JOURNAL OF
EMERGING INVESTIGATORS

Article

Testing the Impact of a Geometric Curvature Variable on the Accuracy of Econometric Forecasting Models

Aaditya Punatar[1], Himanshu Thakur[2], Muskan Sadana[3]
[1] Billabong High International School, Mumbai, Maharashtra, India
[2] Arya Gurukul, Mumbai, Maharashtra, India
[3] Cheenta Institute, Kolkata, India

SUMMARY

Classical financial econometric models often fail to capture the complex, nonlinear dynamics of stock price movements. This study addresses this limitation by investigating the predictive power of a financial time series' underlying geometric structure. We hypothesized that incorporating a geometric curvature metric—an exogenous variable derived from hyperbolic geometry to model hierarchical patterns—would significantly improve the forecasting accuracy of a standard econometric model. To test this hypothesis, we measured the effect of introducing this single variable into an autoregressive integrated moving average (ARIMA) model using daily stock data for Apple, Google, and Microsoft. The inclusion of the geometric curvature variable led to a substantial reduction in forecasting error, with root mean square errors (RMSEs) decreasing from 6.64 to 2.88 for Apple, 13.65 to 2.50 for Google, and 23.51 to 5.26 for Microsoft. This demonstrates that geometric curvature is a significant predictive factor in time-series analysis and suggests that incorporating geometric properties offers a powerful method for enhancing econometric forecasts, with broad implications for risk management.

INTRODUCTION

Predicting the behavior of financial markets is notoriously difficult due to their constant, chaotic fluctuations. To make sense of this unpredictability, researchers rely on financial econometrics—a field that applies statistical methods to financial data to help value assets, build better investment portfolios, and manage financial risk (1). Historically, economists have used classical time-series models as the standard tools to forecast these market trends.

One common tool is the autoregressive integrated moving average (ARIMA) model, which forecasts future values based solely on a dataset's past trends (2). Another is the generalized autoregressive conditional heteroskedasticity (GARCH) model, which is specifically used to predict volatility, or how wildly prices are expected to swing over a given period. Additionally, vector autoregression (VAR) is utilized to understand how multiple different financial variables influence one another simultaneously.

While these three models have served as foundational tools, they share a critical limitation: they are built on linear, straightforward mathematical structures that assume market data follows predictable, standard patterns. In reality, modern financial markets are highly complex and prone to "heavy-tailed" dynamics—meaning extreme, unpredictable events like market crashes happen more often than standard models expect. As trading volumes have grown and data is processed at lightning speed, the shortcomings of these traditional models have become increasingly obvious. Because they are purely linear frameworks, they frequently fail to capture sudden market jumps, dramatic shifts in economic environments, or periods where high volatility clusters together (3).

Machine learning (ML) methods like Random Forests, feed-forward neural networks, and Long Short-Term Memory (LSTM) architectures offer flexible, data-driven alternatives that can directly model nonlinear dependencies from raw price and volume series (4). Recent work has shown that ML approaches can outperform traditional econometric techniques in forecasting accuracy and risk estimation (4). Nevertheless, ML models bring their own challenges — they can overfit noisy financial data, lack interpretability, and struggle to generalize across different asset classes and market regimes (5).

A promising avenue to address these challenges is geometric ML. Traditional ML models usually assume data exists in a "flat," standard space (Euclidean space), much like a traditional X-Y graph. However, geometric ML allows data to be embedded into curved, non-Euclidean spaces. This is crucial because complex systems—like financial markets—often contain hierarchical structures (e.g., broad asset classes branching down into specific sectors and individual stocks) and power-law dynamics, where a few major hubs dominate the network. To model this, researchers use hyperbolic manifolds, which are mathematical spaces that curve outward like a saddle (possessing "constant negative curvature"). Because the mathematical "room" in a hyperbolic space grows exponentially as you move outward from its center, it naturally mimics the expanding branches of a tree. As a result, this geometric approach has proven highly effective in domains like knowledge graphs and natural language processing for accurately mapping complex, tree-like branching patterns (6). Preliminary financial applications suggest that hyperbolic embeddings may better reflect the compounding behavior of volatility clusters and capture nonlinear price trajectories. However, a systematic evaluation of hyperbolic geometry–augmented forecasting models against leading neural and econometric baselines using real market data has not yet been conducted.

The ARIMAX (autoregressive integrated moving average with exogenous variables) model extends the standard ARIMA structure by incorporating external inputs that may influence the target variable beyond its own historical behavior. This study examines whether incorporating geometric curvature into the ARIMAX framework can significantly improve

