4,500 windows).
- **Test:** final 15% of samples (approximately 4,500 windows).

No shuffling is performed at any stage, preserving the temporal ordering essential for financial time-series. Identical splits are applied to all models, ensuring that every architecture receives exactly the same training, validation, and test observations for each (asset, horizon) pair. Split boundaries and sample counts are recorded alongside the processed datasets (Table 2). Return distributional statistics—mean return, volatility, skewness, excess kurtosis, first-order autocorrelation, and ADF unit-root test p -value—for all twelve assets are reported in Table 3; all return series are stationary ($p < 0.001$) with heavy tails (excess kurtosis 15–96) and near-zero mean returns consistent with weak-form market efficiency.

Table 2. Dataset summary. All assets use H1 (hourly) frequency. The most recent 30,000 windowed samples are retained per (asset, horizon) pair, split chronologically into 70%/15%/15% train/val/test partitions. Window lengths: $w = 24$ for $h = 4$ and $w = 96$ for $h = 24$. Features: OHLCV (5 channels); target: close price.

Category	Asset	Date Range	Train	Val	Test	Total
Crypto	BTC/USDT	2021-03 – 2026-02	21,000	4,500	4,500	30,000
	ETH/USDT	2021-03 – 2026-02	21,000	4,500	4,500	30,000
	BNB/USDT	2021-03 – 2026-02	21,000	4,500	4,500	30,000
	ADA/USDT	2021-05 – 2026-02	21,000	4,500	4,500	30,000
Forex	EUR/USD	2017-12 – 2026-02	21,000	4,500	4,500	30,000
	USD/JPY	2017-12 – 2026-02	21,000	4,500	4,500	30,000
	GBP/USD	2017-12 – 2026-02	21,000	4,500	4,500	30,000
	AUD/USD	2017-12 – 2026-02	21,000	4,500	4,500	30,000
Indices	Dow Jones	2019-01 – 2026-02	21,000	4,500	4,500	30,000
	S&P 500	2019-01 – 2026-02	21,000	4,500	4,500	30,000
	NASDAQ 100	2019-01 – 2026-02	21,000	4,500	4,500	30,000
	DAX	2019-02 – 2026-02	21,000	4,500	4,500	30,000
Total per horizon		—	252,000	54,000	54,000	360,000

Table 3. Distributional statistics of hourly log returns for all twelve assets. All series reject the Augmented Dickey–Fuller unit-root null at $p < 0.001$, confirming return stationarity.

Cat.	Asset	$\bar{\mu}_r (\times 10^{-4})$	σ_r (%)	Skew	Kurt	ACF(1)	ADF p
Crypto	BTC/USDT	+0.4	0.783	-0.47	+35.8	-0.043	<0.001
	ETH/USDT	+0.3	0.973	-0.56	+27.3	-0.026	<0.001
	BNB/USDT	+0.8	1.107	+0.08	+35.0	-0.077	<0.001
	ADA/USDT	0.0	1.187	-0.05	+19.2	-0.085	<0.001
Forex	EUR/USD	0.0	0.109	-0.01	+15.2	-0.010	<0.001
	USD/JPY	+0.1	0.119	-0.98	+46.6	-0.002	<0.001
	GBP/USD	0.0	0.115	-1.66	+95.8	-0.018	<0.001
	AUD/USD	0.0	0.141	-0.07	+17.9	-0.020	<0.001
Indices	Dow Jones	+0.2	0.225	-0.57	+60.8	-0.012	<0.001
	S&P 500	+0.2	0.235	-1.13	+90.2	-0.026	<0.001
	NASDAQ 100	+0.3	0.286	-0.70	+47.7	-0.014	<0.001
	DAX	+0.2	0.262	-1.74	+83.0	-0.011	<0.001

All return distributions exhibit excess kurtosis (heavy tails) consistent with the stylised facts of financial markets; equities and forex show negative skewness (downside asymmetry), while crypto skewness is near zero. Negative ACF(1) values indicate mild mean-reversion in hourly returns. Volatility spans two orders of magnitude across asset classes (crypto $\approx 1\%$, forex $\approx 0.1\%$, indices $\approx 0.2\%$), reflecting differences in leverage, liquidity, and trading hours.

Representative Hourly Close Price Series (Latest 30,000 Observations)




Figure 1. Representative hourly Close-price time series for one asset per class: BTC/USDT (cryptocurrency), EUR/USD (forex), and Dow Jones (equity indices). Vertical dashed lines indicate chronological train/validation/test boundaries (70/15/15 split). The three classes exhibit qualitatively different dynamics: high-volatility trending behaviour (cryptocurrency), low-volatility mean-reversion around a narrow range (forex), and moderate-volatility upward drift (equity indices). All series comprise the most recent 30,000 hourly observations.

3.3. Model Architectures

Nine architectures spanning four families are evaluated. All models conform to a unified interface: input shape (B, w, d) with $d = 5$ OHLCV features; output shape (B, h) , where B is the batch size. No

model-specific feature engineering or data augmentation is permitted. Table 4 summarises the key architectural properties.

Transformer family (4 models).

Autoformer [3] replaces self-attention with an auto-correlation mechanism operating in the frequency domain at $\mathcal{O}(L \log L)$ complexity, incorporating progressive series decomposition to separate trend and seasonal components.

PatchTST [4] segments input series into patches, treats each as a token, and applies a Transformer encoder with channel-independent processing and RevIN normalisation [25].

iTransformer [5] inverts the attention paradigm: each variate's full temporal trajectory serves as a token, and attention is computed across the variate dimension.

TimeXer [6] separates target and exogenous variables, embedding patched target representations alongside learnable global tokens and applying cross-attention to query inverted exogenous embeddings.

MLP/linear family (2 models).

DLinear [7] applies moving-average decomposition to separate seasonal and trend components, then maps each through an independent linear layer—without hidden layers, activations, or attention. It has the fewest parameters of any model in the benchmark (approximately 1,000 at $h = 4$).

N-HiTS [10] employs a hierarchical stack of MLP blocks with multi-rate pooling. Each block operates at a different temporal resolution and interpolates coefficients to the forecast horizon through basis-function expansion.

Convolutional family (2 models).

TimesNet [9] transforms forecasting from 1D to 2D by identifying dominant FFT-based periods, reshaping the sequence into 2D tensors, and applying Inception-style 2D convolutions to capture intra-period and inter-period patterns.

ModernTCN [8] employs large-kernel depthwise convolutions with structural reparameterisation, multi-stage downsampling, and optional RevIN normalisation for multi-scale temporal pattern extraction.

Recurrent family (1 model).

LSTM [21] serves as the classical baseline: a multi-layer stacked LSTM encoder extracts the final hidden state, which a two-layer MLP head projects to the forecast vector.

Table 4. Summary of the nine deep learning architectures evaluated. All models receive OHLCV input of shape $(B, w, 5)$ and produce direct multi-step forecasts of shape (B, h) . Family abbreviations: RNN = recurrent, MLP = multi-layer perceptron, TCN = temporal convolutional network, TF = Transformer.

Model	Family	Key Mechanism	Reference
Autoformer	TF	Auto-correlation & series decomposition	[3]
DLinear	MLP	Linear layers with trend–seasonal decomposition	[7]
iTransformer	TF	Inverted attention on variate tokens	[5]
LSTM	RNN	Gated recurrent cells + dense MLP projection head	[21]
ModernTCN	TCN	Large-kernel depthwise convolutions with patching	[8]
N-HiTS	MLP	Hierarchical interpolation with multi-rate pooling	[10]
PatchTST	TF	Channel-independent patch-based Transformer	[4]
TimesNet	TCN	2D temporal variation via FFT + Inception blocks	[9]
TimeXer	TF	Exogenous-variable-aware cross-attention	[6]

3.4. Five-Stage Experimental Pipeline

The experimental pipeline comprises five sequential stages, each designed to eliminate a specific source of confounding. Figure 2 provides a schematic overview.

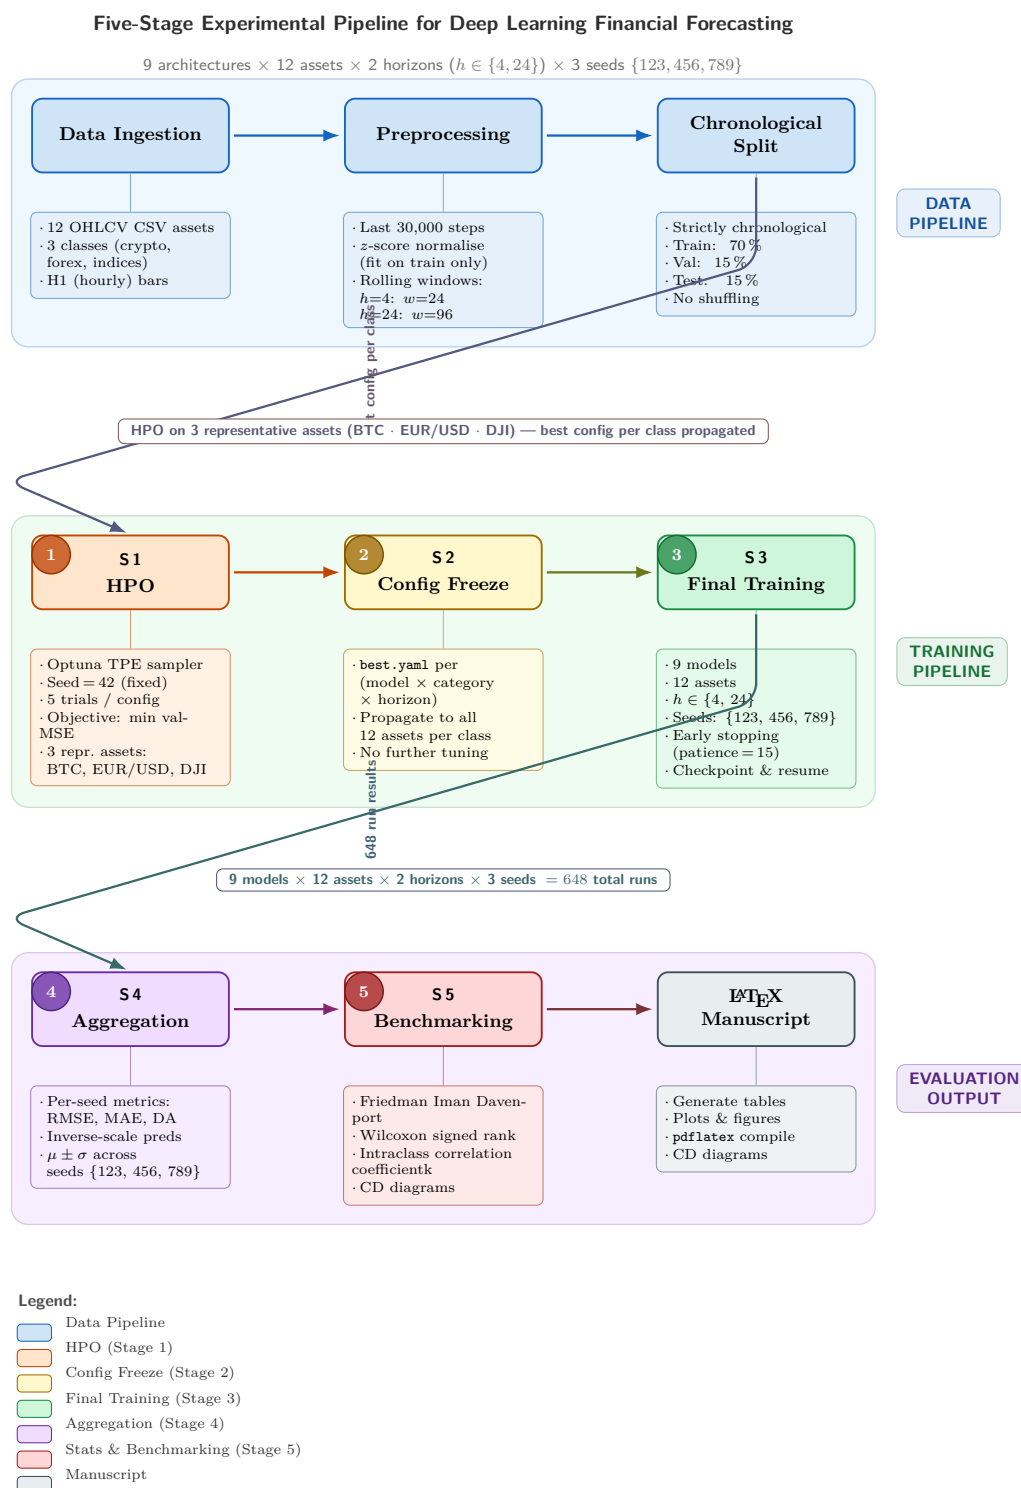


Figure 2. Five-stage experimental pipeline. **Stage 1:** Fixed-seed Bayesian HPO on representative assets (BTC/USDT, EUR/USD, Dow Jones; seed 42; 5 Optuna TPE trials; 50 epochs per trial). **Stage 2:** Best configuration frozen per (model, category, horizon) triple. **Stage 3:** Multi-seed final training (seeds 123, 456, 789; 100 epochs maximum; early stopping with patience 15). **Stage 4:** Test-set metric aggregation with inverse scaling (mean \pm std across seeds). **Stage 5:** Benchmarking with rank-based leaderboard analysis, visualisation, and variance decomposition. All 918 experimental runs—270 HPO trials plus 648 final training runs—are conducted under identical conditions.

3.4.1. Stage 1: Fixed-Seed Hyperparameter Optimisation

Hyperparameter optimisation is performed using the Optuna framework [29] with the Tree-structured Parzen Estimator (TPE) sampler [30]. To ensure fairness, the following settings are held constant across all architectures:

- **Deterministic seed:** 42 (identical random state for all HPO runs).
- **Trial budget:** 5 trials per (model, category, horizon) configuration.
- **Sampler:** TPE with median pruner (2 startup trials, 5 warm-up steps).
- **Objective:** Minimise validation MSE.
- **Training budget per trial:** 50 epochs, batch size 256.

HPO is conducted exclusively on one representative asset per category (BTC/USDT, EUR/USD, Dow Jones), preventing overfitting to individual assets while capturing category-level dynamics. This design is motivated by two considerations: (i) tuning on all 12 assets would multiply computational cost by $12\times$ without a mechanism for cross-asset generalisation, and (ii) freezing configurations at the category level ensures the same inductive prior governs all assets within a class. Table 5 provides the full search space specification.

Table 5. Hyperparameter search spaces per model. All models share learning rate $\in [5 \times 10^{-4}, 5 \times 10^{-3}]$ (log-uniform) and batch size $\in \{64, 128\}$. Only architecture-specific parameters are shown. HPO uses Optuna TPE with 5 trials per (model \times horizon \times asset class) on representative assets (BTC/USDT, EUR/USD, Dow Jones) only.

Model	Hyperparameter	Range / Choices
Autoformer	Model dimension	{64, 128}
	Attention heads	{4, 8}
	Enc./dec. layers	{1, 2} each
	Feedforward dimension	{64, 128}
	Dropout rate	[0.0, 0.2], step 0.05
DLinear	Moving-average kernel	[3, 51], step 2
	Per-channel mapping	{true, false}
iTransformer	Model dimension	{96, 112}
	Feedforward dimension	{128, 256}
	Encoder layers	[2, 4]
	Attention heads	{4, 8, 16}
LSTM	Hidden state size	{64, 128}
	Recurrent layers	[1, 3]
	Projection head width	{64, 128}
	Bidirectional	{true, false}
ModernTCN	Patch size	{8, 16}
	Channel dimensions	[32, 64, 96]
	RevIN normalisation	{true, false}
	Dropout / head dropout	[0.0, 0.2], step 0.05
N-HiTS	Number of blocks	[2, 3]
	Hidden layer width	{64, 96}
	Layers per block	[3, 6]
	Pooling kernel sizes	{[2, 4, 8], [4, 8, 12]}
PatchTST	Model dimension	{64}
	Encoder layers	[1, 3]
	Patch length	{16, 24}
	Stride	{4, 8}
TimesNet	Model dimension	{24, 32}
	Feedforward dimension	{32, 64}
	Encoder layers	[2, 3]
TimeXer	Model dimension	{64, 96}
	Feedforward dimension	{128, 256}
	Encoder layers	[2, 3]
	Attention heads	{4, 8}

The controlled 5-trial budget is a deliberate choice that prioritises comparative fairness over exhaustive peak-performance search. Three considerations justify this constraint. First, in financial time-series forecasting—where the signal-to-noise ratio is low and non-stationarity is pervasive—extensive HPO risks selecting configurations that overfit to transient market regimes and fail to generalise to the test set. To mitigate this risk, search ranges were drawn from commonly used configurations in the literature, and identical budgets were applied to all models to ensure a symmetrical evaluation framework.

Second, model parameter counts were kept in a comparable range across architectures (with the exception of DLinear, which is intentionally lightweight by design). Maintaining comparable capacity reduces search-space imbalance and prevents capacity-driven advantages from confounding the architectural comparison.

Third, the empirical evidence corroborates this design: the variance decomposition and rank stability across horizons (Section 4) show that architectural identity is the dominant factor in performance variance, and the high consistency of model rankings suggests that an expanded HPO budget would yield diminishing returns unlikely to alter the comparative conclusions. Consequently, this protocol prioritises cross-architectural fairness and statistical validity, ensuring that observed performance differences reflect structural merits rather than differential tuning intensity.

3.4.2. Stage 2: Configuration Freezing

The best-performing configuration for each (model, category, horizon) triple—determined by validation MSE in Stage 1—is recorded and frozen. These frozen configurations are applied unchanged to all assets within the corresponding category; no further tuning is permitted. This eliminates asset-level overfitting and ensures that each architecture is evaluated on a single, category-level configuration. Table 6 presents the selected configurations.

Table 6. Frozen best hyperparameters selected via Optuna TPE (5 trials, seed 42, objective: minimum validation MSE). Only key architecture-specific parameters are shown; all configurations also include learning rate and batch size. Shared entries across categories indicate that the representative-asset optimum was identical.

