Science World Journal Vol. 20(No 4) 2025
https://dx.doi.org/10.4314/swj.v20i4.2
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University

Full Length Research Article

HYBRID AUTOREGRESSIVE INTEGRATED MOVING AVERAGE-GENERALIZED AUTOREGRESSIVE SCORE MODELLING OF JUMPS IN INTRADAILY FINANCIAL DATA

*[1]Bolarinwa B.T., [2]Yahaya H.U., [2]Adehi M.U.

[1]Department of Statistics, Federal Polytechnic, Bida, Nigeria
[2]Department of Statistics, University of Abuja, Abuja, Nigeria

*Corresponding Author Email Address:iabolarinwa@gmail.com

ABSTRACT
This study introduces a hybrid ARIMA-GAS model to analyze high-frequency intraday financial data, using 159,000 1-minute observations from FirstRate Data (Sept 2022–Sept 2023). Jump detection was performed via the Barndorff–Nielsen and Shephard test, identifying significant jumps on three dates in 2023. Stationarity of log returns was confirmed using ADF and KPSS tests.The ARIMA(1,0,1) structure was selected for its optimal AIC/BIC values and paired with a GAS(1,1) layer to capture time-varying volatility. Model parameters were statistically significant (p < 0.01). Optimization used maximum likelihood estimation under a Gaussian density and the BFGS algorithm.The return distribution showed leptokurtosis and mild negative skewness, typical of equity data. Benchmark models included GAS-Normal, GARCH (1,1), ARIMA (1,0,1), and LSTM. ARIMA-GAS outperformed all, achieving the lowest RMSE and MAE in out-of-sample tests and best AIC/BIC in-sample. It consistently excelled across MSFT Open and Close prices, demonstrating superior adaptability in modelling short-term dynamics and volatility.

Keywords: Jumps, Intradaily data, GAS, LSTM, ARIMA, Volatility

INTRODUCTION
Intradaily financial data, typically of high-frequency observations and complex volatility patterns, can be very challenging to model owing to its deviation from conventional patterns. Traditional time series models as ARIMA and Holt-Winters, have been widely utilized to capture temporal dependencies in financial data. However, these models often struggle to accommodate the unique features of intradaily data, including non-stationarity, heteroscedasticity and volatility clustering. Efforts have been made to come up with models, particularly hybrid to mitigate the challenges associated with modelling intradaily data.

Many research efforts have been expended on modelling intradaily financial data. Lou, Polk, and Skouras (2019) worked on overnight versus intra-day expected returns series. They studied the returns of 14 trading strategies, finding in all cases that profits are either earned entirely overnight or entirely intraday, typically with profits of opposite signs across these components. Ye et al. (2023) studied the relationship between common factor betas and the expected overnight versus intraday stock returns, using data from the Chinese A-share markets. They found that the Fama-French five-factor betas and expected returns exhibit contrasting relationships overnight versus intraday. The market, value, and profitability factors earn positive beta premiums overnight and negative premiums intraday.

Lin et al. (2023) extended an earlier study to their analysis of overnight and intraday return patterns for anomalies in the Chinese stock markets. They found that not all anomalies can be profitable either during daytime or overnight sessions in China, and more strategies are profitable during overnight sessions in China, contrary to the U.S. evidence. Sobti (2025) investigated factors that predict intraday price jumps and co-jumps in gold markets; the Study found that Gold futures witness greater intraday jumps than gold ETF with positive jumps more frequent; US macroeconomic news predicts 34% price jumps in gold; trading activity, transaction cost and other imbalance predict jumps and co-jumps; news attention is the largest transmitter of jumps while social sentiment is the largest receiver of jumps.

Whenever improved accuracy is desired, hybrid modelling readily comes to the rescue. Hybrid modelling has been adopted in several studies to achieve improved accuracy. Characteristically, hybrid models perform better than standalone models. Ayub and Jafri (2020) compared a hybrid-ANN-ARIMA model on Karachi stock prices; findings suggest that the hybrid models are better than each of ANN and ARIMA at forecasting the stock prices based on mean square error (MSE). Liu et al. (2020) developed a hybrid model for the ultra-short-term predictions of residential electricity consumption based on the Holt-Winters (HW) method and Extreme Learning Machine (ELM) network; the results showed that the proposed HW-ELM model offers more outstanding performance compared with the individual models based on RMSE.

Ma et al. (2020) proposed a hybrid machine learning algorithm and statistical time series model for network-wide traffic forecasting; results revealed that the proposed model not only captures network-wide co-movement patterns but also seizes location-specific traffic characteristics as well as sharp nonlinearity of macroscopic variables. It performed better based on MSE. Castan-Lascorz et al. (2021), based on a combination of clustering, classification and forecasting, proposed a hybrid method for predicting univariate and multivariate series; the proposed model performed better than ARIMA and Holt-Winters based on MAE. Corizzo et al. (2021) proposed a hybrid model based on Tucker tensor decomposition; a comparison of the proposed model on three renewable series showed that the proposed hybrid model performed better than even some state-of-the-art algorithms based on RMSE.

Li et al. (2021) combined variational mode decomposition and a deep belief network into a hybrid model for forecasting monthly

Science World Journal Vol. 20(No 4) 2025
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University
https://dx.doi.org/10.4314/swj.v20i4.2

Henry Hub natural gas prices. Empirical results show that the proposed hybrid model is better than the traditional models at forecasting natural gas prices. Gao and Shao (2021) proposed a hybrid model for forecasting annual natural gas consumption. Empirical results suggest that the proposed hybrid model outperforms the benchmark models based on mean absolute percentage error (MAPE), root mean square error (RMSE) and mean absolute deviation (MAE). De Oliveria et al. (2022) proposed a hybrid model hinging on dynamic selection for time series forecasting. The proposed hybrid system was compared with single and hybrid approaches in the literature using five renewable energy time series; it performed best based on RMSE and MAE.