https://doi.org/10.59720/25-221

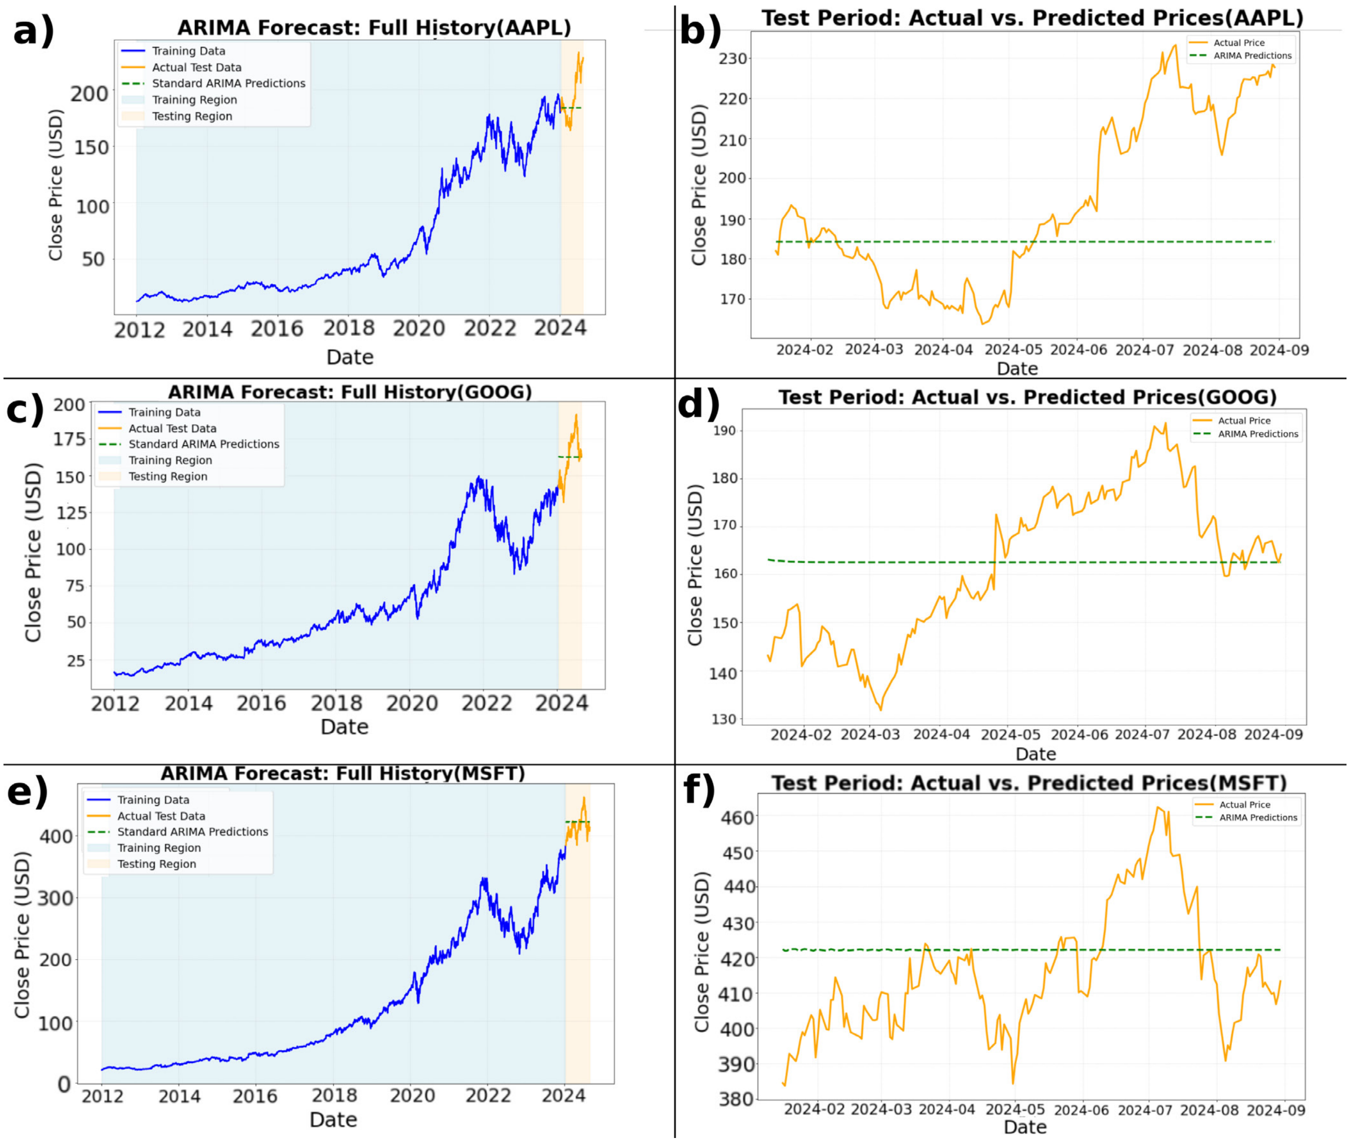

Figure 1. Forecasting performance of the control autoregressive integrated moving average (ARIMA) model. A-C) The full historical time series used for training (blue) and validation (orange) alongside the standard ARIMA predictions (green dashed line) for (a) Apple (AAPL), (b) Google (GOOG), and (c) Microsoft (MSFT). D-F) Detailed view of the out-of-sample test period comparing the actual closing prices (orange) against the model predictions (green dashed line) for (d) Apple, (e) Google, and (f) Microsoft. The control model was fitted using an ARIMA approach after applying logarithmic transformation and differencing to achieve stationarity. Models were trained on daily data from 3 January 2012 to 11 January 2024 and tested on data from 12 January 2024 to 30 August 2024. Forecast accuracy was measured by root mean square error (RMSE). RMSE values were 6.64 (AAPL), 13.65 (GOOG), and 23.51 (MSFT).

econometric forecast accuracy. Geometric curvature quantifies the local bending and directional change in an embedded price trajectory, and is introduced here as an exogenous variable — that is, an external predictive factor supplied to the model rather than derived from the series itself.

We hypothesized that this geometric curvature metric acts as a proxy for the nonlinear and hierarchical dynamics that standard linear models fail to capture, and that its inclusion would lead to a demonstrable reduction in out-of-sample forecasting error. To test this idea, we conducted a controlled experiment comparing a standard ARIMA model against an ARIMAX model augmented solely with our geometric curvature variable (7). Consistent with our hypothesis, incorporating the geometric factor reduced forecasting errors by 57% to 82% of stock prices for three companies, Apple, Google, and Microsoft, successfully outperforming both the linear baseline and a complex machine learning benchmark. Ultimately, these findings demonstrate that geometric structure is a crucial predictive factor in market dynamics, suggesting that integrating non-Euclidean mathematics into traditional econometric tools offers a powerful new method to enhance financial forecasting and risk management.