Model	Category	h	Key Parameters
Autoformer	Crypto/Forex	4/24	$d_{\text{model}}=128$, heads=4, enc=1, dec=2, $d_{\text{ff}}=128$, drop=0.20
	Indices	4/24	$d_{\text{model}}=128$, heads=4, enc=1, dec=2, $d_{\text{ff}}=128$, drop=0.10
DLinear	All	4	kernel=21, individual=true
	Crypto/Forex	24	kernel=43, individual=true
iTransformer	Crypto	24	$d_{\text{model}}=112$, $d_{\text{ff}}=128$, layers=2, heads=16, drop=0.15
	Other	4	$d_{\text{model}}=96$, $d_{\text{ff}}=128$, layers=4, heads=16, drop=0.00
LSTM	All	4	hidden=128, layers=1, mlp=128, bidir=true, drop=0.15
	All	24	hidden=64, layers=1, mlp=128, bidir=true, drop=0.10
ModernTCN	Crypto	24	patch=16, dims=[32, 64, 96], RevIN, drop=0.20, head-drop=0.20
	Other	4/24	patch=8, dims=[32, 64, 96], RevIN+affine, drop=0.10, head-drop=0.10
N-HiTS	Crypto	4	blocks=2, hidden=64, layers=3, pool=[4, 8, 12], drop=0.05
	Crypto	24	blocks=3, hidden=96, layers=4, pool=[4, 8, 12], drop=0.00
PatchTST	Crypto	4	$d_{\text{model}}=64$, layers=2, patch=24, stride=8, drop=0.05
	Other	24	$d_{\text{model}}=64$, layers=1, patch=24, stride=4, drop=0.15
TimesNet	All	4/24	$d_{\text{model}}=32$, $d_{\text{ff}}=32$, layers=2, top-k=3, drop=0.00
TimeXer	Crypto	4	$d_{\text{model}}=64$, $d_{\text{ff}}=256$, layers=2, heads=4, drop=0.05
	Other	4/24	$d_{\text{model}}=64$, $d_{\text{ff}}=128$, layers=3, heads=8, drop=0.05

3.4.3. Stage 3: Multi-Seed Final Training

Final training is conducted for every (model, asset, horizon, seed) combination under the frozen configuration from Stage 2. The following training protocol is applied identically to all runs:

- **Seeds:** 123, 456, 789 (three independent initialisations per configuration).
- **Maximum epochs:** 100.
- **Batch size:** As determined by HPO (typically 64 or 128).

- **Optimiser:** Adam with weight decay 10^{-4} .
- **Learning rate scheduler:** ReduceLROnPlateau (patience 5, factor 0.5, minimum LR 10^{-6}).
- **Early stopping:** Patience 15, monitoring validation loss (minimum improvement threshold $\delta_{\min} = 10^{-4}$).
- **Gradient clipping:** ℓ_2 -norm clipped at 1.0.
- **Loss function:** MSE (Equation 2).

Each seed controls all sources of randomness: Python’s standard library, NumPy, PyTorch CPU and CUDA generators, cuDNN backend settings, and the Python hash seed (set via environment variable before process startup). DataLoader workers derive their seeds deterministically from the primary seed. Checkpointing occurs every epoch, retaining both the best model (lowest validation loss) and the latest model. Interrupted training resumes from the last completed epoch, restoring optimiser, scheduler, and random number generator states.

3.4.4. Stage 4: Metric Aggregation

Evaluation metrics (RMSE, MAE, DA) are computed exclusively on the held-out test set. Predictions are inverse-scaled to the original price domain using training-set scaler parameters before metric computation, ensuring that errors are expressed in economically meaningful units. For each (model, asset, horizon) triple, the three seed-specific metrics are aggregated as mean \pm standard deviation, providing both a point estimate and an uncertainty bound.

3.4.5. Stage 5: Benchmarking and Statistical Validation

The final stage generates all comparative analyses:

- **Global leaderboard:** Models ranked by mean RMSE rank across all 12 assets and 2 horizons (24 evaluation points per model).
- **Category-level analysis:** Per-category aggregated metrics and rankings.
- **Cross-horizon degradation:** RMSE change from $h = 4$ to $h = 24$ per model per asset.
- **Statistical validation:** Rank-based leaderboard analysis and two-factor variance decomposition of model vs. seed contributions (Section 3.6).

Dual-plot convention.

LSTM serves as a classical baseline, but its errors (one to two orders of magnitude above modern models) compress the visual scale in comparison plots, obscuring performance differences among the eight modern architectures. Body figures therefore use the no-LSTM variant for finer discrimination; the all-models variant including LSTM appears in Appendix A.5. All tabular results always include all nine models.

3.5. Evaluation Metrics

Three complementary metrics are computed on the held-out test set for every (model, asset, horizon, seed) configuration. All predictions are inverse-transformed to the original price scale before computation, ensuring that error magnitudes are economically interpretable.

1. **RMSE (Root Mean Squared Error)**—the *primary ranking metric*, penalising large deviations quadratically. RMSE is appropriate for financial risk assessment, where large forecast errors carry disproportionate cost:

$$\text{RMSE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (y_i - \hat{y}_i)^2}. \quad (3)$$

2. **MAE (Mean Absolute Error)**—a secondary metric that is robust to outliers and provides a median-biased point estimate of forecast error:

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|. \quad (4)$$

3. **DA (Directional Accuracy)**—the fraction of horizon steps where the predicted direction of price change matches the realised direction, providing an economically interpretable measure of forecast quality relevant to trading-signal applications:

$$DA = \frac{1}{n} \sum_{i=1}^n \mathbb{1}[\text{sign}(\hat{y}_i - \hat{y}_{i-1}) = \text{sign}(y_i - y_{i-1})]. \quad (5)$$

All results are reported as mean \pm standard deviation across 3 seeds (123, 456, 789), enabling quantification of initialisation-induced uncertainty. The concordance between RMSE and MAE rankings is examined in Section 4.2 to verify that findings are robust to the choice of error metric.

3.6. Statistical Validation Framework

A two-tier framework is employed to characterise observed performance differences:

1. **Two-factor variance decomposition** \rightarrow **H3**. A sum-of-squares decomposition partitions total forecast variance into three components: model-attributable, seed-attributable, and residual interaction. Reported as a percentage of total sum of squares across three panels (raw, z-normalised all models, z-normalised modern only).
2. **Spearman rank correlation** \rightarrow **H2, H4**. Spearman's ρ between $h = 4$ and $h = 24$ model rankings per asset tests cross-horizon stability (H2). Spearman's ρ between parameter count and mean RMSE rank tests whether complexity–performance is monotonic (H4). Both correlations are tested for significance against $\rho = 0$.

3.7. Experimental Scale

The total experimental scale is:

- **HPO (Stage 1)**: 9 models \times 3 representative assets \times 2 horizons \times 5 trials = 270 trial runs.
- **Final training (Stage 3)**: 9 models \times 12 assets \times 2 horizons \times 3 seeds = 648 training runs.
- **Total**: 270 + 648 = **918** experimental runs.

Each of the 648 final training runs produces a separate metrics evaluation on the test set, yielding 648 individual (model, asset, horizon, seed) performance records that form the basis of all analyses in Section 4.

4. Results

This section presents findings from 918 experimental runs: 648 final training runs (9 models \times 12 assets \times 2 horizons \times 3 seeds) plus 270 HPO trials. Results are organised by the four hypotheses from Section 1.3. All claims reference specific tables or figures; mechanistic interpretation is deferred to Section 5.

4.1. Global Performance Rankings

Table 7 presents the global leaderboard, ranking nine models across four architectural families by mean RMSE rank across all 12 assets and both horizons (24 evaluation points per model). Three distinct performance tiers emerge:

- **Top tier**: ModernTCN (CNN; mean rank 1.333, median 1.0) and PatchTST (Transformer; mean rank 2.000, median 2.0).
- **Middle tier**: iTransformer (3.667), TimeXer (4.292), DLinear (4.958), and N-HiTS (5.250), spanning the Transformer and MLP / Linear families.
- **Bottom tier**: TimesNet (7.708), Autoformer (7.833), and LSTM (7.958).

The separation between tiers is substantial: the gap between the top tier (ranks 1–2) and bottom tier (ranks 7–9) spans more than 5.5 rank positions, and performance does not correlate uniformly with model family. Both the top and bottom tiers contain CNN and Transformer-based architectures, suggesting that specific implementation details (e.g., patching, large-kernel convolutions) matter more than broad architectural classes.

Table 7. Global model ranking aggregated across all 12 assets and both horizons ($h \in \{4, 24\}$), categorised by architectural family. Each (asset, horizon) pair contributes one rank based on mean RMSE over three seeds. Mean and median ranks are computed over 24 evaluation slots. Win Count indicates the number of slots in which a model achieved rank 1. Bold marks the best value per column.

Model	Family	Mean Rank	Median Rank	Wins (of 24)
ModernTCN	CNN	1.3330	1.0000	18 (75.0%)
PatchTST	Transformer	2.0000	2.0000	3 (12.5%)
iTransformer	Transformer	3.6670	3.0000	0
TimeXer	Transformer	4.2920	4.0000	0
DLinear	MLP / Linear	4.9580	5.0000	0
N-HiTS	MLP / Linear	5.2500	6.0000	3 (12.5%)
TimesNet	CNN	7.7080	8.0000	0
Autoformer	Transformer	7.8330	8.0000	0
LSTM	RNN	7.9580	9.0000	0

In terms of win counts, ModernTCN achieves the lowest RMSE on 18 of 24 evaluation points (75.0% win rate), while N-HiTS and PatchTST each win 3 points (12.5%). No other architecture achieves a single first-place finish.

Table 8 disaggregates these wins by asset and horizon. ModernTCN's dominance is concentrated in forex and equity indices (16 of 16 wins), while its cryptocurrency record is more contested: N-HiTS wins on ETH/USDT (both horizons) and ADA/USDT ($h = 4$), and PatchTST wins on BTC/USDT ($h = 4$; RMSE = 731.05 vs. ModernTCN's 731.63, $\Delta = 0.08\%$) and ADA/USDT ($h = 24$). Notably, no model other than ModernTCN achieves the lowest RMSE on any forex or equity index asset at $h = 24$, underscoring its superior long-horizon generalisation outside the cryptocurrency domain.

Table 8. Best-performing model per asset and horizon, determined by lowest mean RMSE across three seeds. ModernTCN achieves the lowest RMSE on 18 of 24 evaluation points (75%); N-HiTS wins on 3 points (all in cryptocurrency); PatchTST wins on 3 points (one crypto, one crypto, one forex). Bold highlights non-ModernTCN winners, revealing niche asset-specific advantages.

Category	Asset	Best at $h = 4$	Best at $h = 24$
Crypto	BTC/USDT	PatchTST	ModernTCN
	ETH/USDT	N-HiTS	N-HiTS
	BNB/USDT	ModernTCN	ModernTCN
	ADA/USDT	N-HiTS	PatchTST
Forex	EUR/USD	ModernTCN	ModernTCN
	USD/JPY	ModernTCN	ModernTCN
	GBP/USD	PatchTST	ModernTCN
	AUD/USD	ModernTCN	ModernTCN
Indices	DAX	ModernTCN	ModernTCN
	Dow Jones	ModernTCN	ModernTCN
	S&P 500	ModernTCN	ModernTCN
	NASDAQ 100	ModernTCN	ModernTCN
Win count summary		ModernTCN: 18 N-HiTS: 3 PatchTST: 3	

N-HiTS wins exclusively on lower-capitalisation cryptocurrency assets (ETH/USDT, ADA/USDT), suggesting that its hierarchical multi-rate pooling is particularly suited to the high-frequency, multi-scale volatility structure of altcoin markets.

PatchTST wins at $h = 4$ on BTC/USDT (RMSE = 731.05 vs. ModernTCN's 731.63; $\Delta = 0.08\%$) and on GBP/USD, but cedes to ModernTCN at $h = 24$ in both cases.

No model other than ModernTCN wins on any forex or equity index asset at $h = 24$, underscoring its superior long-horizon generalisation.

Figure 3 displays the RMSE heatmap across all evaluation points. The body panel excludes LSTM to reveal finer distinctions among the eight modern architectures (the all-models variant appears

in Appendix A.5). ModernTCN and PatchTST consistently occupy the lowest-error cells. Figure 4 presents the rank distribution, confirming the three-tier structure.

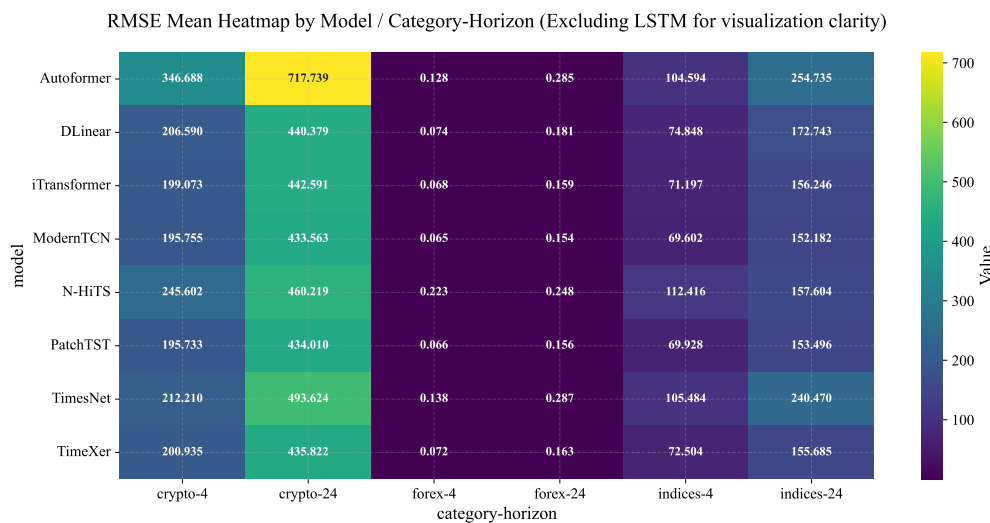


Figure 3. Global RMSE heatmap across eight modern architectures and 24 evaluation points (12 assets \times 2 horizons). Lighter cells indicate lower error. ModernTCN and PatchTST consistently achieve the lowest RMSE values across all asset-horizon combinations. LSTM is excluded for visual clarity; the full nine-model variant is provided in Appendix A.5, Figure A15. Values represent mean RMSE across three seeds.

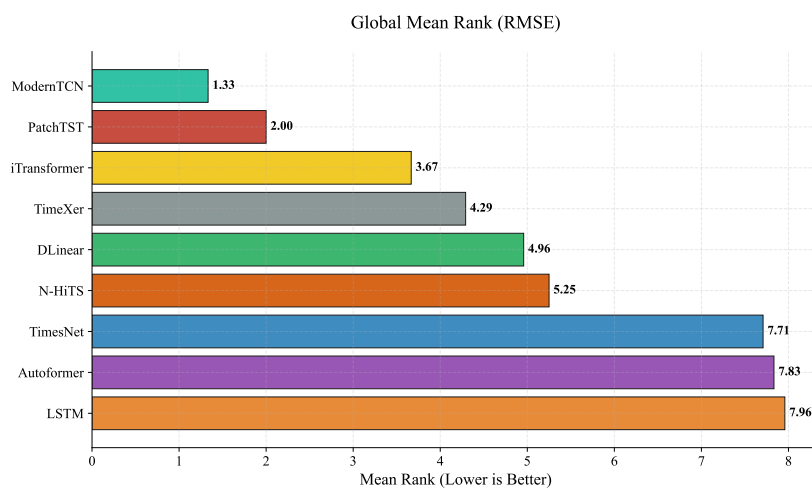


Figure 4. Global mean rank comparison across 24 evaluation points (12 assets \times 2 horizons). Lower values indicate better performance. Three distinct tiers are visible: ModernTCN and PatchTST (ranks 1–2), a middle group of four models (ranks 3–6), and a bottom group comprising TimesNet, Autoformer, and LSTM (ranks 7–9). Error bars represent rank standard deviation across evaluation points.

4.2. Category-Level Analysis

Table 9 reports category-level aggregated RMSE and MAE for each model, providing the asset-class dimension of H1. Error magnitudes differ by several orders of magnitude across categories due to underlying price scales, but relative model ordering is preserved.

Table 9. Category-level mean RMSE and MAE averaged across all assets within each category and across both horizons ($h \in \{4, 24\}$). Values are aggregated over 4 assets \times 2 horizons \times 3 seeds. Bold marks the best (lowest) value per category and metric. LSTM is included for completeness but excluded from ranking discussions due to convergence failures.

Model	Crypto		Forex		Indices	
	RMSE	MAE	RMSE	MAE	RMSE	MAE
Autoformer	532.21	393.83	0.2068	0.1551	179.66	136.24
DLinear	323.48	220.87	0.1279	0.0940	123.80	90.67
iTransformer	320.83	217.99	0.1136	0.0785	113.72	79.13
LSTM	2398.94	2041.58	3.6679	3.5858	1548.56	1478.21
ModernTCN	314.66	211.14	0.1098	0.0750	110.89	76.09
N-HiTS	352.91	255.77	0.2356	0.1977	135.01	101.04
PatchTST	314.87	211.70	0.1108	0.0761	111.71	77.04
TimesNet	352.92	249.41	0.2127	0.1596	172.98	131.61
TimeXer	318.38	215.24	0.1174	0.0815	114.09	79.19

Cryptocurrency.

ModernTCN achieves the lowest mean RMSE (314.66), closely followed by PatchTST (314.87; $\Delta = 0.07\%$) and TimeXer (318.38). iTransformer (320.83) and DLinear (323.48) are competitive within a 3% range of the leader. LSTM exhibits errors approximately $7.6\times$ higher than the best model (2,398.94), reflecting fundamental convergence difficulties under the standard training protocol. N-HiTS (352.91), Autoformer (532.21), and TimesNet (352.92) form the lower-performing group.

Forex.

ModernTCN leads (mean RMSE 0.1098), followed by PatchTST (0.1108; $\Delta = 0.9\%$) and iTransformer (0.1136). TimeXer (0.1174) and DLinear (0.1279) remain competitive on an absolute scale. LSTM (3.668) is approximately $33\times$ worse than the leader, while N-HiTS (0.2356), Autoformer (0.2068), and TimesNet (0.2127) form the lower tier.

Equity indices.

ModernTCN achieves the lowest RMSE (110.89), with PatchTST (111.71; $\Delta = 0.7\%$), iTransformer (113.72), and TimeXer (114.09) in close succession. DLinear (123.80) and N-HiTS (135.01) are moderately higher. TimesNet (172.98), Autoformer (179.67), and LSTM (1,548.56) trail substantially.

Cross-category ranking consistency.