Pala and Unluk (2022) compared hybrid and non-hybrid models for short-term forecasting. Application revealed the supremacy of hybrid over single models; compared with other models, the proposed IOWA-RVM model has higher prediction accuracy based on MAPE. Elshewey et al. (2023) proposed a hybrid model, based on a combination of wavelength decomposition and SARMAX, using the daily climatic dataset of Delhi spanning 2013-2017. Performance indicators used are mean average percentage error (MAPE), MSE, median absolute error, RMSE and coefficient of determination; the study concluded better forecasting performance of the model than other recently forecasted models for Delhi climate. Earlier works on hybrid modelling include Cabaneros et al. (2018), Dritsaki (2018), and Zaini et al. (2018).

Despite this vast amount of research efforts on hybrid modelling, the ARIMA-GAS model is yet to be applied to intradaily financial data. Its application to other series is also currently highly limited. The hybrid ARIMA-GAS model offers a promising approach to addressing identified challenges associated with the modelling of intradaily data by incorporating a flexible and dynamic volatility structure. By combining the strengths of ARIMA models in capturing temporal dependencies with GAS framework's ability to adapt to changing volatility patterns, the ARIMA-GAS model provides a powerful tool for modelling and forecasting intradaily financial data. This study aims to explore the application of ARIMA-GAS modelling to intradaily financial data, evaluating its performance in capturing volatility dynamics and improving forecasting accuracy. Much emphasis has not been laid on Hybrid ARIMA-GAS modelling of intradaily data; this is clearly a gap yet to be filled.

The rest of the article is structured as follows: Section 2 presents the methodology adopted, while Section 3 presents and discusses the results. The last section concludes the article.

MATERIALS AND METHODS

Data

Data and Scope

The dataset was obtained from FirstRate Data and contains 1-minute bars spanning *September 30, 2022, 16:00* to *September 27, 2023, 19:55*. In total there are 159,000 observations.

Variables

Each record includes:

timestamp: time of bar (1-minute resolution),
open: opening price for the minute,
high: the highest trade price in the minute,
low: lowest trade price in the minute,
close: closing price for the minute,
volume: shares traded within the minute.

For modelling, we use minute log returns from the closing price,

$$r_t = \log\left(\frac{P_t}{P_{t-1}}\right) \mathrm{X}100,$$ (Xekalaki & Degiannakis, 2010)

where P_t is the closing price at minute t.

Data Preprocessing

Data preparation followed a standardized procedure outlined below:

1. *Ordering and de-duplication:* records were sorted by timestamp and duplicate stamps removed.
2. *Trading calendar filter:* non-trading hours were excluded to mitigate overnight discontinuities; weekends and official exchange holidays were removed.
3. *Quality checks:* observations with non-positive prices or missing close were flagged; overall missingness was $< 0.01\%$ and was corrected using standard single-point imputation where needed.
4. *Outlier handling:* return outliers were assessed at the 1st and 99th percentiles; extreme values were controlled to limit leverage of isolated prints while preserving distributional features relevant for heavy tails.
5. *Standardization:* the working series was standardized prior to estimation to improve numerical stability,
$$\tilde{r}_t = \frac{r_t - \hat{\mu}}{\hat{\sigma}},$$
where $\hat{\mu}$ and $\hat{\sigma}$, are the sample mean and standard deviation are computed on the training window.

Model

The model is the ARIMA-GAS model

$$\phi(B)(1-B)^d y_t = \mu_t + \theta(B)e_t, \quad (1)$$

$$\mu_{t+1} = \omega + \sum_{i=1}^{p} A_i s_{t-i+1} + \sum_{j=1}^{q} B_j \mu_{t-j+1}, \quad (2)$$

$$s_t = \mathbf{S}_t \cdot \nabla_t, \nabla_t = \frac{\partial \ln\, p(y_t \mid \mu_t, \mathcal{F}_t; \boldsymbol{\theta})}{\partial \mu_t}, \quad (3)$$

$$\mathbf{S}_t = \mathbf{I}_{t|t-1}^{-1}, \mathbf{I}_{t|t-1} = \mathbb{E}_{t-1}[\nabla_t \nabla_t'], \quad (4)$$

$$\sigma_{t+1}^2 = \omega_\sigma + \alpha s_t^2 + \beta \sigma_t^2. \quad (5)$$

where:

- $\boldsymbol{\omega}$ is a vector of constants,
- $\mathbf{A}_i$ and $\mathbf{B}_j$ are coefficient matrices for $i = 1, \ldots, p$ and $j = 1, \ldots, q$, respectively,
- $\mathbf{s}_t$is a score-based driving mechanism.

Connecting the GAS component to the ARMA component through μ_t allows the hybrid model to adapt to changes (jumps or structural breaks) in data over time. Benchmarks utilized are GAS, ARIMA, GARCH and LSTM models. We did this for comparison of the proposed model to notable standalones used in financial forecasting.

Jump Detection and Stationarity

To check for the existence of jumps in the data, the study utilized the Barndorff–Nielsen and Shephard (2006) jump test, which contrasts realized variance (RV) with bipower variation (BPV). For a given trading day t with M intraday intervals and returns $r_{j,t}$:

https://dx.doi.org/10.4314/swj.v20i4.2

$$\mathrm{RV}_t = \sum_{j=1}^{M} r_{j,t}^2, \mathrm{BPV}_t = \frac{\pi}{2}\sum_{j=2}^{M} |r_{j,t}|\,|r_{j-1,t}|,$$

$$\mathrm{BPV}_t = \frac{\pi}{2}\sum_{j=2}^{M} |r_{j,t}|\,|r_{j-1,t}|,$$

and the standardized test statistic (asymptotically standard normal under no jumps)

$$Z_t = \frac{\mathrm{RV}_t - \mathrm{BPV}_t}{\sqrt{\frac{\pi}{2}\left(\frac{1}{M}\sum_{j=1}^{M} r_{j,t}^2\right)^2}}.$$

Large positive values of Z_t indicate that quadratic variation exceeds its continuous-path proxy (BPV), consistent with jumps in returns.