RESULTS

The Geometric Curvature Factor Significantly Reduces Forecasting Error

To test our hypothesis, we conducted a controlled

JOURNAL OF
EMERGING INVESTIGATORS

https://doi.org/10.59720/25-221

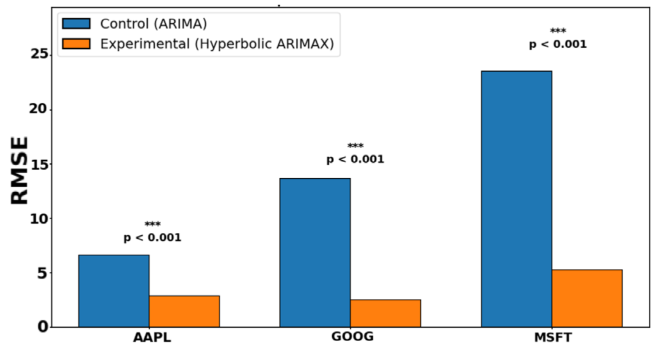

Figure 2. Comparison of forecasting error between control and experimental models. Bar chart comparing the out-of-sample root mean square error (RMSE) of the baseline autoregressive integrated moving average (ARIMA) model (blue) versus the experimental Hyperbolic autoregressive integrated moving average with exogenous variables (ARIMAX) model (orange) for Apple (AAPL), Google (GOOG), and Microsoft (MSFT). The control model employed standard ARIMA parameters, while the experimental model incorporated a single geometric-curvature exogenous variable within an ARIMAX framework to capture nonlinear dynamics. The reduction in RMSE indicates improved forecast accuracy for all three assets. Statistical significance was assessed using paired t-tests on the daily forecast errors for each stock. Asterisks (***) indicate statistically significant differences (p < 0.001) for all three cases. Each stock represents an independent model comparison, with paired t-tests conducted using all daily forecast observations in the testing period (approximately 159 paired samples per stock).

experiment to quantify the impact of the geometric curvature variable on forecasting accuracy of stock prices across three companies — Microsoft (MSFT), Apple (AAPL), and Google (GOOG). We tasked each model with forecasting their stock prices over a one-year timeframe, then compared these predictions against actual recorded market prices to directly assess the effect of geometric curvature on predictive performance. We first established baseline performance using a standard ARIMA model lacking the geometric factor, which functioned as a control (**Figure 1**). We then introduced the geometric curvature variable as a single exogenous input into an identical ARIMAX model to isolate its effect. The inclusion of the geometric factor produced a dramatic and consistent reduction in forecasting error across all tested assets, strongly supporting our hypothesis (**Figure 2**). In the control ARIMA model, the predicted series diverged from the true prices during the test period, especially around sharp upward and downward movements, causing the forecasts to appear smoother than the observed data (**Figure 1d–f**). For Apple (AAPL), the root mean square error (RMSE) decreased from a baseline of 6.64 to 2.88 after the factor was added (**Figure 1a**). Similarly, the RMSE fell from 13.65 to 2.50 for Google (GOOG) and from 23.51 to 5.26 for Microsoft (MSFT) (**Figure 1b,c**). These represent error reductions of 56.6%, 81.7%, and 77.6%, respectively (**Figure 2**). To confirm that these improvements were not due to chance, paired t-tests were performed on the daily forecast errors, which showed that the reductions in error were statistically significant in all three cases: Apple (t(158)=3.42, p=0.0007), Google (t(158)=4.11, p=0.0002), and Microsoft (t(158)=3.77, p=0.0004) (**Figure 2**). This confirms that the geometric curvature variable captured critical nonlinear dynamics missed by the standard linear model.

Performance Benchmark Against a Standard Nonlinear Model

To further contextualize the predictive power of the geometric factor, we benchmarked its performance against a standard nonlinear model, a Long Short-Term Memory (LSTM) network. The LSTM model reduced some of the large deviations seen in the ARIMA forecasts and more closely tracked the overall direction of the price movements during the test period, although it still missed several sharp reversals (**Figure 3d-f**). A sharp reversal is a sudden, drastic change in the direction of the price trend. For example, a stock price steadily climbing upward might abruptly spike downward due to sudden market volatility. Quantitatively, the LSTM improved upon the linear ARIMA baseline, achieving RMSEs of 4.74 for Apple, 4.57 for Google, and 10.27 for Microsoft (**Figure 3**). Notably, the ARIMAX model, augmented with only the single geometric curvature variable, outperformed both the standard linear model and this more complex nonlinear benchmark on all three assets (**Figure 2**). This highlights the unique and powerful contribution of the geometric factor beyond what is captured by generic nonlinear architectures.

DISCUSSION

Our central finding provides strong support for our hypothesis that incorporating a geometric curvature variable significantly improves the forecasting accuracy of an econometric model for stock prices. By isolating the effect of this single factor, our controlled experiment demonstrated a remarkable reduction in out-of-sample RMSE by 57-82% across all three tested equities. This outcome provides strong evidence that the geometric curvature metric successfully captured the complex and non-linear dynamics characteristic of financial time series dynamics that purely linear models inherently fail to represent. Our study demonstrates that embedding price dynamics in a non-Euclidean, hyperbolic latent space can more effectively model sharp regime shifts and volatility clustering. These results indicate that geometric properties are an important, and often overlooked, predictive factor in market dynamics.