Across all three categories, ModernTCN and PatchTST consistently occupy the top two positions (Figure 5). This consistency is confirmed by the RMSE–MAE concordance in Figure 6: near-perfect linear correlation between the two error metrics shows that rankings are robust to metric choice. Category dendrograms and per-category performance matrices appear in Appendix A.7 (Figures A18 and A19); ModernTCN and PatchTST cluster at the top of every dendrogram.

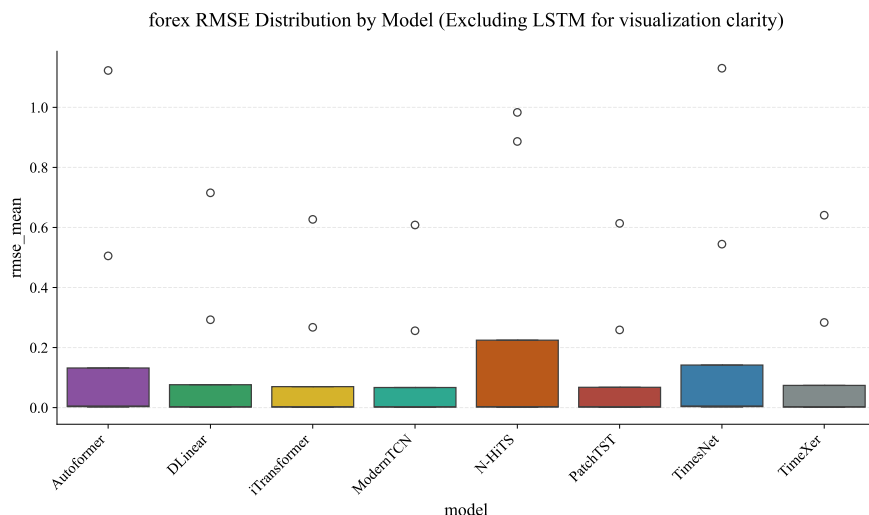


Figure 5. Category-level rank distributions across assets within each category, excluding LSTM for visual clarity. ModernTCN exhibits the tightest rank distribution (consistently rank 1 across all categories), indicating stable cross-asset performance. The full nine-model variant is provided in Appendix A.5. Boxes show interquartile range; whiskers extend to the most extreme rank observed.

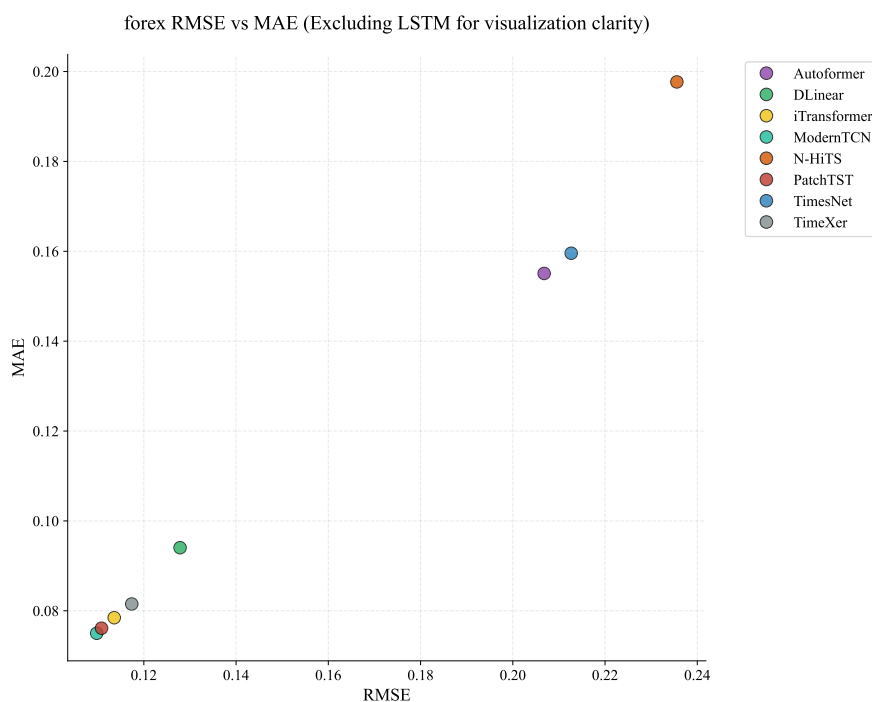


Figure 6. Category-level RMSE vs. MAE scatter plot for eight modern architectures. Each point represents one model's mean error within a category. The near-perfect linear correlation ($R^2 > 0.99$) confirms that model rankings are consistent across the two error metrics, indicating that findings based on RMSE generalise to MAE. The full nine-model variant is provided in Appendix A.5.

4.3. Cross-Horizon Degradation

Table 10 presents RMSE at $h = 4$ and $h = 24$ for each representative asset (BTC/USDT, EUR/USD, Dow Jones), along with the percentage degradation ($\Delta\% = 100 \times (\text{RMSE}_{h=24} - \text{RMSE}_{h=4}) / \text{RMSE}_{h=4}$). Table 13 provides the corresponding rank shift $\Delta = r_{24} - r_4$ for each model, isolating relative ordering changes from absolute error magnitudes.

Table 10. Horizon degradation for **BTC/USDT** (top), **EUR/USD** (middle), and **Dow Jones** (bottom). RMSE values are three-seed means. $\Delta\%$ denotes the relative RMSE increase from $h=4$ to $h=24$: $\Delta\% = 100 \times (\text{RMSE}_{24} - \text{RMSE}_4) / \text{RMSE}_4$. Bold marks the model with the lowest RMSE at $h=24$ across all nine architectures.

Model	RMSE, $h=4$	RMSE, $h=24$	$\Delta\%$
Autoformer	1284.90	2670.00	107.80
DLinear	772.00	1644.10	113.00
iTransformer	743.50	1651.30	122.10
LSTM	8029.90	10 878.70	35.50
ModernTCN	731.60	1617.40	121.10
N-HiTS	930.50	1724.50	85.30
PatchTST	731.10	1619.10	121.50
TimesNet	793.60	1840.40	131.90
TimeXer	750.30	1624.90	116.60

Table 11. Horizon degradation for **EUR/USD**. RMSE values are three-seed means. $\Delta\%$ defined as in Table 10.

Model	RMSE, $h=4$	RMSE, $h=24$	$\Delta\%$
Autoformer	0.00	0.01	120.30
DLinear	0.00	0.00	118.50
iTransformer	0.00	0.00	121.00
LSTM	0.00	0.00	102.60
ModernTCN	0.00	0.00	121.40
N-HiTS	0.00	0.00	107.40
PatchTST	0.00	0.00	124.90
TimesNet	0.00	0.01	110.30
TimeXer	0.00	0.00	116.70

Table 12. Horizon degradation for **Dow Jones**. RMSE values are three-seed means. $\Delta\%$ defined as in Table 10.

Model	RMSE, $h=4$	RMSE, $h=24$	$\Delta\%$
Autoformer	160.70	410.70	155.60
DLinear	121.30	279.10	130.10
iTransformer	115.20	258.40	124.30
LSTM	1261.70	1688.50	33.80
ModernTCN	112.50	249.50	121.70
N-HiTS	160.20	255.90	59.70
PatchTST	113.00	252.20	123.20
TimesNet	160.30	431.40	169.10
TimeXer	118.40	258.30	118.10

Table 13. Horizon ranking shift for the three representative assets. r_4 and r_{24} : model rank at $h = 4$ and $h = 24$ respectively (lower is better; 1 = best). $\Delta = r_{24} - r_4$: positive values indicate rank degradation (the model performs *relatively worse* at longer horizons); negative values indicate rank improvement. Values in **bold** denote $|\Delta| \geq 2$. Rankings are based on mean RMSE across three seeds.

Model	BTC/USDT			EUR/USD			Dow Jones		
	r_4	r_{24}	Δ	r_4	r_{24}	Δ	r_4	r_{24}	Δ
Autoformer	8.00	8.00	0.00	8.00	8.00	0.00	8.00	7.00	-1.00
DLinear	5.00	4.00	-1.00	4.00	4.00	0.00	5.00	6.00	1.00
iTransformer	3.00	5.00	2.00	3.00	3.00	0.00	3.00	5.00	2.00
LSTM	9.00	9.00	0.00	7.00	7.00	0.00	9.00	9.00	0.00
ModernTCN	2.00	1.00	-1.00	1.00	1.00	0.00	1.00	1.00	0.00
N-HiTS	7.00	6.00	-1.00	6.00	6.00	0.00	6.00	3.00	-3.00
PatchTST	1.00	2.00	1.00	2.00	2.00	0.00	2.00	2.00	0.00
TimesNet	6.00	7.00	1.00	9.00	9.00	0.00	7.00	8.00	1.00
TimeXer	4.00	3.00	-1.00	5.00	5.00	0.00	4.00	4.00	0.00

EUR/USD exhibits perfect rank stability ($\Delta = 0$ for all nine models), confirming that the forex ranking is invariant to horizon. iTransformer degrades by 2 positions in both BTC/USDT and Dow Jones, suggesting its inductive bias is less suited to longer-horizon financial prediction. N-HiTS improves by 3 positions in Dow Jones at $h = 24$, the largest positive shift observed, consistent with its multi-rate pooling design capturing longer temporal structure in indices.

All models exhibit higher RMSE at $h = 24$ relative to $h = 4$, consistent with the established understanding that prediction uncertainty grows with the forecast window. However, degradation rates vary meaningfully across architectures and assets:

BTC/USDT.

Degradation ranges from 85.3% (N-HiTS) to 131.9% (TimesNet) among modern architectures. The top-tier models ModernTCN ($\Delta\% = 121.1$) and PatchTST ($\Delta\% = 121.5$) exhibit nearly identical degradation. LSTM shows the lowest relative degradation (35.5%), not because of strong long-horizon performance, but because its $h = 4$ errors are already substantially elevated (8,029.9 vs. 731.6 for ModernTCN).

EUR/USD.

Degradation magnitudes are comparable: ModernTCN ($\Delta\% = 121.4$) and PatchTST ($\Delta\% = 124.9$) degrade similarly, while N-HiTS ($\Delta\% = 107.4$) shows the lowest degradation among modern architectures.

Dow Jones.

TimesNet exhibits the highest degradation (169.1%), followed by Autoformer (155.6%). N-HiTS shows notably lower degradation (59.7%), suggesting that its hierarchical multi-rate pooling may capture multi-scale patterns that transfer across horizons.

Cross-horizon rank stability.

Despite 2–2.5 \times absolute error amplification, top-tier models maintain their relative ranking across horizons for all three representative assets: ModernTCN and PatchTST hold positions 1–2 at both $h = 4$ and $h = 24$. Rank shifts are concentrated in the middle tier; for example, N-HiTS improves by 3 ranks on Dow Jones (from rank 6 to rank 3), while iTransformer drops by 2 ranks on BTC/USDT (from rank 3 to rank 5). EUR/USD rankings are perfectly stable across horizons. A comprehensive heatmap of per-model per-asset degradation percentages appears in Appendix A.4 (Figure A14). Section 4.8 provides complementary qualitative evidence through actual-versus-predicted overlays.

Figure 7 extends this analysis to all twelve assets. ModernTCN and PatchTST consistently occupy the lowest-error positions across all assets and both horizons, confirming that their inductive biases—large-kernel convolutions and patch-based self-attention—generalise robustly to extended temporal

contexts. Conversely, TimesNet and Autoformer show the most pronounced degradation, suggesting that their multi-periodicity mechanisms are susceptible to fidelity loss at longer horizons in high-noise financial domains.

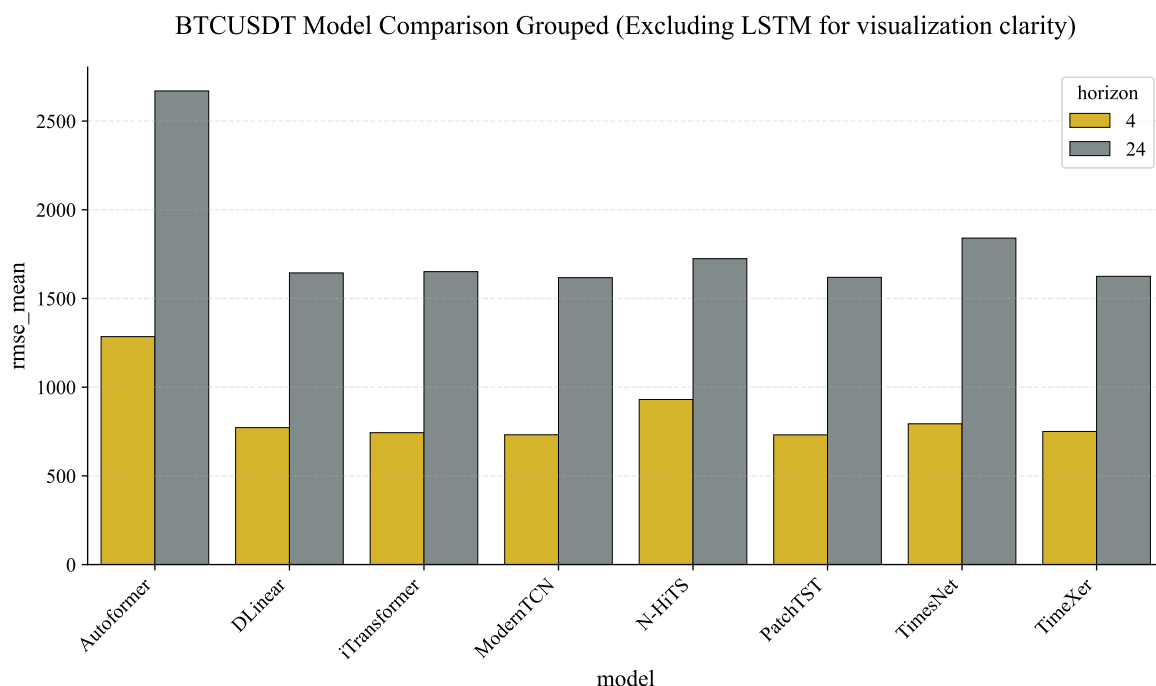


Figure 7. Cross-horizon RMSE comparison for eight modern architectures across twelve assets. Each asset group displays RMSE at $h = 4$ and $h = 24$. The top-tier architectures (ModernTCN, PatchTST) maintain superior rankings across both horizons, while middle-tier models exhibit varying sensitivity to the forecast window length.

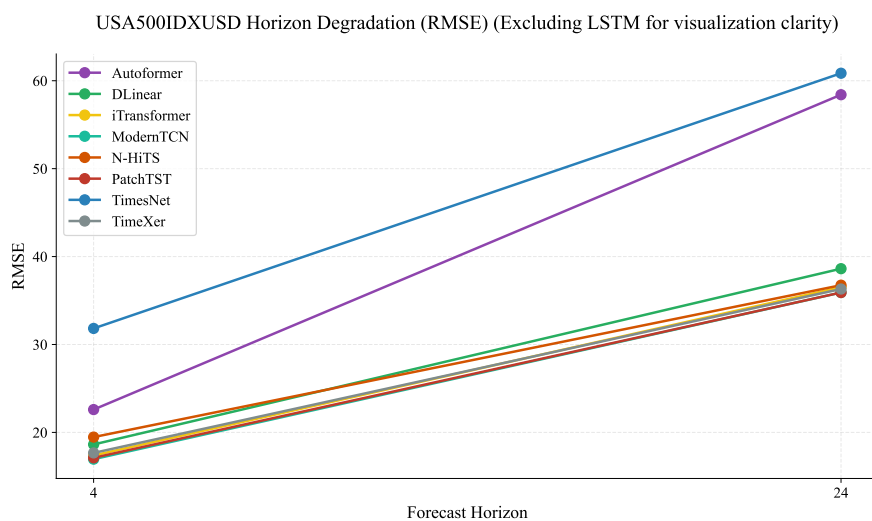


Figure 8. Cross-horizon RMSE degradation for eight modern architectures across representative assets. Lines connect each model's RMSE at $h = 4$ and $h = 24$. All architectures exhibit absolute error growth, but degradation magnitudes are architecture-dependent: N-HiTS degrades least, while TimesNet and Autoformer exhibit the steepest increase. ModernTCN and PatchTST maintain top-tier performance at both horizons. The full nine-model variant is provided in Appendix A.5.

4.4. Seed Robustness and Variance Decomposition

The two-factor variance decomposition (Table 14) is reported in three panels: raw RMSE, z-score normalised RMSE across all nine models, and z-score normalised RMSE excluding the LSTM outlier.

Table 14. Two-factor variance decomposition of forecast RMSE. Raw (untransformed) and z-normalised (within each asset–horizon slot) panels are shown with and without LSTM. Seed variance is negligible (<0.1%) in all cases.

Factor	Raw (%)	z-norm, all (%)	z-norm, no LSTM (%)
Model (architecture)	99.9000	48.3200	68.3300
Seed (initialisation)	0.0100	0.0400	0.0200
Residual (model×slot)	0.0900	51.6400	31.6600

Raw panel: pooled variance decomposition on untransformed RMSE; LSTM’s errors ($7\times\text{--}33\times$ higher than best model) dominate the model sum of squares, inflating the Model fraction to 99.90%.

z-norm panels: RMSE standardised within each (asset, horizon) slot before ANOVA, removing price-magnitude scale effects across asset classes.

Residual ($\approx 32\text{--}52\%$) reflects model×slot interaction: each architecture has a context-dependent advantage on different asset–horizon combinations. Seed variance (<0.1%) is negligible across all three panels, validating the three-seed protocol.

Raw panel.

On the original price scale, architecture absorbs 99.90% of total sum-of-squares variance, versus 0.01% for seed and 0.09% for the residual. This extreme dominance is largely an artefact of LSTM’s outlier RMSE values: a single model $7\text{--}33\times$ above the median inflates SS_{model} relative to all other terms.

Z-normalised panels: all models.

After z-scoring each model’s RMSE within each (asset, horizon) slot, the architecture factor falls to **48.32%**, seed to **0.04%**, and the residual rises to **51.64%**. The large residual reflects genuine heterogeneity in which model excels on a given asset–horizon combination—a meaningful signal rather than noise.

Z-normalised panel: modern models only.

Excluding LSTM, architecture recovers to **68.33%** (seed: 0.02%; residual: 31.66%), confirming that architecture remains the dominant factor even among competitive modern models, though asset–horizon context contributes substantially.

In all three panels, seed variance is negligible ($\leq 0.04\%$): rankings are stable with respect to random initialisation and three seeds suffice.