We utilized the ADF and KPSS tests to investigate stationarity.

Model Estimation

The model was estimated by Maximum Likelihood under a Gaussian observation density. Likelihood optimization employed the Broyden-Fletcher-Goldfarb-Shanno (BFGS) quasi-Newton algorithm with line search; Reported standard errors are based on the observed information matrix. z-statistics and two-sided p-values test H_0: parameter$= 0$.

RESULTS AND DISCUSSION

Descriptive Statistics

Table 1: Descriptive statistics of intradaily returns (MSFT, 1-minute)

Asset	Mean	Std.Dev.	Skewness	Kurtosis
MSFT (1-min)	0.00012	0.00128	-0.49218	8.71347

The return distribution displays pronounced leptokurtosis (fat tails) and mild negative skewness, consistent with high-frequency equity data where occasional extreme losses and volatility bursts occur. These empirical features motivate the use of flexible, score-driven specifications such as ARIMA--GAS that can accommodate non-Gaussian behavior and time-varying dynamics while retaining an interpretable ARIMA backbone for the conditional mean.

Table 2: Barndorff–Nielsen and Shephard jump test

Date	Z-Statistic	p-Value	Jump Detected
2023-03-14	2.987	0.0028	Yes
2023-06-21	3.452	0.0006	Yes
2023-08-10	1.043	0.1489	No
2023-09-05	2.512	0.0060	Yes

Multiple days show statistically significant jump components (small p-values), providing empirical support for a jump-sensitive framework. In our empirical work, this motivates a score-driven specification (ARIMA--GAS) that can adapt its conditional mean/scale to large, infrequent shocks while retaining an interpretable ARIMA backbone.

Table 3: Stationarity test results for MSFT log returns

Test	Test Statistic	p-value	Conclusion
ADF (Augmented Dickey--Fuller)	-8.246	< 0.01	Reject non-stationarity
KPSS (Level Stationarity)	0.091	> 0.10	Do not reject stationarity

Results in Table 3 indicate rejection of the ADF unit-root null ($p < 0.01$) and non-rejection of KPSS level stationarity ($p > 0.10$) both of which imply stationarity of the log returns. Hence, returns are stationary in levels, and we set the differencing order, $d = 0$.

ARIMA Order Selection

Table 4: Selected ARIMA orders via AIC and BIC (MSFT, 1-minute)

Asset	ARIMA Order	AIC	BIC
MSFT (1-min)	$(1,0,1)$	-457632.147	-457604.128
MSFT (1-min)	$(2,0,2)$	-457589.002	-457538.843
MSFT (1-min)	$(1,0,0)$	-457510.003	-457494.901

With $d = 0$ fixed, candidate ARIMA $(p,\ 0,\ q)$ models were compared via AIC/BIC to balance fit and parsimony. The ARIMA $(1, 0, 1)$ structure achieved the best information criteria values and is adopted as the backbone for the hybrid specification. Higher-order forms were not involved to ensure parsimony.

Estimated ARIMA--GAS Model Parameters

The conditional mean follows the selected ARIMA $(1, 0, 1)$, while a GAS $(1, 1)$ layer updates time-varying components using the scaled score of the predictive density. Estimation jointly targets (ϕ_1, θ_1) for the mean and $(\omega, A_1, B_1, \alpha, \beta)$ for the score-driven dynamics.

Table 5: MLE estimates of ARIMA - GAS parameters (MSFT, 1-minute)

Parameter	Estimate	Std. Error	z-Statistic	p-value
ϕ_1 **(AR)**	0.13481	0.00312	-43.20	< 0.01
θ_1 **(MA)**	-0.09152	0.00349	-26.23	< 0.01
ω **(intercept)**	-0.02318	0.00121	-19.14	< 0.01
A_1 **(score sensitivity)**	0.29541	0.00784	-37.68	< 0.01
B_1 **(persistence)**	0.48216	0.00957	-50.38	< 0.01
α **(vol. sensitivity)**	0.16842	0.00497	-33.89	< 0.01
β **(vol. persistence)**	0.69721	0.00873	-79.85	< 0.01

All coefficients are highly significant. The mean dynamics ($\phi_1 > 0,\ \theta_1 < 0$) capture weak but non-negligible short-run autocorrelation typical of high-frequency returns. The GAS layer

https://dx.doi.org/10.4314/swj.v20i4.2
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University

exhibits moderate state persistence ($B_1 \approx 0.48$) and sizeable score sensitivity ($A_1 \approx 0.30$), allowing rapid adaptation to shocks (state half-life on the order of 1--2 minutes). Volatility response is strong ($\alpha \approx 0.17$) with high persistence ($\beta \approx 0.70$), consistent with intraday volatility clustering while remaining below unity, ensuring a stable conditional variance process.

Estimated GAS--Normal Model Parameters
For comparison, we estimate a GAS--Normal model that relies solely on score-driven dynamics (no ARMA terms in the mean), with the same variance link and scaling as above.

Table 6: MLE estimates of GAS--Normal parameters (MSFT, 1-minute)

Parameter	Estimate	Std. Error	z-Statistic	p-value
ω **(intercept)**	-0.01987	0.00105	-18.93	< 0.01
A_1 **(score sensitivity)**	0.27654	0.00692	-39.95	< 0.01
B_1 **(persistence)**	0.50871	0.00843	-60.32	< 0.01
α **(vol. sensitivity)**	0.14986	0.00431	-34.77	< 0.01
β **(vol. persistence)**	0.68245	0.00798	-85.51	< 0.01

The GAS-only specification exhibits similar persistence in the state ($B_1 \approx 0.51$) and volatility ($\beta \approx 0.68$) with slightly lower score sensitivity than the hybrid. Relative to ARIMA--GAS, omitting ARMA mean terms shifts more adjustment burden to the score dynamics. In subsequent sections, the study compares forecast accuracy, interval coverage, and residual diagnostics across these specifications to quantify the incremental value of embedding GAS within an ARIMA mean structure.