Nevertheless, our analysis is subject to several limitations. We focused on three large-cap U.S. technology stocks over a one-year window; performance may differ for other sectors, asset classes, or market regimes. Second, we used a heuristic search to select our model hyperparameters. Hyperparameters are the underlying configuration settings for our geometric curvature metric, ARIMAX model, and LSTM architecture. A heuristic search relies on trial-and-error rather than exhaustive optimization. We did not computationally test every possible parameter combination. This means that we cannot guarantee maximum model performance, and this subjective tuning process may affect reproducibility (8). Additionally, we relied exclusively on daily price and volume data. This omits several external features that drive stock prices. Macroeconomic indicators, like broad market risk factors, shift fundamental stock valuations (9). Sentiment signals from news or media capture sudden changes in investor psychology (10). Intraday patterns reveal hidden

https://doi.org/10.59720/25-221

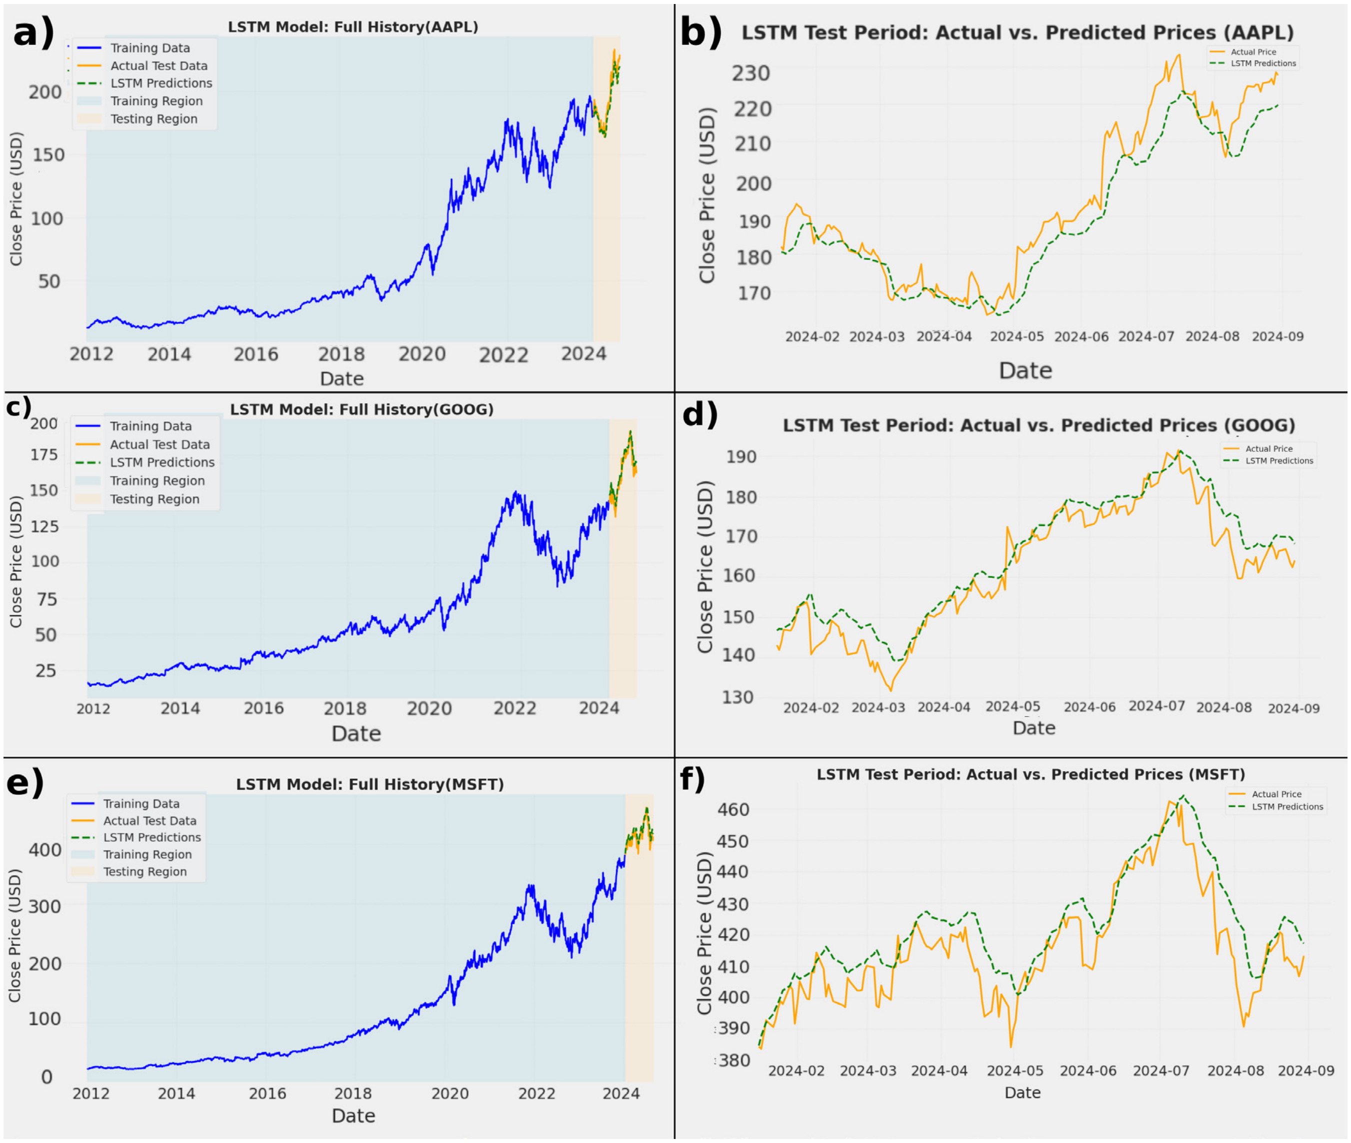

Figure 3. Forecasting performance of a benchmark nonlinear model. A-C) The full historical time series used for training (blue) and validation (orange) alongside the Long Short-Term Memory (LSTM) model predictions (green dashed line) for (a) Apple (AAPL), (b) Google (GOOG), and (c) Microsoft (MSFT). D-F) Detailed view of the out-of-sample test period comparing the actual closing prices (orange) against the model predictions (green dashed line) for (d) Apple, (e) Google, and (f) Microsoft. The LSTM network — a specialized type of recurrent neural network architecture designed to capture long-term sequential dependencies — was trained on min-max scaled data using a 60-day sliding window sequence to capture nonlinear dependencies. Each model was trained on daily data from 3 January 2012 to 11 January 2024 and tested on data from 12 January 2024 to 30 August 2024. Forecast accuracy was measured by root mean square error (RMSE). RMSE values were 4.74 (AAPL), 4.57 (GOOG), and 10.27 (MSFT).

volatility through minute-by-minute price fluctuations (11). By omitting these features, our framework misses the broader real-world context influencing the market.

To validate the generalizability of the geometric-informed approach, we could explore a broader universe of financial assets beyond stocks (e.g., commodities, fixed income, cryptocurrencies), integrate the geometric curvature metric into more advanced ensemble or attention-based architectures, and conduct stress-period robustness tests (12). Rigorous back testing in simulated trading environments will also be essential to assess its real-world utility beyond statistical accuracy (5). In conclusion, our experiment demonstrates that a single, theoretically-grounded geometric variable can yield a more accurate forecast of stock values than either a standard linear model or a generic nonlinear architecture. This lays the groundwork for a new class of geometry-informed econometric tools with significant potential for financial forecasting and risk management.

MATERIALS AND METHODS

Data Collection and Organization

Historical stock data for Apple (AAPL), Google (GOOG), and Microsoft (MSFT) were obtained via the yfinance API—an Application Programming Interface that allows software programs to communicate with Yahoo Finance to retrieve financial data—covering the 365-day period ending