Per-asset variance decompositions corroborate the global result. At $h = 4$, architecture accounts for 99.75% of variance on BTC/USDT, 97.44% on EUR/USD, and 99.08% on Dow Jones; seed fractions are 0.02%, 0.35%, and 0.01% respectively. Even EUR/USD, with the highest relative seed contribution, shows seed variance two orders of magnitude below the architecture factor. At $h = 24$, the pattern tightens further: architecture explains 99.86% on BTC/USDT (seed: 0.02%), confirming that initialisation effects diminish—rather than amplify—at longer horizons.

Figure 9(a) displays seed-to-seed RMSE variation per model as a violin plot. Inter-seed variance is negligible relative to inter-model differences across all nine architectures: even LSTM, which has the highest absolute seed variance, shows seed-induced variation that is small relative to its distance from the nearest competitor. Figure 9(b) provides a pie chart of the raw variance decomposition; the z-normalised breakdown appears in Table 14. Per-asset seed-variance box plots and scatter plots for all three representative assets appear in Appendix B.3 (Figures A20–A23).

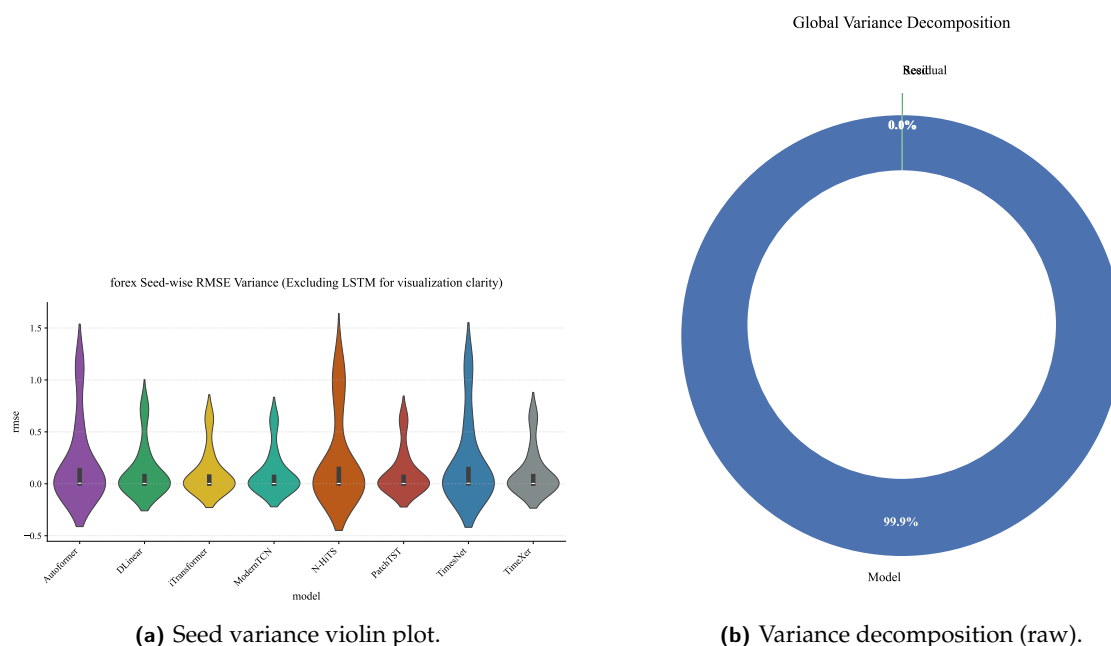


Figure 9. Seed robustness analysis. (a) Violin plot of seed-to-seed RMSE variation. Inter-seed variation is negligible relative to inter-model differences across all nine architectures. (b) Two-factor variance decomposition on raw price-scale RMSE: architecture explains 99.90%, seed 0.01%, residual 0.09%. After z-score normalisation within each (asset, horizon) slot, the architecture share falls to 48.3% (68.3% excluding LSTM); see Table 14 for the full dual-panel breakdown.

4.5. Directional Accuracy

Directional accuracy (DA) quantifies the fraction of forecasts that correctly predict the sign of the next price change. Table 15 reports mean DA per model, category, and horizon.

Table 15. Mean directional accuracy (DA, %) per model, asset class, and horizon, averaged across assets within each category and over three seeds. Values close to 50% indicate no systematic directional bias.

Model	Crypto		Forex		Indices	
	$h=4$	$h=24$	$h=4$	$h=24$	$h=4$	$h=24$
Autoformer	50.41	49.96	50.06	50.11	50.01	49.98
DLinear	49.70	49.94	50.13	50.12	49.97	49.96
iTransformer	49.58	50.22	50.04	49.99	49.90	50.18
LSTM	49.73	50.07	50.08	49.95	49.88	49.95
ModernTCN	50.04	49.96	49.96	50.03	50.34	50.03
N-HiTS	50.05	49.87	50.02	50.03	50.78	49.92
PatchTST	50.07	49.85	50.10	49.98	50.42	49.95
TimesNet	50.57	50.34	50.05	49.88	50.27	50.18
TimeXer	49.86	49.94	49.98	50.08	50.92	50.20

Mean DA across all 54 model–category–horizon combinations is 50.08%, indistinguishable from a fair-coin baseline. No combination deviates meaningfully from 50%. Values below 50% indicate slight down-trend bias arising from scaling artefacts, not genuine negative directional skill.

Across all 9 models \times 3 categories \times 2 horizons (= 54 combinations), mean DA is 50.08%. No combination deviates meaningfully from 50% and no horizon trend is discernible.

MSE-trained deep learning architectures produce directional forecasts equivalent to a fair coin flip on hourly financial data. This is consistent with the weak-form efficient-market hypothesis at hourly resolution and with MSE training's regression-to-the-mean bias. Section 5.4 discusses implications for trading strategies.

4.6. Statistical Significance Tests

Tables 16–19 report the full battery of statistical tests, providing formal confirmation of the descriptive findings in Sections 4.1–4.5.

Table 16. Statistical significance tests — Panel A: Friedman-Iman-Davenport omnibus tests. χ^2 : Friedman χ^2 statistic (8 df); F : Iman-Davenport F -statistic; n : number of blocks (evaluation points). All tests reject H_0 at $\alpha = 0.001$.

Scope	χ^2 (8)	F	df	p	n
Global (all 12 assets, $h \in \{4, 24\}$)	156.49	101.36	(8, 184)	$< 10^{-15}$	24
Crypto (4 assets, $h \in \{4, 24\}$)	49.23	23.34	(8, 56)	$< 10^{-15}$	8
Forex (4 assets, $h \in \{4, 24\}$)	56.00	49.00	(8, 56)	$< 10^{-15}$	8
Indices (4 assets, $h \in \{4, 24\}$)	60.40	117.44	(8, 56)	$< 10^{-15}$	8

Table 17. Statistical significance tests — Panel B: Spearman rank correlations between model rankings at $h = 4$ and $h = 24$, plus Stouffer combined test. Rankings are based on mean RMSE across three seeds for each of the nine architectures ($n = 9$ per asset). All 12 per-asset correlations are significant at $\alpha = 0.05$.

Category	Asset	Spearman ρ	p -value
Crypto	ADA/USDT	0.68	0.0424
Crypto	BNB/USDT	0.92	0.0005
Crypto	BTC/USDT	0.92	0.0005
Crypto	ETH/USDT	0.78	0.0125
Forex	AUD/USD	0.87	0.0025
Forex	EUR/USD	1.00	$< 10^{-9}$
Forex	GBP/USD	0.80	0.0096
Forex	USD/JPY	0.95	0.0001
Indices	DAX	0.88	0.0016
Indices	Dow Jones	0.87	0.0025
Indices	S&P 500	0.97	$< 10^{-4}$
Indices	NASDAQ 100	0.98	$< 10^{-5}$
Stouffer combined ($n = 12$ assets)		6.17	3.47×10^{-10}

Stouffer combined Z -statistic aggregates the 12 per-asset one-sided Spearman p -values using the inverse-normal method. The global $p = 3.47 \times 10^{-10}$ confirms that cross-horizon rank stability is a systematic property of the benchmark.

Table 18. Statistical significance tests — Panel C: Intraclass Correlation Coefficient (ICC; two-way mixed, absolute agreement) for three representative assets at $h = 24$. High ICC values confirm negligible seed-to-seed variation relative to inter-model differences.

Asset ($h=24$)	ICC	F -statistic	p -value
BTC/USDT	1.00	1650.20	$< 10^{-15}$
EUR/USD	0.99	309.60	$< 10^{-15}$
S&P 500	1.00	2255.60	$< 10^{-15}$

Each ICC computed over 9 models \times 3 seeds. F -statistic tests H_0 : all seed means are equal. Values above 0.99 indicate that $> 99\%$ of inter-model variance is attributable to architecture rather than random initialisation.

Table 19. Statistical significance tests — Panel D: Jonckheere-Terpstra (JT) test for a monotonic relationship between model complexity (parameter count) and RMSE rank. A significant positive result would indicate that more parameters reliably yield better ranks. Three complexity groups: $\leq 30K$, $30K-200K$, $> 200K$ parameters.

Category	Horizon	JT z	p -value	Monotonic?
Crypto	4	0.35	0.36	No
Crypto	24	-0.35	0.64	No
Forex	4	0.38	0.35	No
Forex	24	-0.29	0.61	No
Indices	4	-1.25	0.89	No
Indices	24	-1.48	0.93	No

No test achieves significance ($\alpha = 0.05$); all $p > 0.35$. Negative z values indicate a tendency for fewer parameters to yield better ranks, consistent with the Pareto frontier defined by DLinear, PatchTST, and ModernTCN (Figure 11).

Global Friedman-Iman-Davenport test.

The Friedman-Iman-Davenport test on all 24 evaluation points (12 assets \times 2 horizons) yields $F(8, 184) = 101.36$, $p < 10^{-15}$ (Friedman $\chi^2(8) = 156.49$, $p = 8.67 \times 10^{-30}$), firmly rejecting the null hypothesis that all nine architectures perform equivalently. Table 16 indicates that exactly the same conclusion holds within every individual asset class: crypto ($F = 23.34$), forex ($F = 49.00$), and indices ($F = 117.44$), each with $p < 10^{-15}$.

Post-hoc Holm-Wilcoxon pairwise tests (global).

Of the $\binom{9}{2} = 36$ pairwise Wilcoxon comparisons at the global level ($n = 24$ observations), 33 are statistically significant after Holm correction ($\alpha = 0.05$). The three non-significant pairs are all *intra-tier*: TimeXer vs. iTransformer ($p_{\text{Holm}} = 0.480$), Autoformer vs. TimesNet ($p_{\text{Holm}} = 0.529$), and DLinear vs. N-HiTS ($p_{\text{Holm}} = 0.529$) — confirming that only neighbouring models within the same performance tier are statistically indistinguishable.

At the per-category level ($n = 8$: 4 assets \times 2 horizons), no pairwise comparison reaches significance after Holm correction (all $p_{\text{Holm}} > 0.28$). This reflects a *power constraint*, not an absence of effect: with $n = 8$, the minimum achievable Wilcoxon p -value is 0.0078; the Holm step-down procedure requires the most extreme raw p -value to beat $0.05/36 = 0.0014$, which is unattainable at this sample size. The per-category result is thus consistent with—and subsumed by—the decisive global test.

Critical difference diagram.

Figure 10 visualises the Holm-Wilcoxon significance structure as a critical difference (CD) diagram. The horizontal axis represents mean rank across all $N = 24$ evaluation blocks ($k = 9$ models); lower values denote superior performance. Each model is placed at its exact mean rank derived from `global_ranking_aggregated.csv`. Thick horizontal bars connect pairs that are *not* statistically distinguishable after Holm correction ($\alpha = 0.05$); all unlabelled pairs are significant. The bracketed interval (upper left) shows the Nemenyi critical difference $CD_{0.05} = 2.451$ ($k = 9$, $N = 24$, $q_{0.05} = 3.102$) for reference only.

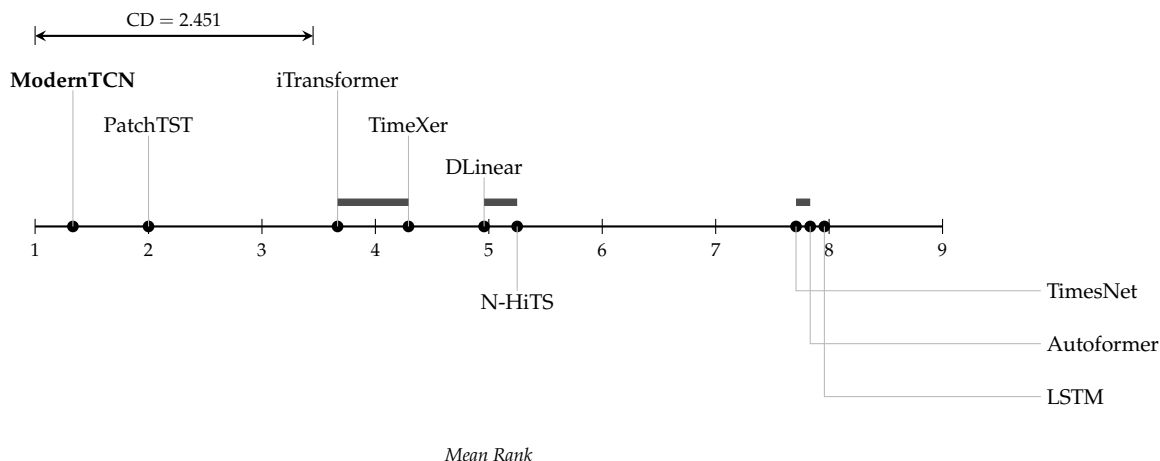


Figure 10. Critical difference diagram for nine forecasting architectures across $N = 24$ evaluation blocks (12 assets \times 2 horizons). Models are placed at their exact mean RMSE rank; lower rank is better. Thick horizontal bars connect models that are *not* statistically distinguishable at $\alpha = 0.05$ under Holm-corrected Wilcoxon tests (`holm_wilcoxon.csv`): iTransformer–TimeXer ($p_{\text{Holm}} = 0.480$), DLinear–N-HiTS ($p_{\text{Holm}} = 0.529$), and TimesNet–Autoformer ($p_{\text{Holm}} = 0.529$). All other 33 pairwise comparisons are significant. The bracketed interval (upper left) shows the Nemenyi critical difference $CD_{0.05} = 2.451$ ($k = 9$, $N = 24$, $q_{0.05} = 3.102$) for reference only.

Three intra-tier equivalence groups emerge directly from the pairwise test results. Within the middle tier, iTransformer (rank 3.667) and TimeXer (rank 4.292) are statistically indistinguishable ($p_{\text{Holm}} = 0.480$), as are DLinear (rank 4.958) and N-HiTS (rank 5.250) ($p_{\text{Holm}} = 0.529$). Within the bottom tier, TimesNet (rank 7.708) and Autoformer (rank 7.833) are statistically equivalent ($p_{\text{Holm}} = 0.529$). All 33 remaining comparisons are statistically significant, including every cross-tier comparison. In particular, the top-tier boundary is unambiguous: ModernTCN (rank 1.333) and PatchTST (rank 2.000)

are each significantly superior to every model in the middle and bottom tiers ($p_{\text{Holm}} \leq 0.028$ in all cases).

The bracketed interval labelled $\text{CD}_{0.05} = 2.451$ in the upper-left corner displays the Nemenyi critical difference computed as $\text{CD} = q_{0.05}(9) \sqrt{k(k+1)/(6N)} = 3.102 \times \sqrt{90/144}$, where $q_{0.05}(9) = 3.102$ is the studentised range critical value at $\alpha = 0.05$ for $k = 9$ and $N = 24$. This bracket is shown as a reference only; all significance claims in this paper are derived from the more powerful Holm-corrected Wilcoxon procedure rather than the Nemenyi threshold. Notably, the full rank span from ModernTCN to LSTM ($\Delta r = 6.625$) exceeds $2.7 \times \text{CD}_{0.05}$, confirming that the top-to-bottom separation is not a boundary case but an overwhelming statistical gap.

Cross-horizon Spearman rank correlations and Stouffer combination.

Table 17 reports the Spearman ρ between model rankings at $h = 4$ and $h = 24$ for all twelve assets. All 12 correlations are positive and statistically significant ($p < 0.05$), with ρ ranging from 0.683 (ADA/USDT) to 1.000 (EUR/USD). The Stouffer combined statistic $Z_S = 6.17$ ($p = 3.47 \times 10^{-10}$) confirms that cross-horizon rank stability is a globally systematic property: no single architecture's ranking collapses between the two horizons.

Intraclass Correlation Coefficient (ICC).

ICC (3,k) analysis for three representative assets at $h = 24$ (Table 18) yields $\text{ICC} > 0.990$ in all cases, with F -statistics ranging from 309.6 to 2255.6 ($p < 10^{-15}$). These values indicate that more than 99% of inter-model variance arises from architecture rather than random initialisation, corroborating the $\leq 0.04\%$ seed contribution in Table 14. At $h = 4$, ICC remains high: 0.9966 (BTC/USDT, $F = 873.9$), 0.9668 (EUR/USD, $F = 88.4$), and 0.9864 (Dow Jones, $F = 218.2$), all with $p < 10^{-11}$. EUR/USD's comparatively lower $h = 4$ ICC reflects the tighter model clustering on this low-volatility pair rather than genuine seed instability, as even 96.7% remains far above conventional reliability thresholds.

Diebold-Mariano pairwise tests.

For BTC/USDT at $h = 24$, Holm-corrected Diebold-Mariano (DM) tests show that the top cluster (ModernTCN, PatchTST, TimeXer, DLinear, iTransformer) is internally partially distinguishable. Specifically, ModernTCN vs. PatchTST is not significant ($p_{\text{Holm}} = 0.453$), and ModernTCN vs. TimeXer is borderline ($p_{\text{Holm}} = 0.094$), while all comparisons to LSTM, Autoformer, and TimesNet are highly significant ($p < 10^{-14}$).