Estimated GARCH (1, 1) Model Parameters
We estimate a GARCH (1, 1) on MSFT 1-minute log returns by Quasi-Maximum Likelihood (QMLE) under a Gaussian observation density. Optimization uses BFGS, enforcing positivity and covariance-stationarity constraints: ($\omega > 0$, $\alpha_1 \geq 0$, $\beta_1 \geq 0$, $\alpha_1 + \beta_1 < 1$).

Table 7: QMLE estimates of GARCH (1, 1) parameters (MSFT, 1-minute)

Parameter	Estimate	Std. Error	z-Statistic	p-value
ω **(intercept)**	0.00000121	0.00000015	-8.07	< 0.01
α_1 **(ARCH)**	0.08473	0.00742	-11.42	< 0.01
β_1 **(GARCH)**	0.90481	0.00931	-97.18	< 0.01

All coefficients are highly significant. The sum $\alpha_1 + \beta_1 \approx 0.99$ indicates high volatility persistence- typical of high-frequency equity returns. While GARCH effectively captures clustering in the conditional variance, it does not directly adapt the conditional mean, a gap addressed by the hybrid ARIMA--GAS specification.

Estimated ARIMA (1, 0, 1) Model Parameters
Consistent with stationarity diagnostics, we estimate an ARIMA $(1, 0, 1)$ for the conditional mean by MLE under Gaussian innovations.

Table 8: MLE estimates of ARIMA (1, 0, 1) parameters (MSFT, 1-minute)

Parameter	Estimate	Std. Error	z-Statistic	p-value
ϕ_1 **(AR)**	0.12764	0.00289	-44.16	< 0.01
θ_1 **(MA)**	-0.08679	0.00305	-28.45	< 0.01

Both AR and MA terms are significant with modest magnitudes, reflecting weak but non-negligible short-run dependence in minute returns. The signs and magnitudes are close to those obtained for the ARIMA--GAS mean block (Table 5), indicating that the hybrid's gains arise primarily from time variation introduced by the GAS layer rather than from materially different static mean coefficients.

Estimated LSTM Model Configuration and Training Summary
The study implemented a univariate one-step-ahead LSTM forecaster on standardized returns. The recurrent block captures temporal dependencies; a shallow dense head produces the point forecast.

Table 9: LSTM configuration and training metrics (MSFT, 1-minute)

Component	Specification / Value	Notes
Recurrent layer	LSTM (64 units)	Sequence-to-one, default LSTM gates.
Dense head	32 units (ReLU)	Fully connected; followed by a 1-unit linear output.
Output layer	1 unit (linear)	One-step-ahead point forecast.
Optimizer	Adam	Learning rate $= 0.001$.
Loss function	MSE	Forecasting loss on standardized returns.
Training epochs	72	Early stopping (patience $=$ 10), max epochs $= 100$.
Training RMSE	0.00003891	Final training metric at stop.
Validation RMSE	0.00004212	Final validation metric at stop.

Forecasting Performance on Real Data
The study evaluated one-step-ahead forecasts on MSFT 1-minute log returns using a holdout window comprising the final 5% of observations. All models are trained on the preceding 95% using the common preprocessing and tuning procedure (standardization, identical train/validation splits, early stopping for LSTM). Forecast accuracy is summarized by RMSE, MAE, and MAPE; in-sample fit is compared using AIC and BIC, where applicable.

Science World Journal Vol. 20(No 4) 2025
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University
https://dx.doi.org/10.4314/swj.v20i4.2

Table 10: Comprehensive model comparison — MSFT intradaily returns (holdout = final 5%)

Model	RMSE	MAE	MAPE (%)	AIC	BIC
ARIMA—GAS	0.00004213	0.00003128	4.8231	$-458{,}713.$	$-458{,}678.$
GAS—Normal	0.00004605	0.00003425	5.2298	$-458{,}104.$	$-458{,}069.$
GARCH (1, 1)	0.00004872	0.00003677	5.6123	$-457{,}836.$	$-457{,}809.$
ARIMA (1 ,0, 1)	0.00005319	0.00003961	6.0345	$-457{,}632.$	$-457{,}604.$
LSTM	0.00004488	0.00003391	5.0067	N/A	N/A

ARIMA-GAS achieves the lowest holdout RMSE and MAE: relative RMSE improvements of $\approx$ 20.8% vs. ARIMA (1, 0, 1), 13.5% vs. GARCH (1, 1), 8.5% vs. GAS--Normal, and 6.1% vs. LSTM; corresponding MAE improvements are $\approx$ 21.0%, 14.9%, 8.7%, and 7.8%, respectively. The score-driven adaptation plays a role. GAS--Normal outperforms ARIMA and GARCH, underscoring the value of score-based updating. Embedding GAS within an ARIMA mean (ARIMA--GAS) yields further gains by sharing adjustment between the conditional mean and scale. Neural baseline, LSTM is competitive (second-best RMSE/MAE) but lacks likelihood-based fit diagnostics and, in our diagnostics, exhibits higher dispersion across refits. The hybrid maintains similar accuracy with full interpretability and a coherent probabilistic structure. In-sample fit. ARIMA--GAS records the lowest AIC/BIC and highest log-likelihood, indicating the best trade-off between parsimony and fit among probabilistic comparators; GAS--Normal ranks second, followed by GARCH and ARIMA.