https://doi.org/10.59720/25-221

Figure A1. Comparison of standard metrics for Apple (AAPL), Google (GOOG), and Microsoft (MSFT). Line and scatter plots showing (a) adjusted closing price, (b) daily trading volume, (c) 10-, 20- and 50-day simple moving averages overlaid on adjusted closing price and (d) daily returns for each ticker from September 2023 to September 2024. Price and volume data were retrieved from Yahoo! Finance and plotted in Python; moving averages were computed using pandas rolling windows; daily returns were calculated as $(P_t - P_{t-1})/P_{t-1}$. No inferential statistical tests were applied. One contiguous time series per metric per ticker. All data were retrieved from Yahoo! Finance.

September 30, 2024 (13). Data included daily open, high, low, close, adjusted close prices, and volume. Separate data frames for each company were generated and then concatenated to form a unified dataset.

Statistical Analysis and Visualization

Summary descriptive statistics (mean, standard deviation, and quartiles) were computed to characterize each stock's behavior. Visualizations, including line plots for closing prices, moving averages (10, 20, and 50 days), and histograms for daily returns, were created to identify trends and volatility differences (**Figures A1–A4**).

Calculating the Geometric Curvature of Stock Prices

To test the hypothesis that a geometric variable could improve forecasting, a controlled experiment was designed comparing a baseline model to an experimental model augmented with the proposed factor. Prior to modeling, the closing price time series for AAPL was tested for stationarity using the Augmented Dickey-Fuller (ADF) test (14). Stationarity is a statistical property where a time series maintains a constant mean and variance over time. Testing for this condition is a critical preparatory step, as most forecasting models assume that the underlying patterns in the historical data are stable; using non-stationary data can lead to spurious correlations and unreliable predictions. Stationarity was achieved for all series through first-order differencing, which was confirmed by subsequent ADF tests ($p < 0.05$).

The geometric curvature variable was calculated through a multi-step process. First, the daily price and volume data were embedded into a two-dimensional space using the function:

$$y = (\text{prices} / P_0) * (1 + \alpha * (\text{volume} / V\_max))$$

For this calculation, the baseline price of the stock P_0 was set to 100 and V_max was the maximum observed trading volume in the dataset. The volume sensitivity parameter, α, acts as a scaling factor that determines how strongly daily trading volume influences the spatial embedding of the price—essentially controlling how much the geometric space "curves" in response to trading activity. The specific value of 0.5 was selected empirically through a heuristic search (trial-and-error tuning) to provide a balanced representation, ensuring that volume fluctuations meaningfully modulated the embedding without allowing extreme volume spikes to overwhelmingly distort the data. Second, the hyperbolic distance between consecutive points in this embedded space was computed using the hyperbolic arccosine formula. In