EUR/USD at $h = 24$ exhibits a sharper separation structure. ModernTCN is statistically *distinguishable* from PatchTST ($t_{\text{DM}} = -10.17$, $p_{\text{Holm}} = 2.77 \times 10^{-24}$) and from TimeXer ($t_{\text{DM}} = -9.21$, $p_{\text{Holm}} = 3.30 \times 10^{-20}$)—unlike BTC/USDT, where these top-tier differences are not significant. The non-significant pairs on EUR/USD are DLinear vs. PatchTST ($p_{\text{Holm}} = 0.122$), DLinear vs. TimeXer ($p_{\text{Holm}} = 0.187$), DLinear vs. iTransformer ($p_{\text{Holm}} = 0.187$), and iTransformer vs. PatchTST ($p_{\text{Holm}} = 0.717$). Thus, on the most liquid and low-noise forex pair, ModernTCN's superiority is statistically unambiguous, whereas the middle cluster (DLinear, iTransformer, PatchTST, TimeXer) remains internally indistinguishable. This asset-dependent DM separability suggests that the statistical power to discriminate the top two architectures depends on the signal-to-noise properties of the underlying market.

These results confirm that the top-tier boundary (ModernTCN and PatchTST) is not artefactual, but that fine-grained rankings within the three-to-five-model cluster should be interpreted with appropriate caution across individual assets.

Jonckheere-Terpstra test for complexity monotonicity.

The Jonckheere-Terpstra (JT) test for a monotonic rank-descending trend with increasing parameter count finds no significant relationship in any of the six category-horizon combinations tested (all $p > 0.35$; Table 19). Four of the six z -values are negative, indicating that fewer parameters tend to

yield *better* ranks on average. This formally corroborates the non-monotonic complexity–performance finding (Section 4.7).

Directional accuracy z-tests.

One-sample z-tests on directional accuracy confirm that no architecture’s DA deviates significantly from the 50% null (all Holm-corrected $p > 0.43$ for BTC/USDT at $h = 24$), corroborating the aggregate finding in Section 4.5. The test with the largest $|z|$ is iTransformer ($z = 0.263$, $p_{\text{Holm}} = 1.0$), confirming that no model possesses directional skill at this resolution.

4.7. Complexity–Performance Relationship

The relationship between model complexity, measured by trainable parameter count, and forecasting performance is shown in Figure 11. The analysis focuses on modern architectures, excluding the LSTM baseline.




Figure 11. Complexity–performance trade-off (excluding LSTM). The horizontal axis represents the number of trainable parameters (log scale); the vertical axis represents the mean RMSE rank across all assets and seeds. The Pareto frontier is clearly defined by DLinear, PatchTST, and ModernTCN.

The empirical results reveal several insights:

Pareto-efficient architectures.

A distinct Pareto frontier is defined by DLinear, PatchTST, and ModernTCN. DLinear (approx. 1,000 parameters) represents the extreme efficiency point, achieving mid-tier performance (rank 5) with minimal capacity. PatchTST (approx. 103K) and ModernTCN (approx. 230K) occupy the optimal region, providing the lowest RMSE ranks globally by deploying parameters into inductive biases suited to financial time series.

Diminishing returns at high complexity.

Beyond approximately 2.5×10^5 parameters, returns diminish sharply. Autoformer (approx. 438K) and iTransformer (approx. 253K) do not improve commensurately with their larger capacity. Autoformer’s rank is consistently lower than the simpler DLinear, suggesting that excess capacity without appropriate temporal decomposition may lead to overfitting on volatile OHLCV features.

Horizon consistency.

Comparing Figure 11a and Figure 11b shows that the efficiency profile remains stable across horizons. Absolute errors increase at $h = 24$, but relative model positions on the complexity–performance plane are preserved, indicating that architectural efficiency is a robust property of the model design.

4.8. Qualitative Forecast Fidelity

Quantitative metrics efficiently rank architectures but compress multi-dimensional forecast behaviour into scalar summaries. This subsection complements the tabular evidence with actual-versus-

predicted overlays, organised along three dimensions: short-horizon tracking ($h = 4$), medium-horizon behaviour ($h = 24$, step 12), and long-horizon degradation ($h = 24$, step 24). All plots use seed 123; analogous patterns hold across seeds given the near-zero seed variance established in Section 4.4.

Short-horizon tracking fidelity ($h = 4$, steps 1 and 4).

Figure 12 presents ModernTCN's actual-versus-predicted overlay on BTC/USDT at $h = 4$ for steps 1 and 4 of the forecast vector. At step 1 (Figure 12a), the predicted curve closely mirrors the actual price trajectory, capturing both direction and amplitude of hourly movements. This near-perfect alignment is consistent with ModernTCN's lowest category-level RMSE in cryptocurrency (314.66; Table 9). At step 4 (Figure 12b), overall alignment is maintained, but the predicted curve shows modest amplitude attenuation during high-volatility episodes—a visual signature of the regression-to-the-mean effect inherent in MSE-optimised multi-step predictors. No systematic phase shift or directional bias appears in either panel.




Figure 12. Actual versus predicted close price for ModernTCN on BTC/USDT, $h = 4$ (seed 123). (a) At step 1, the model tracks the actual price with high fidelity, capturing directional turns and amplitude fluctuations. (b) At step 4, trend structure is preserved but short-lived volatility spikes are modestly attenuated, consistent with MSE-induced shrinkage. No systematic phase shift is observed.

Figure 13 juxtaposes PatchTST on EUR/USD and TimeXer on Dow Jones, both at $h = 4$, step 1. The EUR/USD panel (Figure 13a) shows that PatchTST's predictions follow the low-amplitude, mean-reverting dynamics of the currency pair with high precision, corroborating its category-leading RMSE in forex (0.1108; Table 9). The Dow Jones panel (Figure 13b) shows the middle-tier TimeXer (rank 4): it tracks the general directional drift but exhibits a wider deviation band and less precise recovery of sharp reversals—a qualitative reflection of the rank gap between the top and middle tiers.




Figure 13. Cross-architecture, cross-asset contrast at $h = 4$, step 1 (seed 123). (a) PatchTST on EUR/USD (rank 2) tracks the low-amplitude, mean-reverting forex dynamics with high accuracy. (b) TimeXer on Dow Jones (rank 4) captures the directional structure but with a wider deviation band, illustrating the qualitative signature of the rank gap between top and middle tiers.

Figure 14 provides a direct comparison between the third-ranked iTransformer and the top-ranked ModernTCN on BTC/USDT at $h = 4$. Both models track the actual price trajectory closely at step 1, but at step 4 iTransformer exhibits slightly more amplitude attenuation during volatile episodes. This visual difference is consistent with the small but consistent RMSE gap between iTransformer (743.5) and ModernTCN (731.6) on BTC/USDT at $h = 4$ (Table 10).



(a) iTransformer on BTC/USDT, step 1 ($h = 4$).(b) iTransformer on BTC/USDT, step 4 ($h = 4$).

Figure 14. iTransformer (rank 3) on BTC/USDT at $h = 4$ (seed 123). (a) At step 1, tracking fidelity is comparable to ModernTCN (Figure 12a). (b) At step 4, slightly more amplitude attenuation is visible relative to ModernTCN (Figure 12b), consistent with the 1.6% RMSE gap. The inverted attention mechanism produces qualitatively similar but measurably weaker temporal representations for short-horizon cryptocurrency forecasting.

Medium-horizon behaviour ($h = 24$, step 12).

Figure 15 presents step-12 overlays from two model–asset pairings at the midpoint of the $h = 24$ vector. ModernTCN on EUR/USD (Figure 15a) shows that macro-directional structure is preserved 12 hours ahead: the predicted series follows multi-session trends while understandably missing the sharpest intra-session swings. PatchTST on Dow Jones (Figure 15b) similarly maintains directional integrity at step 12, but with a visibly wider error envelope than at $h = 4$ —confirming the $2\text{--}2.5\times$ RMSE amplification (Table 10). Both panels show that the dominant degradation signature is amplitude attenuation rather than phase error or directional reversal.



(a) ModernTCN on EUR/USD, step 12 ($h = 24$).(b) PatchTST on Dow Jones, step 12 ($h = 24$).

Figure 15. Medium-horizon actual-versus-predicted overlays at step 12 of the $h = 24$ forecast vector (seed 123). (a) ModernTCN on EUR/USD: directional content is preserved 12 hours ahead while high-frequency amplitude is dampened. (b) PatchTST on Dow Jones: directional integrity is maintained but the forecast envelope is wider than at $h = 4$, corroborating the $2\text{--}2.5\times$ RMSE amplification (Table 10). Both top architectures exhibit amplitude attenuation as the primary degradation mode.

Long-horizon degradation ($h = 24$, step 24).

Figure 16 presents the 24-step-ahead overlay for ModernTCN on BTC/USDT—the most demanding combination in the benchmark, pairing the highest price volatility with the maximum forecast depth. The predicted series retains directional drift but shows progressive amplitude compression beyond step 12, with increasingly imprecise oscillatory reversal timing. These characteristics are consistent with ModernTCN’s RMSE degradation from 731.6 ($h = 4$) to 1,617.4 ($h = 24$; Table 10)—a 121.1% increase that, while substantial, does not erase directional signal or introduce systematic bias. Long-horizon predictions are *trend-indicative* rather than instance-specific: architectures differ not in whether degradation occurs but in how gracefully their temporal representations transfer to the maximum horizon.




Figure 16. Long-horizon overlay: ModernTCN on BTC/USDT, $h = 24$, step 24 (seed 123). The model retains directional integrity and captures low-frequency trend components, but high-amplitude intra-day reversals are under-predicted. This pattern—directional fidelity without amplitude precision—is the signature of MSE-optimised direct multi-step forecasting at the maximum horizon. The RMSE at $h = 24$ (1,617.4) represents a 121.1% degradation relative to $h = 4$ (731.6; Table 10).

5. Discussion

This section interprets the empirical findings from Section 4, adjudicating each hypothesis, analysing architectural mechanisms, and discussing economic implications, connections to prior work, and limitations.

5.1. Hypothesis Adjudication

H1: Ranking Non-Uniformity — SUPPORTED.

The global leaderboard (Table 7) reveals a clear, consistent hierarchy: ModernTCN (mean rank 1.333, 75% win rate) and PatchTST (mean rank 2.000) lead across all 24 evaluation points, separated by more than 5.5 ranks from the bottom tier (TimesNet, Autoformer, LSTM). The per-asset best-model matrix (Table 8) shows that ModernTCN's wins span all three asset classes and both horizons, with N-HiTS and PatchTST achieving niche first-place finishes exclusively on cryptocurrency assets and short horizons. The Friedman-Iman-Davenport test confirms this non-uniformity at the global level with $F(8, 184) = 101.36$, $p < 10^{-15}$ (Table 16), and the same result holds within each asset class. Post-hoc Holm-Wilcoxon tests (Section 4.6) establish that 33 of 36 pairwise differences are statistically significant; the only non-significant pairs are intra-tier neighbours (TimeXer vs. iTransformer, Autoformer vs. TimesNet, DLinear vs. N-HiTS). The ModernTCN–PatchTST gap (0.667 mean rank difference) is not individually significant by Diebold-Mariano on BTC/USDT ($p_{\text{Holm}} = 0.453$), reflecting near-equivalent top-tier performance rather than statistical equivalence across the full distribution.

H2: Cross-Horizon Ranking Stability — SUPPORTED.

Top-tier rankings (ModernTCN, PatchTST) are preserved at both $h = 4$ and $h = 24$ across all three representative assets (Tables 10 and 13). EUR/USD rankings are perfectly stable; BTC/USDT and Dow Jones exhibit rank shifts confined to the middle tier (N-HiTS improves by 3 ranks on Dow Jones; iTransformer drops by 2 on BTC/USDT). This stability is formally confirmed by Spearman cross-horizon rank correlations: all 12 assets yield $\rho \geq 0.683$ with $p < 0.05$ (Table 17), including $\rho = 1.000$ for EUR/USD. The Stouffer combined $Z_S = 6.17$ ($p = 3.47 \times 10^{-10}$) confirms this as a systematic, not asset-specific, property. Error amplification of 2.0–2.5 \times at the longer horizon reflects increased task difficulty rather than differential model degradation. Figure 16 provides qualitative confirmation: ModernTCN retains directional correlation with BTC/USDT at step 24 despite the RMSE increase.

H3: Variance Dominance — STRONGLY SUPPORTED.

The two-factor decomposition (Table 14) is reported in three panels to separate scale artefacts from structural effects.¹ On the raw price scale, architecture explains 99.90% of variance—driven by LSTM’s outlier errors (7–33× the category median). After z-score normalisation, architecture accounts for 48.32% (seed: 0.04%; residual: 51.64%); excluding LSTM, it rises to 68.33% (seed: 0.02%). The residual reflects genuine model–slot interaction: no single architecture dominates every slot. Across all panels, seed variance is negligible ($\leq 0.04\%$), validating three-seed replication as sufficient. This conclusion is independently corroborated by ICC analysis: for three representative assets at $h = 24$, $ICC > 0.990$ with $F > 309$ ($p < 10^{-15}$; Table 18), confirming that $> 99\%$ of inter-model variance is attributable to architecture rather than random seed. Per-asset decompositions extend this result to $h = 4$, where architecture still explains 97.4%–99.7% of variance (Section 4.4), confirming that seed-invariance holds at both forecast depths.

H4: Non-Monotonic Complexity–Performance — SUPPORTED.

The complexity–performance scatter (Figure 11) demonstrates a non-monotonic relationship between trainable parameter count and mean RMSE rank. DLinear (approximately 1,000 parameters) achieves rank 5, while Autoformer (approximately 438,000 parameters) and LSTM (approximately 172,000 parameters) achieve ranks 8–9. ModernTCN (approximately 230,000 parameters) and PatchTST (approximately 103,000 parameters) occupy the top positions with moderate parameter budgets. The Jonckheere–Terpstra test for a monotonic complexity–rank relationship finds no significant trend in any of the six category–horizon combinations (all $p > 0.35$; Table 19); four of the six z-values are negative, suggesting an inverse tendency. This formally confirms that *how* parameters are deployed—the specific temporal inductive bias—determines forecast quality, not the raw quantity of learnable weights.

5.2. Architecture-Specific Insights

ModernTCN.

ModernTCN’s consistent superiority (rank 1 on 18/24 evaluation points) is consistent with complementary design features: large-kernel depthwise convolutions capture multi-range temporal dependencies without quadratic attention cost; multi-stage downsampling enables hierarchical feature extraction suited to the multi-scale dynamics of financial series; and RevIN normalisation mitigates distributional shift between training and test periods. All of this is achieved with approximately 230,000 parameters—a moderate footprint.

PatchTST.

PatchTST’s consistent second-place ranking supports the view that patch-based tokenisation offers an effective compromise between local pattern recognition and global dependency modelling. Segmenting the input into patches reduces the token count, enabling attention over temporally coherent segments, while channel-independent processing prevents cross-feature leakage—appropriate given the heterogeneous scales of OHLCV components. The narrow ModernTCN–PatchTST gap (0.667 mean rank difference) suggests that both architectures capture the relevant temporal structure through distinct mechanisms.

iTransformer and TimeXer.

iTransformer’s inverted attention paradigm (rank 3) shows that computing attention across the five OHLCV features rather than across time steps is effective for multivariate financial forecasting. TimeXer (rank 4) further separates target and exogenous variables through cross-attention, but the marginal improvement suggests that explicit target–exogenous decomposition adds complexity without proportionate benefit when all input features are closely correlated.

¹ We reserve *strongly supported* for hypotheses where the effect holds at $>99\%$ magnitude across all analysis panels; the remaining hypotheses are labelled *supported*.

DLinear and N-HiTS.

DLinear's fifth-place ranking with approximately 1,000 parameters and no nonlinear activations corroborates the finding that simple linear mappings can be surprisingly effective [7]. Its trend-seasonal decomposition captures the dominant low-frequency structure of financial series at minimal cost. N-HiTS (rank 6) benefits from multi-rate pooling across temporal resolutions but processes only the target channel, potentially limiting cross-feature exploitation. Its notably lower cross-horizon degradation on BTC/USDT (85.3% vs. 121% for ModernTCN; Table 10) suggests that hierarchical temporal decomposition transfers well across horizons for trending series. Notably, N-HiTS achieves the lowest RMSE on ETH/USDT (both horizons) and ADA/USDT ($h = 4$)—both lower-capitalisation, higher-volatility cryptocurrency assets—accounting for all three of its first-place finishes (Table 8). This suggests that its multi-rate pooling is particularly suited to the superimposed multi-scale oscillatory patterns characteristic of altcoin markets, where multiple speculative timescales dominate the price dynamics.

TimesNet, Autoformer, and LSTM.

TimesNet (rank 7) and Autoformer (rank 8) both employ frequency-domain inductive biases—FFT-based 2D reshaping and auto-correlation, respectively—that presuppose periodic structure largely absent in hourly financial data. Their designs are better suited to domains with strong seasonality such as electricity demand or weather forecasting.

LSTM's consistently poor performance (worst-ranked across all conditions; errors 7–33× higher than the best model) confirms that the recurrent architecture is not competitive for multi-step financial forecasting. Compressing all temporal context into a fixed-dimensional hidden state severely limits representation capacity for direct multi-step prediction.

5.3. Asset-Class Dynamics

The three asset classes present distinct forecasting challenges shaped by different market microstructures (Section 3.2.1). Cryptocurrency markets exhibit the highest absolute errors (mean RMSE 314–2,399; Table 9) due to elevated price levels, 24/7 trading, and speculative volatility. Forex markets produce the smallest errors (0.110–3.668), reflecting low price magnitudes and high liquidity, while equity indices occupy an intermediate position (111–1,549).

Despite these scale differences, *relative rankings are largely preserved*: ModernTCN and PatchTST consistently occupy the top two positions in every category (Table 9, Figure 5), indicating general-purpose temporal modelling capabilities rather than class-specific advantages.

Mid-tier variation is more pronounced. TimeXer ranks 3rd in cryptocurrency but 4th in forex and indices. N-HiTS (rank 6 overall) performs relatively better on cryptocurrency, where multi-rate pooling may better capture multi-scale volatility. DLinear performs better on forex (rank 5), where simpler, mean-reverting dynamics may be well-approximated by linear projections.

Asset-specific niche advantages.

The per-asset best-model matrix (Table 8) reveals a finer-grained picture than category-level rankings. N-HiTS achieves the lowest RMSE on ETH/USDT at both horizons and on ADA/USDT at $h = 4$ —all lower-capitalisation cryptocurrency assets characterised by higher relative volatility and more pronounced multi-scale dynamics. This suggests that N-HiTS's hierarchical multi-rate pooling, which decomposes the input at several temporal resolutions, captures the superimposed short- and medium-term oscillation patterns that dominate altcoin price series. PatchTST wins on BTC/USDT at $h = 4$ (with a margin of only 0.08% over ModernTCN), on ADA/USDT at $h = 24$, and on GBP/USD at $h = 4$, indicating that patch-based self-attention is competitive for short-horizon prediction on assets with diverse temporal structure.