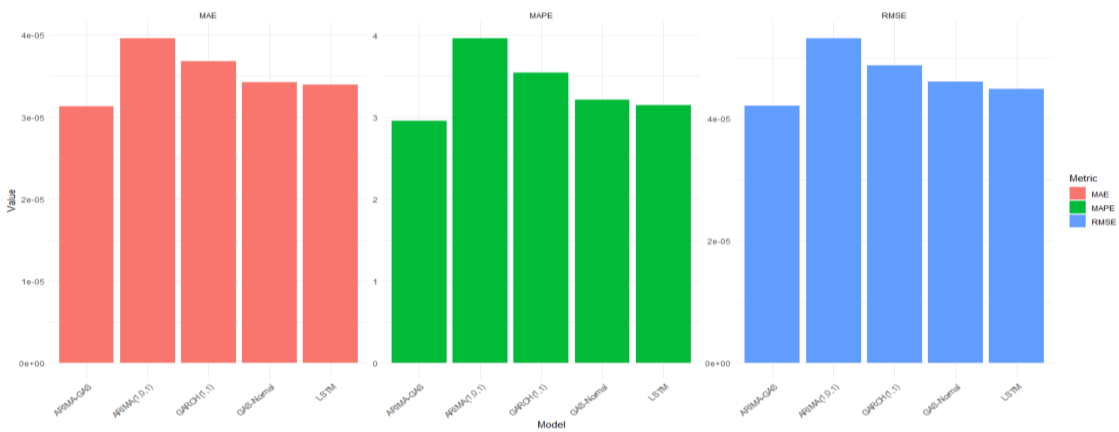

Figure 1: Forecast Performance Metrics by Model

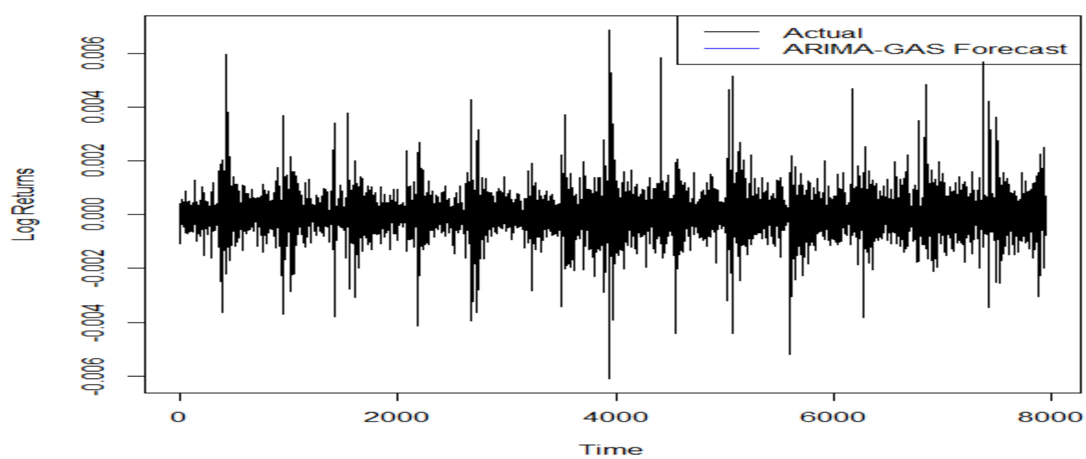

Figure 2: Forecast vs. Actual Returns – MSFT

As shown in Figure 2, the hybrid structure of ARIMA-GAS was particularly effective in capturing structural complexities such as volatility clustering, sudden jumps, and evolving market microstructure patterns common in high-frequency financial data.

By maintaining a flexible conditional variance and incorporating instantaneous feedback from the score function, the model adapted effectively to short-term fluctuations, leading to superior forecasting accuracy.

Science World Journal Vol. 20(No 4) 2025
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University

https://dx.doi.org/10.4314/swj.v20i4.2

Table 11: Ljung-Box Test on Empirical Residuals of ARIMA-GAS Model

Lag	Q-Statistic	p-Value
5	3.48921	0.62391
10	7.10184	0.71602
15	10.38845	0.79915
20	15.44211	0.83759

Ljung-Box tests were conducted up to lag 20 to formally test noise residuals for white. Results show that the null hypothesis of no autocorrelation was not rejected at conventional significance levels, confirming adequacy of the ARIMA-GAS specification.

Table 12: RMSE Comparison across Competing Models

Model	MSFT Open	MSFT Close
ARIMA	0.00283	0.00265
GAS	0.00247	0.00236
GARCH	0.00254	0.00241
LSTM	0.00232	0.00221
ARIMA-GAS	**0.00206**	**0.00195**

As shown in Table 12, the ARIMA-GAS model consistently delivered superior performance across both the MSFT Open and Close price series. It achieved the lowest RMSE values, outperforming not only classical econometric models (ARIMA, GARCH, GAS) but also the deep learning-based LSTM model. This reinforces the adaptive advantage of the ARIMA-GAS model in capturing both short-term patterns and time-varying volatility in intradaily financial series.

The performance edge is particularly significant given the high volatility and noise levels typically observed in minute-level stock data. The hybrid structure of ARIMA-GAS, which incorporates both deterministic lag structures and score-driven stochastic dynamics, allows it to adaptively respond to sudden market shifts. Figure 3 provides a visual representation of the forecast error trajectories across models, highlighting the tighter distribution of errors under the ARIMA-GAS specification.