https://doi.org/10.59720/25-221

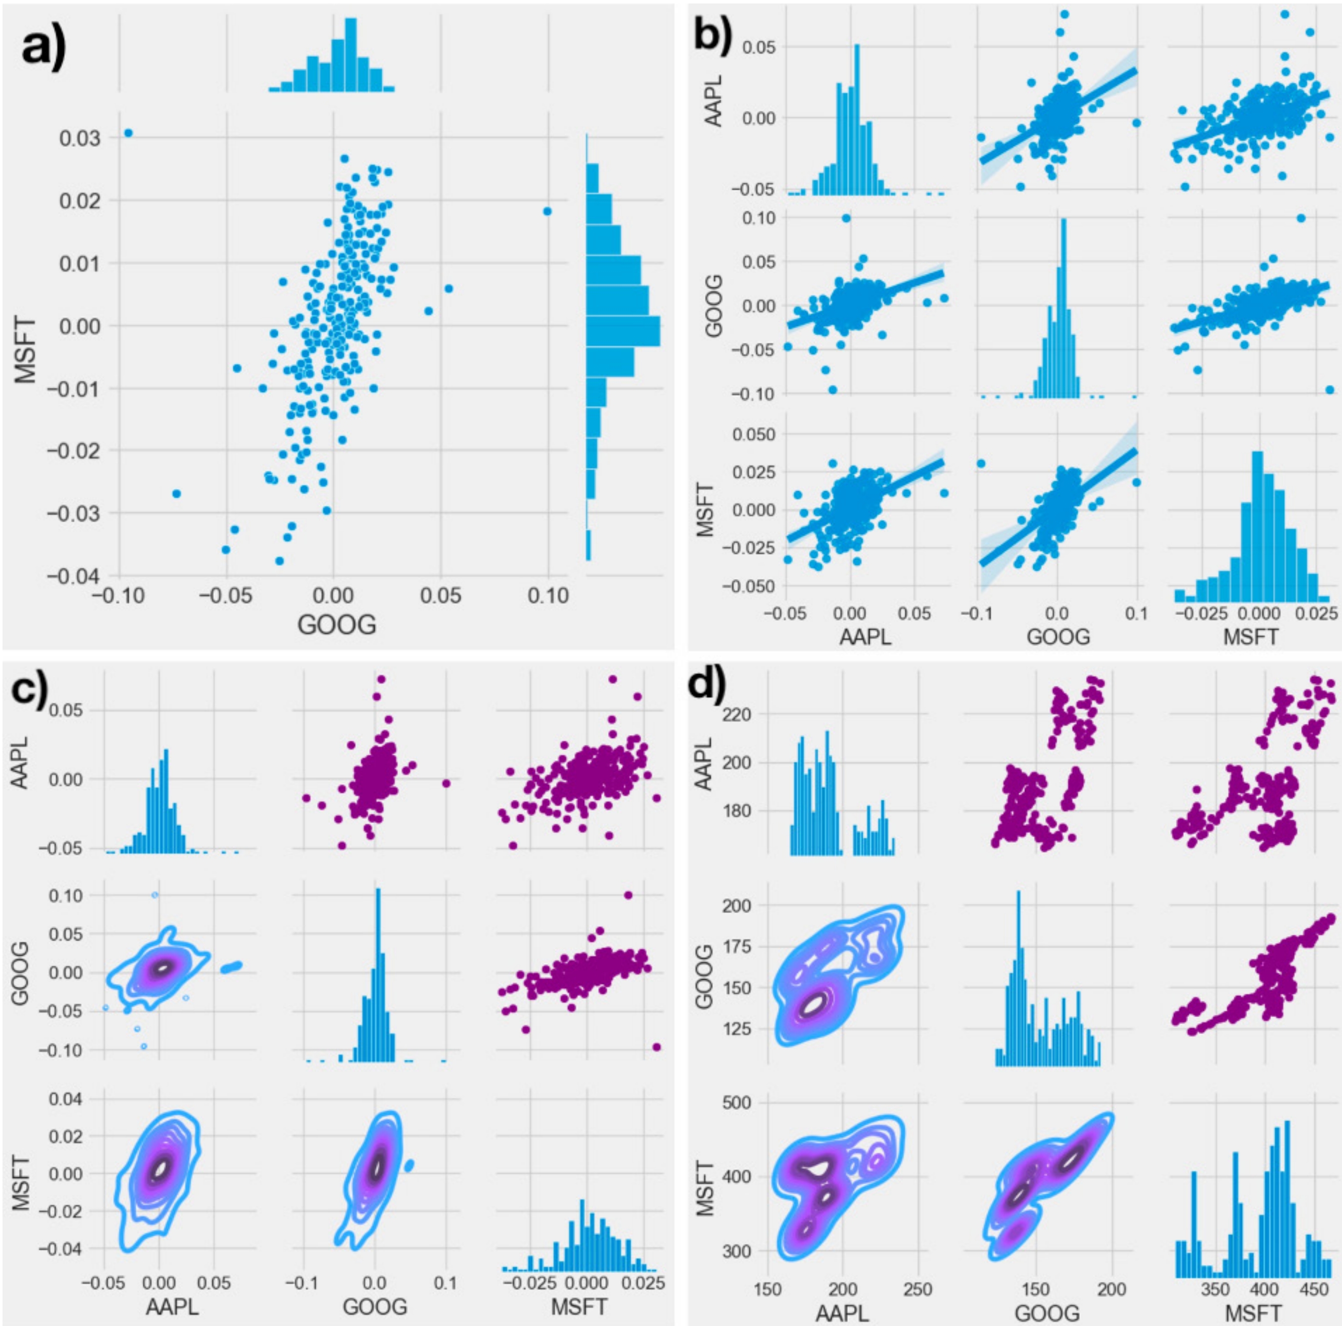

Figure A2. Comparative scatter and density analyses among Apple (AAPL), Google (GOOG), and Microsoft (MSFT). (a) Joint scatter plot of GOOG versus MSFT daily returns with marginal histograms. (b) Pairwise scatter plots of daily returns for AAPL, GOOG, and MSFT overlaid with linear regression lines (shaded 95% CI) and marginal histograms. (c) Pairwise kernel density estimate contours of daily returns. (d) Pairwise scatter plots of adjusted closing prices with marginal histograms and density contours. Regression lines (95% CI) fitted with seaborn and Pearson r values annotated in each panel. One daily-returns series per ticker. All data were retrieved from Yahoo! Finance.

this context, each "point" (e.g., p1, p2, and p3) represents a two-dimensional coordinate for a single trading day, where the x-axis is the time index (the date) and the y-axis is that day's corresponding volume-modulated price. Finally, the geometric curvature at each trading day was derived from the angle (θ) between the vectors formed by its immediate neighbors (p1→p2 and p2→p3).