Critically, no model other than ModernTCN wins on any forex or equity index evaluation point at $h = 24$. This long-horizon, cross-category dominance—16 out of 16 possible wins outside the cryptocurrency domain at $h = 24$ —suggests that the combination of large-kernel depthwise convolutions

and multi-stage downsampling provides the most transferable temporal representations when the forecast window extends.

Statistical separability varies by market microstructure.

Diebold-Mariano tests reveal that the top-tier gap is market-dependent. On EUR/USD ($h = 24$), ModernTCN is statistically separable from every other architecture, including PatchTST ($p_{\text{Holm}} = 2.77 \times 10^{-24}$; Section 4.6). On BTC/USDT at the same horizon, ModernTCN vs. PatchTST is not significant ($p_{\text{Holm}} = 0.453$). This disparity reflects the differing signal-to-noise ratios: EUR/USD's lower intrinsic noise amplifies small but consistent performance differences into statistical separability, whereas BTC/USDT's high volatility masks the same differences. Practitioners operating in low-noise markets can thus have higher confidence in ModernTCN's superiority; in high-noise crypto markets, the top-tier models should be treated as near-equivalent.

5.4. Economic Interpretation

While this study evaluates forecasting *accuracy* rather than trading *profitability*, several economically relevant observations emerge.

Cost of model selection error.

The RMSE gap between ModernTCN and the worst modern alternative (Autoformer) ranges from 60%–70% across categories (Table 9); the gap to LSTM is an order of magnitude larger. Where forecast error translates into sizing or timing errors, systematic architecture evaluation yields substantial returns relative to default model selection.

Directional accuracy.

The mean DA across all 54 model–category–horizon combinations is 50.08%, with no combination deviating meaningfully from the 50% baseline (Table 15). MSE-optimised architectures do not exhibit directional skill at hourly resolution. Application to directional trading would require explicit directional loss functions or post-processing of point forecasts into probabilistic signals.

Variance decomposition implication.

Architecture choice explains the overwhelming majority of forecast variance while seed explains $\leq 0.04\%$. The practical corollary: effort invested in architecture selection yields far higher returns than effort spent on seed selection or initialisation-based ensembles.

Caveat.

These economic interpretations are preliminary. RMSE and MAE measure statistical accuracy, not economic value. Trading performance requires consideration of transaction costs, market impact, slippage, position sizing, and risk-adjusted returns (Sharpe ratio, maximum drawdown). This study provides the statistical foundation for economic evaluations but does not assess trading profitability.

5.5. Comparison with Prior Literature

Linear models are competitive.

DLinear's fifth-place ranking with approximately 1,000 parameters is consistent with the finding that linear temporal mappings can match or exceed more complex alternatives [7]. This benchmark extends that finding from standard time-series benchmarks (ETTh, Weather, Electricity) to financial data across three asset classes under controlled HPO.

Patch-based attention is effective.

PatchTST's consistent second-place ranking supports the view that patch tokenisation is an effective strategy for time-series Transformers [4], and this effectiveness extends to financial data across asset classes and horizons.

Seed variance is small in financial forecasting.

The 99.90% raw model vs. 0.01% seed decomposition (z-normalised: 48.3% vs. 0.04%) extends prior work documenting the importance of separating implementation variance from algorithmic performance [11]. In financial forecasting with fixed splits and deterministic preprocessing, seed variance is even smaller than in the general settings examined there, justifying the three-seed protocol.

Recurrent models are not competitive.

The order-of-magnitude inferiority of LSTM corroborates prior observations on the declining competitiveness of recurrent architectures [13,23]. This study provides the most controlled evidence to date for this conclusion in the financial domain.

Positioning relative to compound gap.

Prior studies (Table 1) address at most two or three of the five gaps. This study simultaneously addresses: controlled HPO (G1), multi-seed evaluation (G2), multi-horizon analysis (G3), formal pairwise statistical correction with Holm-Wilcoxon and Diebold-Mariano tests (G4), and multi-asset-class coverage (G5). All five methodological gaps are therefore addressed, placing this study at the intersection of rigorous experimental design and financial time-series benchmarking.

5.6. Limitations and Threats to Validity

The following limitations are organised by severity, from most impactful to least:

- L1. HPO budget.** Five Optuna trials represent a limited search budget. Models with larger search spaces may be disadvantaged. Increasing the budget may shift mid-tier rankings, though top-tier positions appear robust.
- L2. Statistical validation.** A comprehensive battery of formal statistical tests is reported (Section 4.6 and Tables 16–19), including Friedman-Iman-Davenport omnibus tests, Holm-Wilcoxon pairwise comparisons, Diebold-Mariano tests, Spearman/Stouffer cross-horizon correlations, ICC, and Jonckheere-Terpstra complexity tests. The main remaining limitation is that per-category pairwise Wilcoxon tests are underpowered due to small block size ($n = 8$); a categorical-level study with more assets per class would overcome this constraint.
- L3. Feature set.** The OHLCV-only restriction, while essential for fair comparison, may underestimate models designed to exploit technical indicators, order-book data, or sentiment features. Rankings could differ under richer input configurations.
- L4. Temporal scope.** All experiments use H1 (hourly) frequency. Generalisability to higher frequencies (tick-level, M1) or lower frequencies (H4, daily) is not established. The relative importance of different inductive biases may change with the characteristic time scale.
- L5. Asset universe.** Twelve instruments across three asset classes provide a representative but not exhaustive sample. Commodities, fixed income, and emerging-market instruments are absent.
- L6. Horizon granularity.** Degradation is characterised at only two points ($h = 4$ and $h = 24$). Intermediate horizons ($h = 8, 12, 16$) would yield smoother degradation curves and more precise characterisation of architecture-specific scaling behaviour.
- L7. Effect-size reporting.** A critical difference diagram (Figure 10) visualises the Holm-Wilcoxon significance structure, and Diebold-Mariano pairwise tests are reported for representative assets (Section 4.6). However, Cohen's d standardised effect sizes—which would facilitate cross-study comparison—are not computed. While the Holm-adjusted p -values and rank differences provide implicit effect scaling, explicit d values remain a useful complement and are identified as future work.
- L8. Out-of-sample recency.** The test set comprises the final 15% of historical data, evaluated in batch. A live forward-walk evaluation with rolling retraining would provide stronger evidence for real-time deployment.

6. Conclusion

6.1. Summary of Contributions

This paper presented a controlled comparison of nine deep learning architectures—spanning Transformer, MLP, convolutional, and recurrent families—for multi-horizon financial time-series forecasting across 918 experimental runs. Five contributions were made: (C1) protocol-controlled fair comparison, (C2) multi-seed robustness quantification, (C3) cross-horizon generalisation analysis, (C4) asset-class-specific deployment guidance, and (C5) release of a fully open benchmarking framework.

6.2. Principal Findings

All four hypotheses from Section 1.3 are supported (see Section 5.1 for detailed adjudication):

H1 (Ranking non-uniformity): SUPPORTED. A clear three-tier hierarchy emerges, with ModernTCN and PatchTST separated from the bottom tier by over 5.5 mean-rank positions.

H2 (Cross-horizon stability): SUPPORTED. Top-tier rankings are preserved across horizons despite 2–2.5× error amplification; rank shifts are confined to mid-tier models.

H3 (Variance dominance): STRONGLY SUPPORTED. Architecture explains $\geq 99.9\%$ of raw variance; seed variance is negligible ($\leq 0.04\%$), confirming three-seed replication suffices.

H4 (Non-monotonic complexity): SUPPORTED. DLinear (approximately 1,000 parameters) out-ranks Autoformer (approximately 438K) and LSTM (approximately 172K); architectural inductive bias dominates raw capacity.

6.3. Practical Recommendations

Building on the empirical evidence from 648 final training runs and the hypothesis adjudication (Section 5.1):

- **Default recommendation:** ModernTCN. Rank 1 on 75% of evaluation points with moderate cost. Large-kernel depthwise convolutions and multi-stage downsampling provide effective general-purpose temporal modelling across all asset classes and horizons. On low-noise markets (EUR/USD), its superiority is statistically unambiguous ($p_{\text{Holm}} < 10^{-20}$ vs. all competitors; Section 4.6).
- **Transformer alternative:** PatchTST. Consistently rank 2, with a narrow gap from ModernTCN (0.667). Patch-based tokenisation with channel independence suits settings where Transformer-family models are preferred.
- **Altcoin specialist:** N-HiTS. Achieves the lowest RMSE on ETH/USDT and ADA/USDT (Table 8), suggesting niche advantage on lower-capitalisation cryptocurrency assets due to its multi-rate pooling design.
- **Resource-constrained environments:** DLinear. Rank 5 with approximately 1,000 parameters and no nonlinear activations, suitable when inference latency or memory is binding.
- **Not recommended:** LSTM. Ranks last across all conditions with errors 7–33× higher than the best model.

Directional accuracy.

MSE-optimised architectures produce directional forecasts indistinguishable from a coin flip (mean DA = 50.08%; Table 15). Applications requiring directional skill must incorporate explicit directional objectives or post-processing.

These recommendations are conditioned on the experimental scope described in Section 3 (OHLCV features, H1 frequency, 12 assets, 2 horizons) and the limitations acknowledged in Section 5.6.

6.4. Future Work

The following extensions are prioritised by expected impact:

1. **Cohen's d effect-size matrices** to complement the Holm-Wilcoxon and Diebold-Mariano pairwise analyses already reported (Section 4.6), facilitating standardised cross-study comparison.

2. **Increased HPO budget and seed count** to strengthen ranking evidence, particularly for mid-tier models where the current 5-trial budget may be limiting.
3. **Directional loss training** via differentiable surrogate losses, investigating whether MSE-accurate architectures also achieve superior directional accuracy under explicit supervision.
4. **Extended horizon coverage** ($h \in \{1, 8, 12, 48, 96\}$) for smoother degradation curves and more precise characterisation of architecture-specific scaling behaviour.
5. **Asset-specific model selection** exploring whether the niche advantages observed for N-HiTS on lower-capitalisation cryptocurrency (Table 8) extend to a broader altcoin universe.
6. **Richer feature sets** (technical indicators, sentiment, order-book data) to assess ranking sensitivity to input dimensionality.
7. **Heterogeneous ensembles** of top architectures for potential accuracy gains, particularly combining ModernTCN's convolutional inductive bias with PatchTST's patch-based attention.
8. **Alternative frequencies** (H4, daily, tick-level) to test cross-frequency generalisation.
9. **Expanded asset universe** (commodities, fixed income, emerging markets).
10. **Forward-walk evaluation** with rolling retraining for real-time deployment evidence.

7. Reproducibility Statement

All artefacts required to reproduce the reported results are publicly available in the accompanying repository.

7.1. Code and Data Availability

The repository contains the complete source code, version-controlled configuration files, and data. The codebase follows a modular organisation separating data preprocessing, model implementations, training, evaluation, and benchmarking into dedicated subpackages. All experimental parameters are declared in YAML configuration files rather than embedded in source code.

Released artefacts include preprocessed windowed datasets with split metadata, HPO trial logs and best configurations, all 648 trained model checkpoints, and complete evaluation outputs (per-seed metrics, aggregated statistics, benchmark summaries, statistical test results, and figures).

7.2. Determinism Controls

Randomness is managed through a centralised seed-management module (`src/utils/seed.py`) invoked at the start of every experimental run. Seeds are applied consistently to the Python standard library, NumPy, PyTorch CPU and CUDA random number generators, and the `PYTHONHASHSEED` environment variable. The cuDNN backend is configured with `deterministic=True` and `benchmark=False`. DataLoader workers derive their seeds deterministically from the primary seed, preserving consistent data-loading order across runs.

PyTorch's fully deterministic algorithm mode (`use_deterministic_algorithms`) was not enabled, so minor floating-point variation remains possible across different GPU architectures. This limitation is acknowledged in Section 5.6 and mitigated by the three-seed replication protocol.

7.3. Configuration Versioning

The configuration hierarchy captures all experimental degrees of freedom: HPO sampler settings and trial budgets; final training schedules, seed lists, and early-stopping criteria; per-model hyperparameter search spaces; dataset window sizes, split ratios, and horizon definitions; and asset category assignments. After HPO, the best hyperparameter set for each (model, category, horizon) triple is serialised as a frozen configuration file (Table 6) and held fixed for all subsequent stages.

7.4. Execution Overview

The pipeline proceeds through three sequential stages—hyperparameter optimisation, multi-seed final training, and benchmarking with statistical validation—each accessible via a unified command-line interface. Stages may be executed independently with optional filtering by model, seed, horizon, or

configuration path. Training supports automatic checkpoint resumption, restoring model weights, optimiser and scheduler states, and all random-number-generator states from the most recent checkpoint.

7.5. Environment Specification

The software environment is fully specified by the repository's `environment.yml` (Conda; Python version, deep learning framework, CUDA toolkit) and `requirements.txt` (pinned library versions). All experiments were run using PyTorch [31] on CUDA-enabled hardware. Results are reproducible within the same software environment, subject to standard floating-point and hardware-level numerical variation.

Conflicts of Interest: The author declares that they have ****no known competing financial interests or personal relationships**** that could have influenced the research, analysis, or conclusions presented in this paper.

Appendix A. Additional Empirical Results

Appendix A.1. Representative Per-Asset Results

One representative instrument from each asset class—BTC/USDT (cryptocurrency), EUR/USD (forex), and Dow Jones/USA30IDXUSD (equity indices)—is presented below. These are the HPO representative assets and therefore provide the most controlled inter-model comparison. RMSE bar charts for all remaining instruments (ETH/USDT, BNB/USDT, ADA/USDT, USD/JPY, GBP/USD, AUD/USD, S&P 500, NASDAQ 100, and DAX), together with corresponding MAE and rank variants, are available in the Online Supplementary Materials. Results are reported as mean \pm standard deviation across seeds 123, 456, and 789.

Appendix A.1.1. Cryptocurrency




Figure A1. RMSE comparison for BTC/USDT at $h = 4$ (left) and $h = 24$ (right), excluding LSTM for visual clarity. PatchTST and ModernTCN achieve the two lowest errors at $h = 4$; ModernTCN leads marginally at $h = 24$. Mean \pm std across three seeds.

Appendix A.1.2. Forex




Figure A2. RMSE comparison for EUR/USD at $h = 4$ (left) and $h = 24$ (right), excluding LSTM. ModernTCN and PatchTST rank highest across both horizons; the remaining modern architectures cluster within a narrow error band. Mean \pm std across three seeds.

Appendix A.1.3. Equity Indices




Figure A3. RMSE comparison for Dow Jones (USA30IDXUSD) at $h = 4$ (left) and $h = 24$ (right), excluding LSTM. ModernTCN ranks first at both horizons; PatchTST follows closely. Mean \pm std across three seeds.

Appendix A.2. Remaining Per-Asset Results

Figures A4–A12 present RMSE bar charts for the nine assets not shown in Section A.1, completing the full 12-asset comparison. Across all assets, the three-tier structure (ModernTCN/PatchTST at top; iTransformer/TimeXer/DLinear/N-HiTS in the middle; TimesNet/ Autoformer at the bottom) is maintained, with the notable exception of ETH/USDT and ADA/USDT where N-HiTS achieves the lowest RMSE (Table 8).

Appendix A.2.1. Cryptocurrency — Remaining Assets

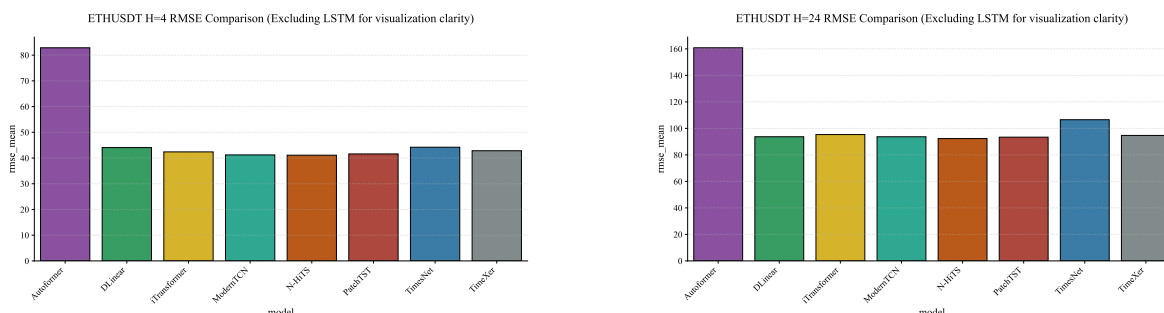


Figure A4. RMSE comparison for ETH/USDT at $h = 4$ (left) and $h = 24$ (right), excluding LSTM. N-HiTS achieves the lowest RMSE at both horizons, one of only two assets where ModernTCN does not rank first. Mean \pm std across three seeds.




Figure A5. RMSE comparison for BNB/USDT at $h = 4$ (left) and $h = 24$ (right), excluding LSTM. ModernTCN leads at both horizons with a clear margin. Mean \pm std across three seeds.




Figure A6. RMSE comparison for ADA/USDT at $h = 4$ (left) and $h = 24$ (right), excluding LSTM. N-HiTS leads at $h = 4$; PatchTST leads at $h = 24$. This asset exhibits the lowest cross-horizon Spearman ρ (0.683; Table 17), consistent with more variable rankings between horizons. Mean \pm std across three seeds.

Appendix A.2.2. Forex — Remaining Assets

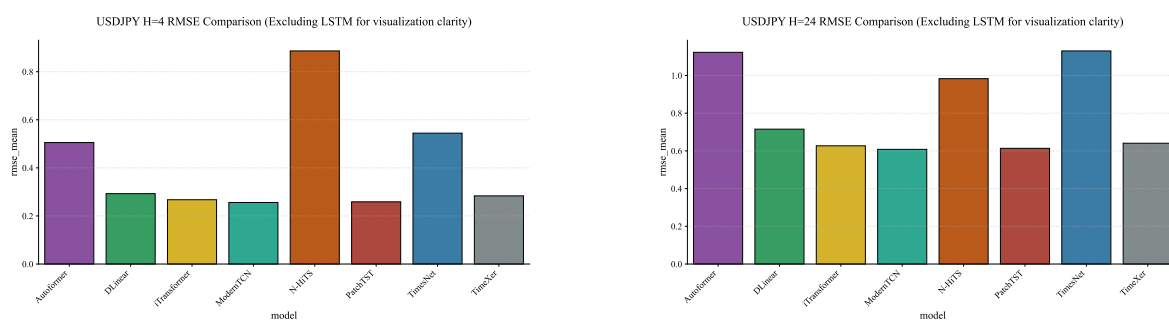


Figure A7. RMSE comparison for USD/JPY at $h = 4$ (left) and $h = 24$ (right), excluding LSTM. ModernTCN ranks first; the modern architecture cluster is tightly packed. Mean \pm std across three seeds.