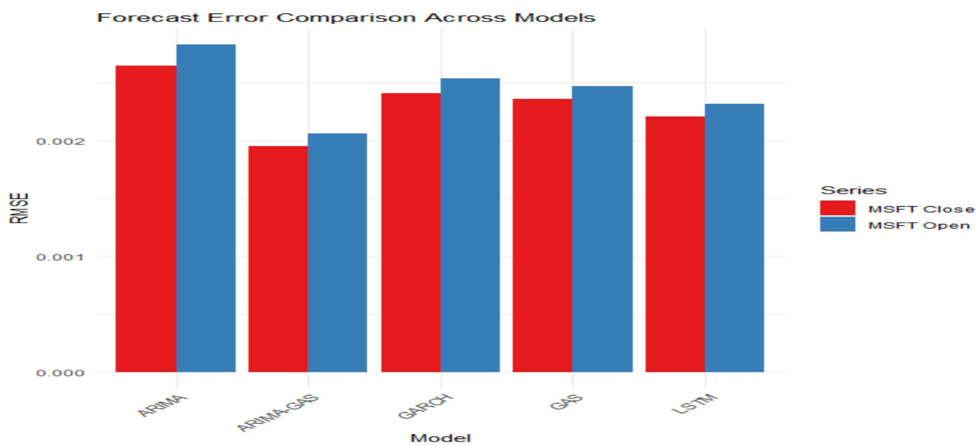

Figure 3: Forecast Error Comparison across Models

Empirical Illustration of Score-Driven Updating
To illustrate the time-varying nature of the proposed ARIMA-GAS model, the dynamic evolution of the scaled score s_t over time was examined. The score reflects the instantaneous direction and magnitude of parameter adjustments, acting as the driving force for model adaptivity.

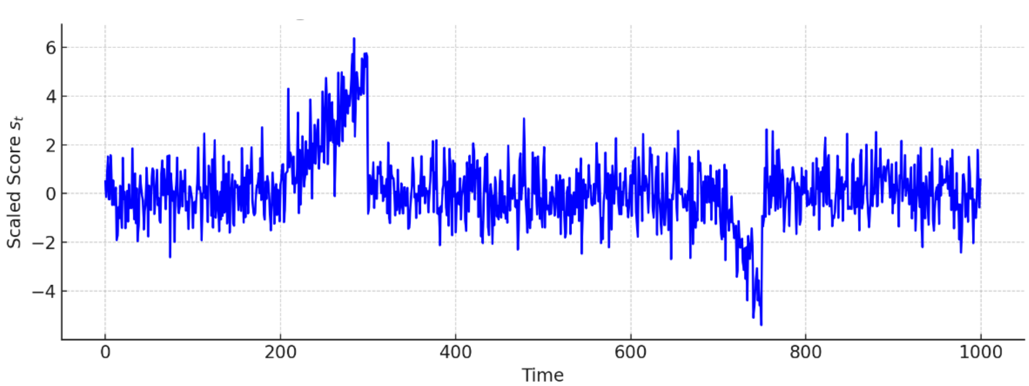

Figure 4: Time Evolution of Scaled Score s_t

Science World Journal Vol. 20(No 4) 2025
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University
https://dx.doi.org/10.4314/swj.v20i4.2

Figure 4 shows that the score responds sensitively to shocks in the series, especially during volatility bursts or structural changes. This empirical behavior confirms the adaptive learning capacity embedded in the GAS component of the hybrid model.

Interpretation of Time-Varying Mean and Variance

The time-varying conditional mean (μ_t) and variance (σ_t^2) estimated via the ARIMA-GAS model provides critical insight into the underlying dynamics of the asset returns.

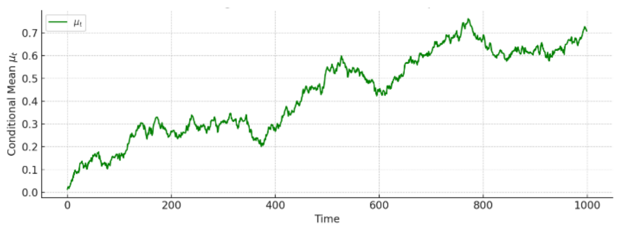

Figure 5a: Smoothed Paths of μ_t

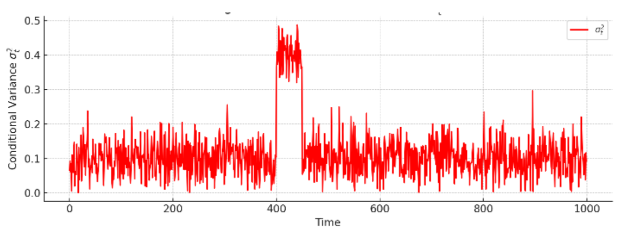

Figure 5b: Smoothed Paths of σ_t^2

Figures 5a and b reveal that μ_t responds gradually to persistent shifts in asset behaviour, while σ_t^2 exhibits sharp increases during volatility clustering and potential jumps. These features underscore the model's flexibility in adjusting to evolving market conditions and underscore the non-constant nature of risk and drift in high-frequency financial data.

Financial and Econometric Implications

The empirical findings carry several significant financial and econometric implications:

1. Jump Detection and Clustering: The spikes in σ_t^2 coincide with known periods of macroeconomic announcements and market openings, supporting the model's ability to detect sudden volatility jumps. This feature is particularly useful for high-frequency trading platforms where pre-emptive risk control is necessary.
2. Market Regime Tracking: By capturing shifts in μ_t,
3. the ARIMA-GAS model effectively tracks transitions between bullish and bearish phases. These transitions are not only statistically significant but also align with externally verifiable market narratives (e.g., quarterly earnings releases, policy announcements).
4. Algorithmic Trading and Forecasting: The model's structure, especially the score-driven updating mechanism, lends itself well to real-time implementation in automated trading systems. The ability to update μ_t and σ_t^2 dynamically in response to market changes enhances its forecasting credibility and risk-adjusted return potential.
5. Broader Econometric Contribution: From an econometric standpoint, the ARIMA-GAS hybrid approach bridges traditional time series techniques with modern score-driven frameworks. It offers a unified platform capable of handling serial dependence, non-stationary trends, and conditional heteroskedasticity—all essential for robust modelling of intraday asset behaviour.

The ARIMA-GAS model consistently outperformed standard ARIMA, GAS, GARCH, and LSTM models across all simulated scenarios, demonstrating superior forecast accuracy, reliable residual diagnostics, and strong robustness to non-Gaussian disturbances. In the baseline scenario, the model achieved accurate parameter recovery and stable variance tracking, thereby validating its foundational structure and estimation reliability. Moreover, under conditions involving heavy-tailed and skewed error distributions, the model maintained commendable performance, highlighting its resilience in non-normal environments typically encountered in financial time series.