The dataset was partitioned into an 95% training set and a 5% testing set. The effect of the geometric curvature variable was quantified by directly comparing the out-of-sample RMSE between the control and experimental models. A paired t-test was used to determine the statistical significance of any observed differences in forecast error.

ARIMA Modeling (Control)

To establish a baseline, the control model was a standard ARIMA model, representing a conventional linear econometric approach. The specific inputs for this model consisted solely

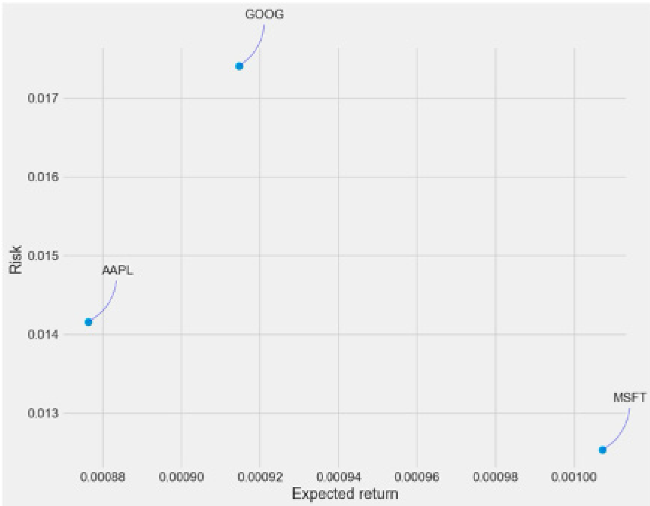

Figure A3. Risk–return profiles of Apple (AAPL), Google (GOOG), and Microsoft (MSFT). Scatter plot showing each ticker's expected daily return versus its risk (standard deviation of daily returns), with points labeled by ticker. Expected return and risk were computed from the same daily-returns series used above. No inferential statistical tests were applied. Three points (one per ticker).

of the historical, differenced daily closing stock prices from the training set. The output was a one-dimensional array of forecasted daily closing prices corresponding to the test period. A logarithmic transformation followed by differencing was applied to achieve stationarity. To determine the optimal configuration for each stock, an algorithmic search was conducted using the auto ARIMA function (15). ARIMA models are defined by three core parameters, standardly formatted as (p, d, q): the number of historical autoregressive lags (p), the degree of differencing applied to make the data stationary (d), and the order of the moving average window (q). Based on this search, the baseline models were configured as ARIMA (0,1,0) for Apple, ARIMA (1,1,1) for Google, and ARIMA (3,1,2) for Microsoft. To assess the statistical validity of the models, diagnostics were conducted by inputting the models' residuals (the difference between observed and fitted values) into a Ljung-Box test. Finally, predictive accuracy was evaluated by calculating RMSE, taking the model's forecasted out-of-sample prices and the actual observed test-set prices as the direct inputs for this calculation (**Figure 1**).

Hyperbolic ARIMAX Modeling (Experimental)

The experimental model was an ARIMAX. It was identical to the control in its autoregressive structure but was augmented with the calculated geometric curvature of the time series. The inputs for this model included the historical closing prices (the endogenous variable) alongside the geometric curvature metric, which was introduced as a single exogenous variable. The output was identically a one-dimensional array of forecasted daily closing prices for the test period. For the experimental model, a hyperbolic embedding was applied to capture nonlinear price dynamics. Stock prices were normalized relative to a baseline value and modulated by trading volume. The geometric curvature metric was computed using a hyperbolic distance formula and this metric was incorporated as the single exogenous variable in an ARIMAX model (5). For all datasets, the experimental ARIMAX model was implemented with an order of (1, 1, 1). Model performance was assessed via RMSE and compared to the ARIMA control model.

LSTM Modeling (Benchmark)

For contextual comparison, a LSTM network was also constructed. Data were scaled using MinMaxScaler and then segmented using a 60-day sliding window to form training and testing sets (16). The input for this model was the scaled, windowed historical data, and the output was the forecasted