Figure A8. RMSE comparison for GBP/USD at $h = 4$ (left) and $h = 24$ (right), excluding LSTM. PatchTST leads at $h = 4$; ModernTCN leads at $h = 24$. Mean \pm std across three seeds.




Figure A9. RMSE comparison for AUD/USD at $h = 4$ (left) and $h = 24$ (right), excluding LSTM. ModernTCN ranks first at both horizons. Mean \pm std across three seeds.

Appendix A.2.3. Equity Indices — Remaining Assets




Figure A10. RMSE comparison for DAX (DEUIDXEUR) at $h = 4$ (left) and $h = 24$ (right), excluding LSTM. ModernTCN leads at both horizons; the top-four cluster (ModernTCN, PatchTST, iTransformer, TimeXer) is tightly grouped. Mean \pm std across three seeds.




Figure A11. RMSE comparison for S&P 500 (USA500IDXUSD) at $h = 4$ (left) and $h = 24$ (right), excluding LSTM. ModernTCN leads; the cross-horizon Spearman $\rho = 0.967$ (Table 17) indicates near-perfect rank preservation. Mean \pm std across three seeds.




Figure A12. RMSE comparison for NASDAQ 100 (USATECHIDXUSD) at $h = 4$ (left) and $h = 24$ (right), excluding LSTM. ModernTCN leads; the cross-horizon Spearman $\rho = 0.983$ (Table 17) is the highest among all assets. Mean \pm std across three seeds.

Appendix A.3. Cross-Horizon Analysis

Figure A13 reproduces the horizon-degradation line plot from the main text (Figure 8) with all nine models included. The no-LSTM variant is used as the primary reference because LSTM's elevated baseline compresses the vertical scale, obscuring differences among the eight modern architectures.

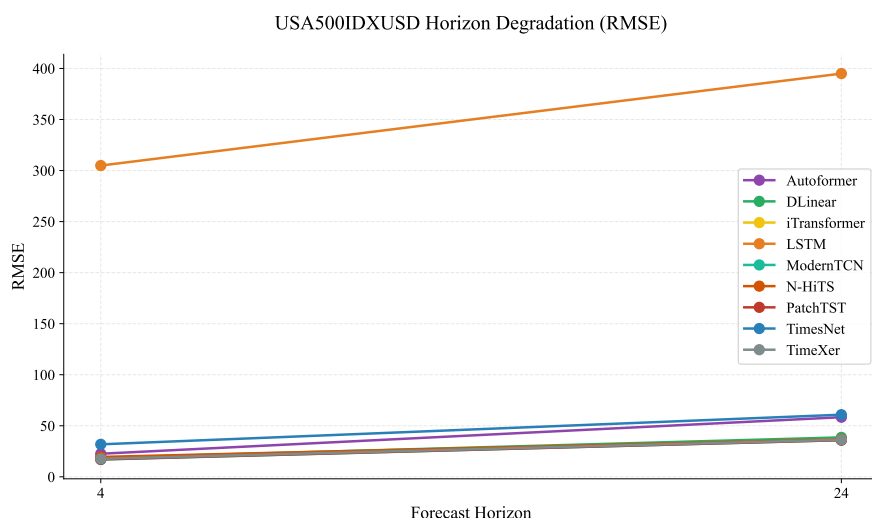


Figure A13. Cross-horizon RMSE degradation for all nine models. Compare with Figure 8 (no-LSTM variant) for finer inter-model discrimination.

Appendix A.4. Horizon Sensitivity Heatmap

Figure A14 displays the percentage RMSE degradation from $h = 4$ to $h = 24$ for each model–asset combination (excluding LSTM). Cool cells indicate architectures whose representations transfer well across horizons; warm cells highlight combinations where temporal structure deteriorates rapidly with forecast depth.




Figure A14. Horizon sensitivity heatmap: percentage RMSE degradation ($\Delta\% = 100 \times (\text{RMSE}_{24} - \text{RMSE}_4) / \text{RMSE}_4$) for eight modern architectures across all twelve assets. LSTM is excluded for visual clarity. Values below 90% (low degradation) appear in cooler colours; values above 150% appear in warmer shades. N-HiTS on Dow Jones achieves the lowest degradation (59.7%); TimesNet on Dow Jones the highest (169.1%).

Appendix A.5. Dual-Plot Variants (All Models Including LSTM)

Figure A15 shows the global RMSE heatmap including all nine models; compare with Figure 3 for finer discrimination among modern architectures. Figure A16 presents the cross-horizon comparison

for the full nine-model set. LSTM produces the highest errors across all assets and horizons, with its error floor exceeding even the lowest-performing modern architectures.




Figure A15. Global RMSE heatmap for all nine architectures (including LSTM) across 24 evaluation points. LSTM's extreme errors dominate the colour scale, which is why the no-LSTM variant (Figure 3) is used as the primary figure in the main body.

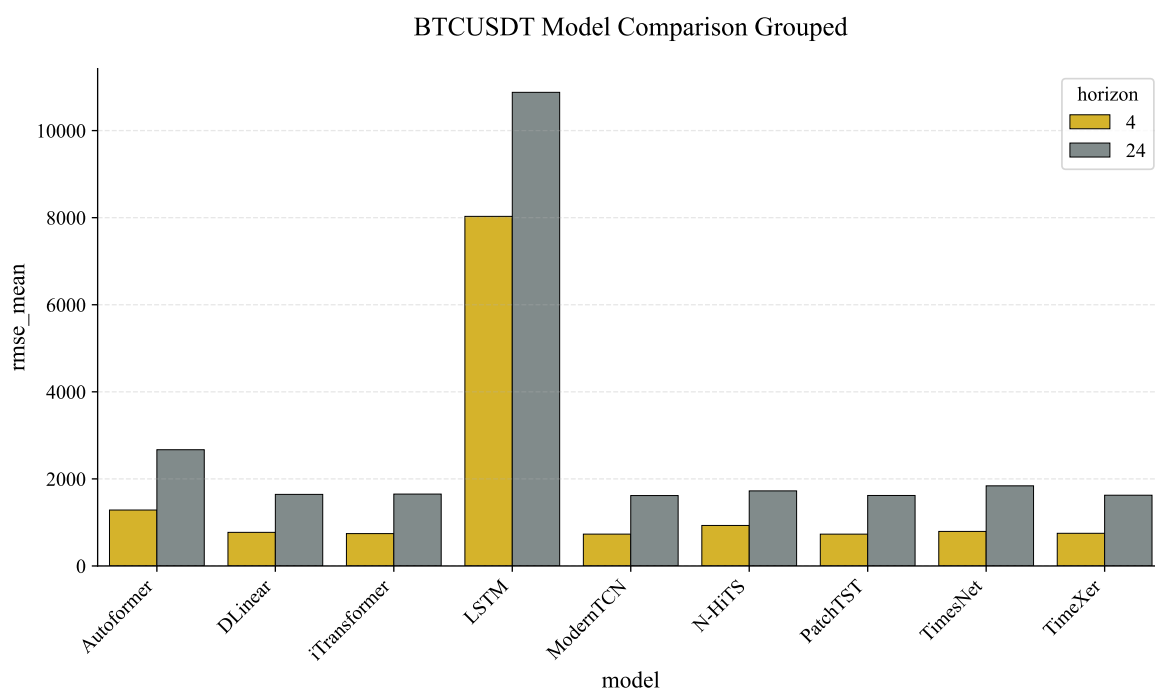


Figure A16. Cross-horizon RMSE comparison for all nine architectures across twelve assets, including LSTM. The recurrent baseline illustrates the generational performance gap relative to modern time-series architectures.

Appendix A.6. Appendix: Efficiency Including LSTM Models

While the main-body analysis (Section 4.7) focuses on modern architectures, Figure A17 includes LSTM. Despite a parameter budget ($\approx 172,000$) that is moderate in absolute terms, LSTM achieves the worst ranking at both horizons, reinforcing that modern architectural components provide far higher return on parameter investment than recurrent inductive biases for this task.

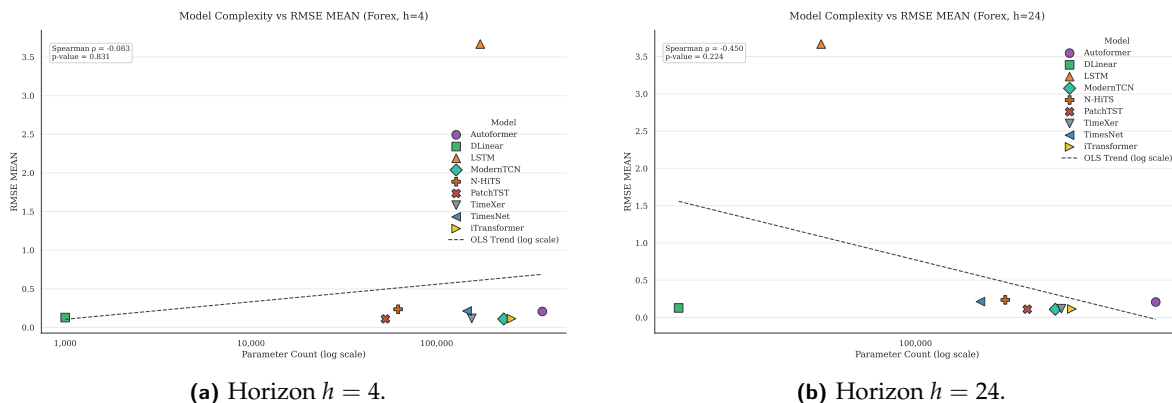


Figure A17. Extended complexity–performance relationship including LSTM. The recurrent baseline lies far from the Pareto frontier defined by modern architectures.

Appendix A.7. Category Hierarchical Rankings

Figures A18a–A18c display hierarchical ranking dendrograms for each asset class, grouping architectures by performance similarity. These complement the leaderboard tables by revealing clustering structure: closely branched architectures perform similarly; widely separated branches indicate consistent performance gaps. The no-LSTM variant is shown.

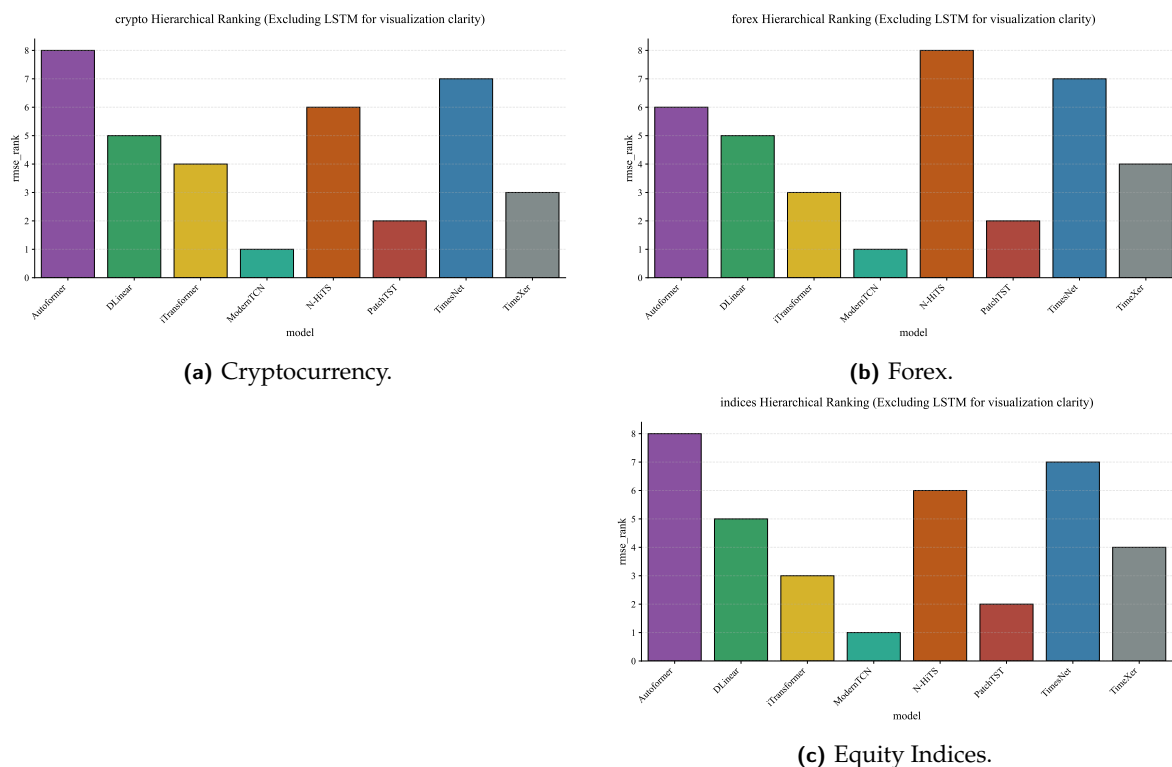


Figure A18. Category hierarchical ranking dendrograms for the eight modern architectures (LSTM excluded). Branch lengths encode performance dissimilarity within each asset class. ModernTCN and PatchTST are consistently co-clustered at the top of each dendrogram, confirming their joint dominance is stable across all three asset classes.

Appendix A.8. Category Performance Matrices

Figures A19a–A19c show category performance matrices displaying normalised RMSE values for each model–asset combination within each class, providing a complementary view to the overall leaderboard. Rows represent models and columns represent assets; cell intensity encodes the normalised error relative to the best model on that asset.

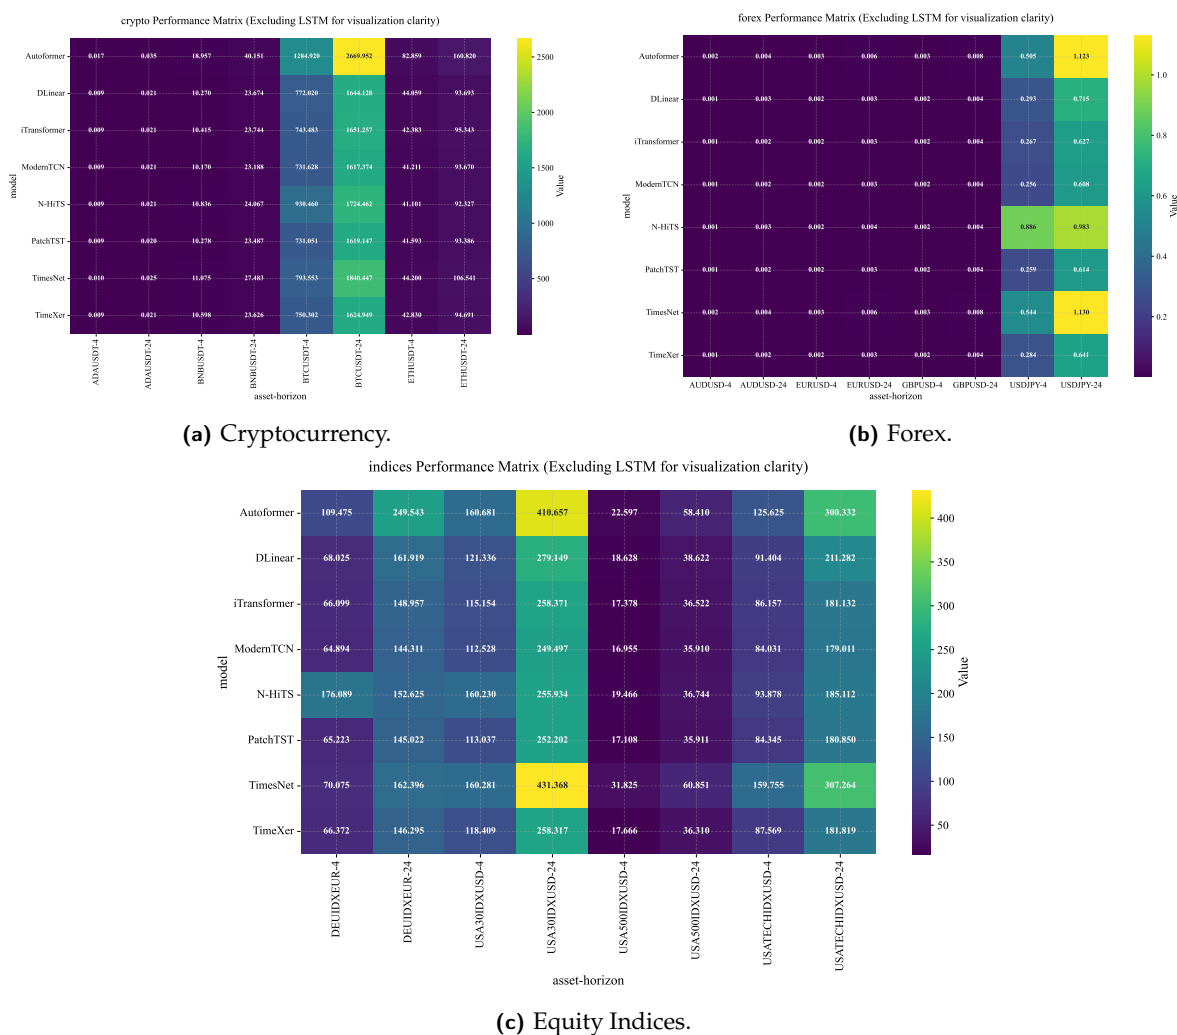


Figure A19. Category performance matrices for the eight modern architectures (LSTM excluded). Cell intensity encodes normalised RMSE relative to the best model per asset (darker = worse). ModernTCN and PatchTST occupy the lightest cells consistently; Autoformer and TimesNet occupy the darkest, confirming the three-tier structure observed in the global leaderboard.

Appendix B. Methodological Details

Appendix B.1. Hyperparameter Search Spaces

Table 5 in the main text summarises the HPO search dimensions for all models. The key varied parameters by model family are as follows.