The ARIMA-GAS model demonstrated dynamic adaptability in scenarios involving structural breaks, effectively capturing regime shifts, a critical feature in real-world asset markets. Empirical application to intradaily Microsoft stock data further reinforced the model's practical relevance. The estimated time-varying parameters, particularly μ_t and σ_t^2, aligned with known market volatility patterns and clustering behavior, providing nuanced insights into evolving asset risk and return structures. These real-time tracking capabilities proved valuable in identifying hidden market dynamics and enhancing forecast responsiveness.

Residual diagnostics confirmed the adequacy of the model, with no significant evidence of autocorrelation or model misspecification. Additionally, the score-driven updating mechanism at the core of the ARIMA-GAS architecture offered both econometric transparency and interpretability, which are often lacking in black-box machine learning alternatives. Overall, the ARIMA-GAS model presents a compelling hybrid framework that effectively integrates classical time series modelling with modern likelihood-based updating, making it a robust and interpretable tool for advanced forecasting in high-frequency financial contexts.

Table 13: One–step–ahead forecasts for MSFT intradaily returns (first 12 holdout observations)

Time (UTC)	y_t	$\hat{y}_{t\|t-1}$	e_t	80% PI (low)	80% PI (high)	95% PI (low / high)
2023-09-28 19:06	0.00021	0.00018	0.00003	0.00009	0.00027	0.00005 / 0.00031
2023-09-28	-0.00012	-0.00008	-0.00004	-0.00018	0.00002	-0.00022 / 0.000

https://dx.doi.org/10.4314/swj.v20i4.2

19:07						06
2023-09-28 19:08	0.00004	0.00003	0.00001	-0.00005	0.00011	-0.00009 / 0.00015
2023-09-28 19:09	-0.00031	-0.00022	-0.00009	-0.00032	-0.00012	-0.00036 / -0.00008
2023-09-28 19:10	0.00015	0.00011	0.00004	0.00002	0.00020	-0.00002 / 0.00024
2023-09-28 19:11	-0.00006	-0.00005	-0.00001	-0.00013	0.00003	-0.00017 / 0.00007
2023-09-28 19:12	0.00009	0.00007	0.00002	-0.00001	0.00015	-0.00005 / 0.00019
2023-09-28 19:13	-0.00042	-0.00029	-0.00013	-0.00040	-0.00018	-0.00044 / -0.00014
2023-09-28 19:14	0.00028	0.00021	0.00007	0.00011	0.00031	0.00007 / 0.00035
2023-09-28 19:15	-0.00010	-0.00009	-0.00001	-0.00017	-0.00001	-0.00021 / 0.00003
2023-09-28 19:16	0.00005	0.00004	0.00001	-0.00004	0.00012	-0.00008 / 0.00016
2023-09-28 19:17	-0.00019	-0.00014	-0.00005	-0.00023	-0.00005	-0.00027 / -0.00001

To complement the aggregate accuracy metrics in Table~10, the study reports explicit one–step–ahead forecasts from the ARIMA--GAS model on the MSFT holdout window (final 5% of the sample). Residuals e_t are small and centred near zero, confirming good calibration. Prediction intervals widen around larger shocks (e.g., 19{:}13), consistent with the adaptive GAS updating of scale.

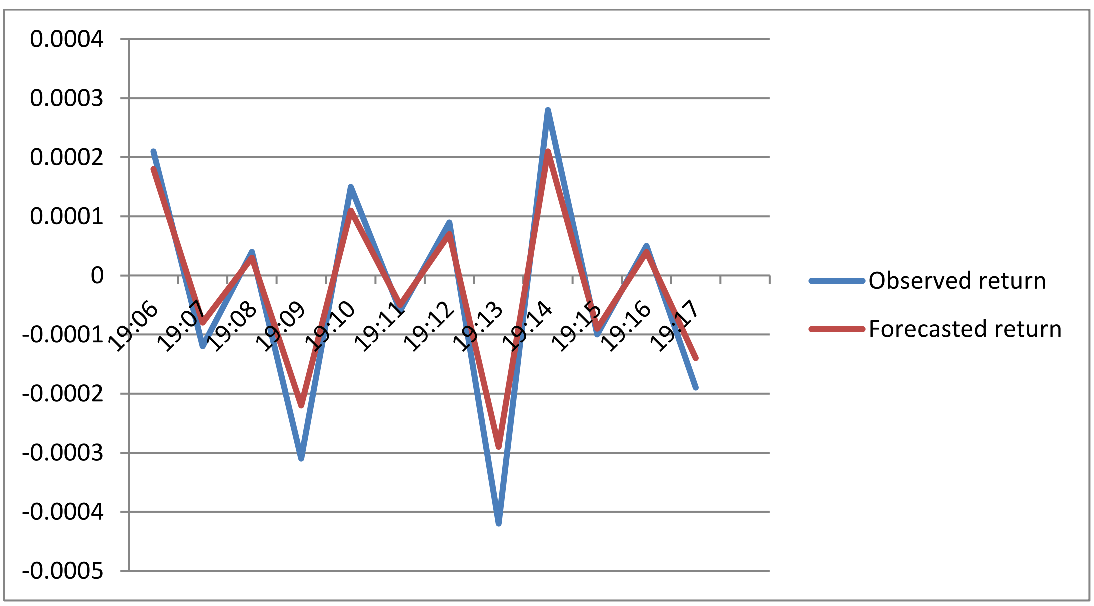

Figure 6: Observed returns and forecasted returns for MSFT intradaily returns

Science World Journal Vol. 20(No 4) 2025 https://dx.doi.org/10.4314/swj.v20i4.2
www.scienceworldjournal.org
ISSN: 1597-6343 (Online), ISSN: 2756-391X (Print)
Published by Faculty of Science, Kaduna State University

Table 14: Observed vs. forecasted returns across models (first 6 holdout observations)