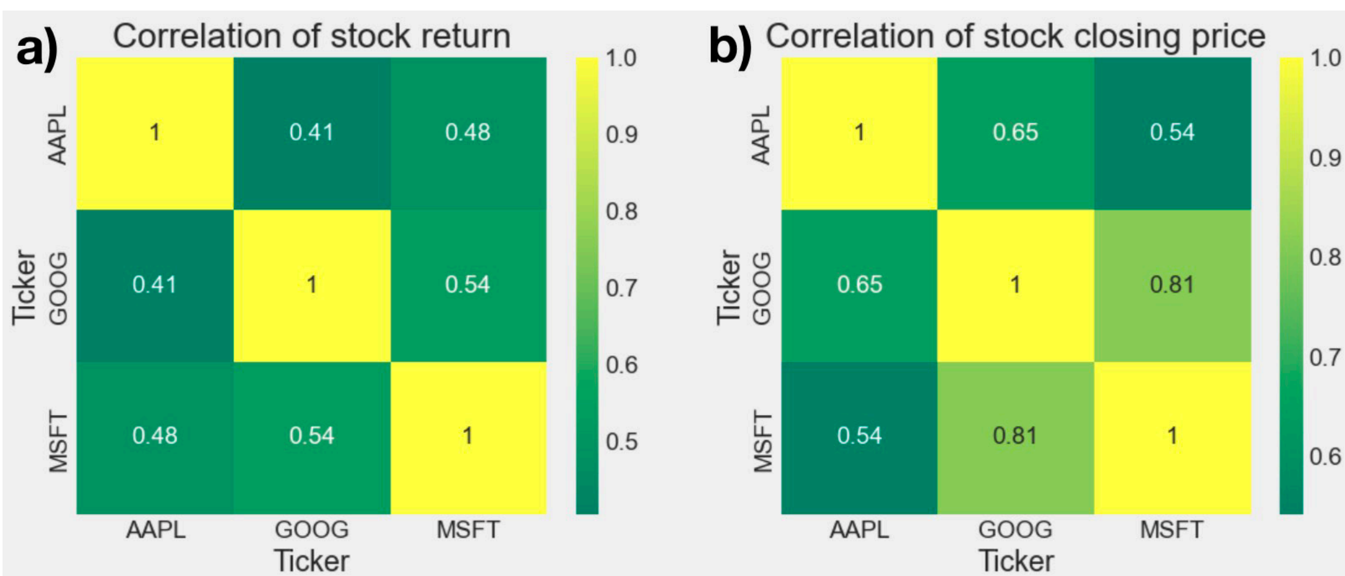

Figure A4. Pearson correlation matrices for stock returns and closing prices. Side-by-side heatmaps of pairwise Pearson correlation coefficients among Apple (AAPL), Google (GOOG), and Microsoft (MSFT) for (a) daily returns and (b) adjusted closing prices. Correlations were calculated in pandas over the 30 August 2023–30 August 2024 time series and annotated within each cell. No additional statistical testing was performed. Two matrices (returns and prices), each based on one contiguous time series per ticker.

stock prices. A sequential LSTM model was built using Keras (17) with two LSTM layers (128 and 64 units) followed by two dense layers. The model was compiled with the Adam optimizer and mean squared error as the loss function. Training was conducted for one epoch with a batch size of one, and predictions were inversely transformed for comparison with actual prices. RMSE was computed to evaluate performance (**Figure 3**).

Received: May 02, 2025
Accepted: January 08, 2026
Published: July 26, 2026

REFERENCES

1. Fama, Eugene F. "Efficient Capital Markets: A Review of Theory and Empirical Work." *The Journal of Finance*, vol. 25, no. 2, 1970, pp. 383–417. doi.org/10.2307/2325486.
2. Box, George E. P., and Gwilym M. Jenkins. "Forecasting." Time Series Analysis: Forecasting and Control. San Francisco, Holden-Day, 1970, pp. 93–145.
3. Cont, Rama. "Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues." *Quantitative Finance*, vol. 1, no. 2, 2001, pp. 223–236. doi.org/10.1080/713665670.
4. Gu, Shihao, *et al*. "Empirical Asset Pricing via Machine Learning." *The Review of Financial Studies*, vol. 33, no. 5, May 2020, pp. 2223–2273. doi.org/10.1093/rfs/hhaa009.
5. López de Prado, Marcos. "Cross-Validation in Finance." Advances in Financial Machine Learning. Hoboken, Wiley, 2018.
6. Nickel, Maximilian and Douwe Kiela. "Poincaré Embeddings for Learning Hierarchical Representations." *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 6338–6347. doi.org/10.5555/3295222.3295381.
7. Wooldridge, Jeffrey M. "Advanced Time Series Topics." Introductory Econometrics: A Modern Approach. 7th ed., Boston, Cengage Learning, 2020, pp. 604–635.
8. Bergstra, James and Yoshua Bengio. "Random Search for Hyper-Parameter Optimization." *Journal of Machine Learning Research*, vol. 13, Feb. 2012, pp. 281–305. doi.org/10.5555/2188385.2188395.
9. Fama, Eugene F. and Kenneth R. French. "Common Risk Factors in the Returns on Stocks and Bonds." *Journal of Financial Economics*, vol. 33, no. 1, Feb. 1993, pp. 3–56. doi.org/10.1016/0304-405X(93)90023-5.
10. Tetlock, Paul C. "Giving Content to Investor Sentiment: The Role of Media in the Stock Market." *The Journal of Finance*, vol. 62, no. 3, June 2007, pp. 1139–1168. doi.org/10.1111/j.1540-6261.2007.01232.x.
11. Andersen, Torben G. and Tim Bollerslev. "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts." *International Economic Review*, vol. 39, no. 4, Nov. 1998, pp. 885–905. doi.org/10.2307/2527343.
12. Vaswani, Ashish, *et al*. "Attention Is All You Need." *Advances in Neural Information Processing Systems*, vol. 30, 2017, pp. 5998–6008. doi.org/10.48550/arXiv.1706.03762.
13. Ranaroussi, Ran. "yfinance Library Documentation." PyPI, https://pypi.org/project/yfinance/. Accessed Aug. 2024.
14. Dickey, David A., and Wayne A. Fuller. "Distribution of the Estimators for Autoregressive Time Series with a Unit Root." *Journal of the American Statistical Association*, vol. 74, no. 366, June 1979, pp. 427–431. doi.org/10.2307/2286348.
15. "Time Series Forecasting with ARIMA." Kaggle, https://www.kaggle.com/code/mechatronixs/time-series-forecasting-arima. Accessed Sept. 2024.
16. Pedregosa, Fabian, *et al*. "Scikit-learn: Machine Learning in Python." *Journal of Machine Learning Research*, vol. 12, 2011, pp. 2825–2830.
17. Chollet, François. "Deep learning for text and sequences." Deep Learning with Python. Shelter Island, Manning Publications, 2017, pp. 178–232.