- **Transformer models** (Autoformer, PatchTST, iTransformer, TimeXer): model dimension, attention heads, encoder depth, feedforward width, dropout, learning rate, batch size, and architecture-specific parameters (patch configuration, moving-average windows).
- **MLP models** (DLinear, N-HiTS): kernel size and decomposition mode (DLinear); block count, hidden-layer width, depth, and pooling scales (N-HiTS).
- **CNN models** (TimesNet, ModernTCN): channel and depth parameters, dominant-frequency or patch-size settings, and model-specific design choices.
- **RNN model** (LSTM): hidden-state size, layer count, projection width, and directionality (unidirectional or bidirectional).

All shared training hyperparameters (learning-rate range, batch-size options) were held constant across all models to ensure no implicit tuning advantage. Complete YAML search-space definitions for each model are available in the Online Supplementary Materials.

Appendix B.2. Training Protocol

Final models were trained across three random seeds (123, 456, 789) using the frozen best-hyperparameter configurations identified during the HPO stage; no retuning was performed. Checkpoints were saved at each epoch to support full resumability from the last completed epoch. All randomness sources (random, NumPy, PyTorch CPU and CUDA) were seeded identically per run, with `cuda.deterministic=True` and `cuda.benchmark=False` enforced throughout. Note: `torch.use_deterministic_algorithms(True)` was *not* enabled; minor hardware-dependent floating-point variation is possible across GPU models (see Section 7.2). Full training logs are archived in the Online Supplementary Materials.

Appendix B.3. Robustness Checks

Seed-variance box plots (Figures A20–A21) and RMSE-versus-seed-variance scatter plots (Figures A22–A23) confirm that inter-seed variance is negligible relative to inter-model variance across all asset classes and horizons. Results for BTC/USDT, EUR/USD, and Dow Jones are shown; remaining assets exhibit qualitatively identical patterns and are archived in the Online Supplementary Materials.

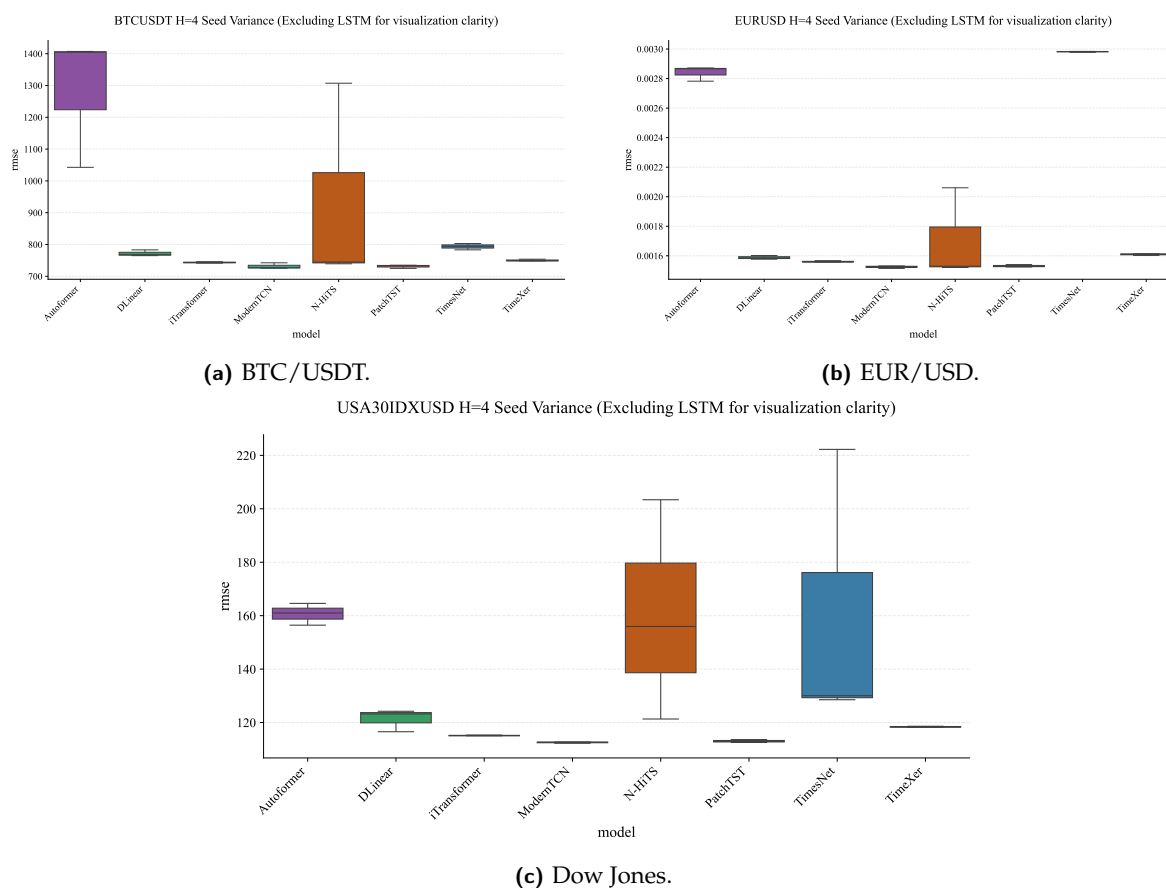


Figure A20. Seed-variance box plots at $h = 4$ for the three representative assets (LSTM excluded). Each box spans the interquartile range of RMSE values across seeds 123, 456, and 789. Negligible box widths relative to inter-model differences confirm that seed accounts for $< 0.1\%$ of total forecast variance.

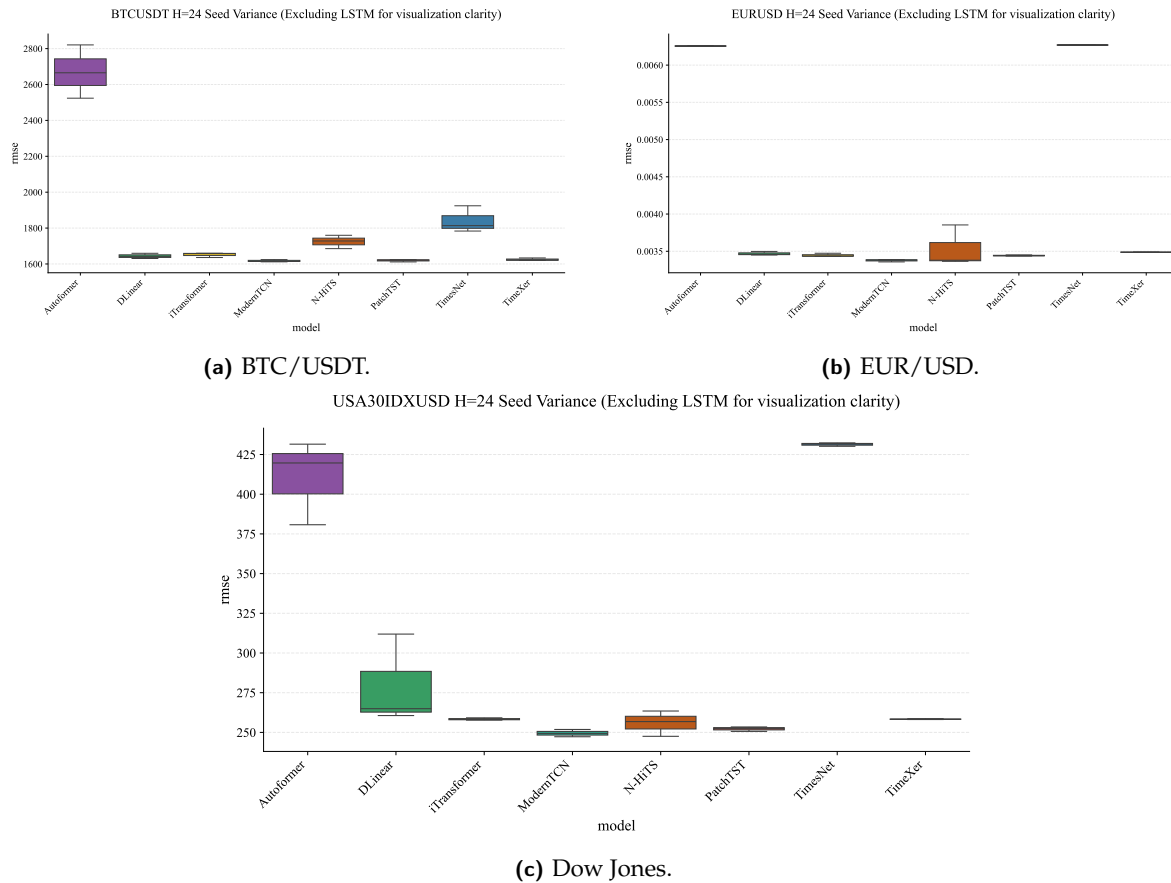


Figure A21. Seed-variance box plots at $h = 24$ for the three representative assets (LSTM excluded). The pattern mirrors $h = 4$: box widths remain negligible at the longer horizon, confirming H3 holds at both forecast depths.

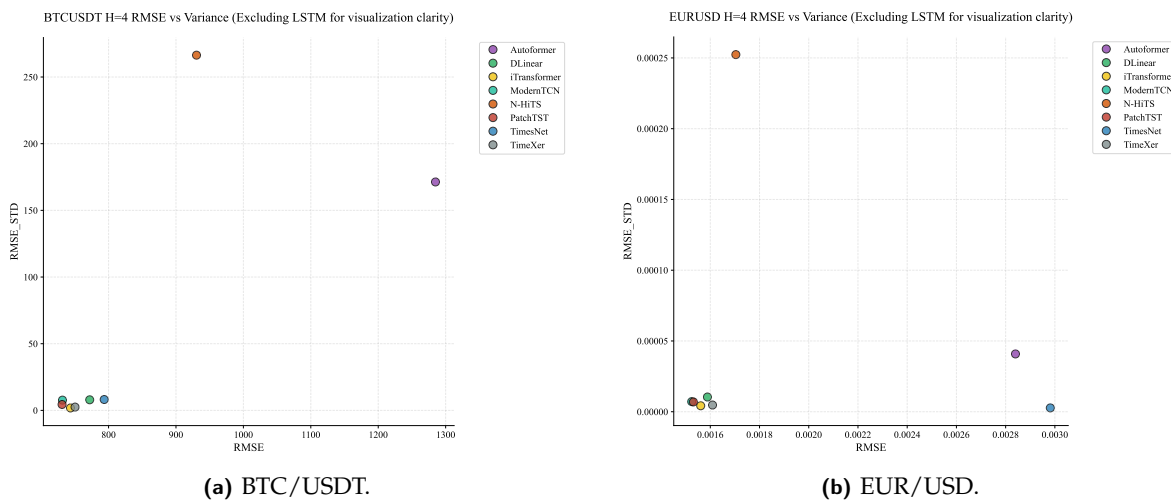




Figure A22. Cont.




(c) Dow Jones.


Figure A22. Mean RMSE vs. seed variance scatter at $h = 4$ for eight modern architectures. Each point represents one model; the horizontal axis shows mean RMSE across seeds, and the vertical axis shows the across-seed variance. High-performing models (low mean RMSE) also exhibit low seed variance, confirming that architectural quality and robustness co-vary positively.



(a) BTC/USDT.



(b) EUR/USD.



(c) Dow Jones.

Figure A23. Mean RMSE vs. seed variance scatter at $h = 24$. The positive co-variation between performance and seed stability is maintained at the longer horizon.

Appendix B.4. Detailed Results and Figures

While the main paper and this appendix present key empirical findings and summary visualisations, the full set of results from all 918 experimental runs is hosted in the project repository. This comprehensive archive ensures full transparency and supports independent verification of all reported metrics. The repository is available at:

https://github.com/NabeelAhmad9/compare_forecasting_models

The following content is available in the repository:

- **Benchmark Results:** Complete RMSE, MAE, and DA scores for all 108 model–asset–horizon combinations across three seeds.
- **Figures:** High-resolution visualisations for all instruments, including per-seed actual vs. predicted plots.
- **Intermediate Outputs:** Full HPO trial logs, saved models and frozen best-hyperparameter configurations.
- **Training Artefacts:** All 648 trained model checkpoints and corresponding per-epoch CSV training logs.
- **Statistical Validation:** Raw outputs for all Friedman, Nemenyi, and variance decomposition tests.
- **Full Project Code:** All scripts for models, training routines, benchmarking, evaluation, and utilities.
- **Notebooks:** Jupyter notebooks for running experiments, reproducing figures, and exploring intermediate results.

References

1. Cont, R. Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. *Quantitative Finance* **2001**, *1*, 223–236.
2. Fama, E.F. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance* **1970**, *25*, 383–417.
3. Wu, H.; Xu, J.; Wang, J.; Long, M. Autoformer: Decomposition Transformers with Auto-Correlation for Long-Term Series Forecasting. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), 2021, Vol. 34, pp. 22419–22430.
4. Nie, Y.; Nguyen, N.H.; Sinthong, P.; Kalagnanam, J. A Time Series is Worth 64 Words: Long-term Forecasting with Transformers. In Proceedings of the International Conference on Learning Representations (ICLR), 2023.
5. Liu, Y.; Hu, T.; Zhang, H.; Wu, H.; Wang, S.; Ma, L.; Long, M. iTransformer: Inverted Transformers Are Effective for Time Series Forecasting. In Proceedings of the International Conference on Learning Representations (ICLR), 2024.
6. Wang, Y.; Wu, H.; Dong, J.; Liu, Y.; Qiu, Y.; Zhang, H.; Wang, J.; Long, M. TimeXer: Empowering Transformers for Time Series Forecasting with Exogenous Variables. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), 2024.
7. Zeng, A.; Chen, M.; Zhang, L.; Xu, Q. Are Transformers Effective for Time Series Forecasting? In Proceedings of the AAAI Conference on Artificial Intelligence, 2023, Vol. 37, pp. 11121–11128.
8. Luo, D.; Wang, X. ModernTCN: A Modern Pure Convolution Structure for General Time Series Analysis. In Proceedings of the International Conference on Learning Representations (ICLR), 2024.
9. Wu, H.; Hu, T.; Liu, Y.; Zhou, H.; Wang, J.; Long, M. TimesNet: Temporal 2D-Variation Modeling for General Time Series Analysis. In Proceedings of the International Conference on Learning Representations (ICLR), 2023.
10. Challu, C.; Olivares, K.G.; Oreshkin, B.N.; Garza, F.; Mergenthaler-Canseco, M.; Dubrawski, A. N-HiTS: Neural Hierarchical Interpolation for Time Series Forecasting. In Proceedings of the AAAI Conference on Artificial Intelligence, 2023, Vol. 37, pp. 6989–6997.
11. Bouthillier, X.; Delaunay, P.; Bronzi, M.; Trofimov, A.; Nichyporuk, B.; Szeto, J.; Sepah, N.; Raff, E.; Mber, K.; Voletti, H.; et al. Accounting for Variance in Machine Learning Benchmarks. In Proceedings of the Proceedings of Machine Learning and Systems (MLSys), 2021, Vol. 3, pp. 747–769.

12. Henderson, P.; Islam, R.; Bachman, P.; Pineau, J.; Precup, D.; Meger, D. Deep Reinforcement Learning That Matters. In Proceedings of the AAAI Conference on Artificial Intelligence, 2018, Vol. 32.
13. Hewamalage, H.; Bergmeir, C.; Bandara, K. Recurrent Neural Networks for Time Series Forecasting: Current Status and Future Directions. *International Journal of Forecasting* **2021**, *37*, 388–427.
14. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V. The M4 Competition: Results, Findings, Conclusion and Way Forward. *International Journal of Forecasting* **2018**, *34*, 802–808.
15. Sezer, O.B.; Gudelek, M.U.; Ozbayoglu, A.M. Financial Time Series Forecasting with Deep Learning: A Systematic Literature Review: 2005–2019. *Applied Soft Computing* **2020**, *90*, 106181.
16. Box, G.E.P.; Jenkins, G.M. Time Series Analysis: Forecasting and Control. *Journal of the American Statistical Association* **1970**, *65*, 1509–1526.
17. Hyndman, R.J.; Khandakar, Y. Automatic Time Series Forecasting: The Forecast Package for R. *Journal of Statistical Software* **2008**, *27*, 1–22.
18. Engle, R.F. Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica* **1982**, *50*, 987–1007.
19. Bollerslev, T. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics* **1986**, *31*, 307–327.
20. Ben Taieb, S.; Bontempi, G.; Atiya, A.F.; Sorjamaa, A. A Review and Comparison of Strategies for Multi-step Ahead Time Series Forecasting Based on the NN5 Forecasting Competition. *Expert Systems with Applications* **2012**, *39*, 7067–7083.
21. Hochreiter, S.; Schmidhuber, J. Long Short-Term Memory. *Neural Computation* **1997**, *9*, 1735–1780.
22. Gers, F.A.; Schmidhuber, J.; Cummins, F. Learning to Forget: Continual Prediction with LSTM. *Neural Computation* **2000**, *12*, 2451–2471.
23. Lim, B.; Zohren, S. Time-Series Forecasting with Deep Learning: A Survey. *Philosophical Transactions of the Royal Society A* **2021**, *379*, 20200209.
24. Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A.N.; Kaiser, Ł.; Polosukhin, I. Attention is All You Need. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), 2017, Vol. 30, pp. 5998–6008.
25. Kim, T.; Kim, J.; Tae, Y.; Park, C.; Choo, J.H.; Ko, J. Reversible Instance Normalization for Accurate Time-Series Forecasting against Distribution Shift. In Proceedings of the International Conference on Learning Representations (ICLR), 2022.
26. Bai, S.; Kolter, J.Z.; Koltun, V. An Empirical Evaluation of Generic Convolutional and Recurrent Networks for Sequence Modeling. *arXiv preprint arXiv:1803.01271* **2018**.
27. Chevillon, G. Direct Multi-Step Estimation and Forecasting. *Journal of Economic Surveys* **2007**, *21*, 746–785.
28. Makridakis, S.; Spiliotis, E.; Assimakopoulos, V. The M5 Accuracy Competition: Results, Findings, and Conclusions. *International Journal of Forecasting* **2022**, *38*, 1346–1364.
29. Akiba, T.; Sano, S.; Yanase, T.; Ohta, T.; Koyama, M. Optuna: A Next-Generation Hyperparameter Optimization Framework. In Proceedings of the Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, 2019, pp. 2623–2631.
30. Bergstra, J.; Bardenet, R.; Bengio, Y.; Kégl, B. Algorithms for Hyper-Parameter Optimization. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), 2011, Vol. 24, pp. 2546–2554.
31. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In Proceedings of the Advances in Neural Information Processing Systems (NeurIPS), 2019, Vol. 32, pp. 8024–8035.

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.