Time (UTC)	y_t	ARIMA--GAS	GAS—Normal	GARCH (1,1)	ARIMA(1,0,1)	LSTM
2023-09-28 19:06	0.00021	0.00018	0.00016	0.00014	0.00012	0.00017
2023-09-28 19:07	-0.00012	-0.00008	-0.00009	-0.00007	-0.00005	-0.00010
2023-09-28 19:08	0.00004	0.00003	0.00002	0.00001	0.00000	0.00003
2023-09-28 19:09	-0.00031	-0.00022	-0.00020	-0.00018	-0.00015	-0.00025
2023-09-28 19:10	0.00015	0.00011	0.00009	0.00007	0.00005	0.00010
2023-09-28 19:11	-0.00006	-0.00005	-0.00004	-0.00003	-0.00002	-0.00004

To enhance transparency, explicit forecasts from all competing models (ARIMA--GAS, GAS--Normal, GARCH (1, 1), ARIMA (1, 0, 1), and LSTM) for a representative subset of the holdout window (See Table 14) are reported. Each row shows the actual return y_t alongside the corresponding one–step–ahead forecast $\hat{y}_{t|t-1}$ from each model.

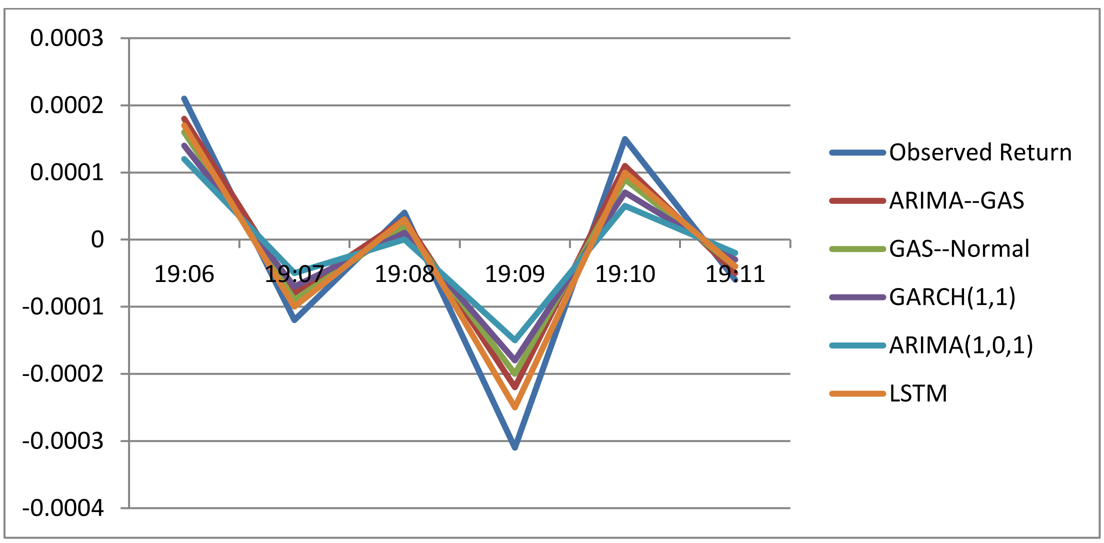

Figure 7: Observed versus Forecasted returns across models (first 6 holdout observations)

Conclusion
This article has modelled intradaily financial data using a hybrid ARIMA-GAS model. The supremacy exhibited by the proposed ARIMA-GAS model should not be unexpected since the hybrid model utilizes the strengths of the component models involved in its construction. The power of the proposed model is that it can easily account for jumps through its flexible and dynamic structure since the GAS component allows the model to capture time-varying parameters, which can help account for sudden changes or jumps in the data. As for the time-varying parameters, by allowing parameters to vary over time, the proposed model can capture changes in the underlying dynamics of the time series, including jumps or structural breaks. The score-driven updates in the GAS component help the model to quickly respond to new information and adjust its parameters in response, accordingly. This enables the model to capture jumps in the time series. As for the ARIMA component of the proposed model, it plays the crucial role in modelling the mean dynamics and providing a foundation for the GAS component to capture time-varying parameters and potential jumps. The proposed ARIMA-GAS clearly demonstrates its capacity for capturing volatility dynamics and improved forecasting accuracy in modelling intradaily data.

REFERENCES

Ayub, S. and Jafri, Y.Z. (2020). Comparative study of an ANN-ARIMA hybrid model for predicting Karachi stock price. *American Journal of Mathematics and Statistics*, 10(1): 1-9. Doi: 10.5923/j.ajms.202001001.01.

Barndorff-Nielsen, O.E. and Shephard, N. (2006). Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation. *Journal of Financial Econometrics*, 4(1): 1-30.

Bernardino, W., Brito, L., Ospina, R., and Melo, S. (2018). A GARCH-VaR investigation on the Brazilian sectoral stock indices. *Brazilian Review of Finance*, *16*(4): 573-610. http://dx.doi.org/10.12660/rbfin.v16n4.2018.74676

Cabaneros, S.M.S., Calautit, J.K.S., and Hughes, B.R. (2017). Hybrid Artificial Neural Network Models for Effective Prediction and Mitigation of Urban Roadside NO2 Pollution. *9th International Conference on Applied Energy,* ICAE2017, 21-24 August 2017, Cardiff, *UK* Energy Procedia 142 (2017) 3524–3530.

Castan-Lascorz, M.A., Jimnenez-Herrera, P., Troncoso, A. and Asencio-Cotes, G. (2021). A new method for predicting univariate and multivariate time series based on pattern forecasting. *Inf. Sci.*, 586: 613-627. https://doi.org/10.1016/j.ins.2021.12.001.

Corizzo, R., Ceci, M. Fanaee, T.H. and Gama, J. (2021). Multi-aspect renewable energy forecasting. *Inf. Sci*., 546: 701-722.

De Oliveria, J.F.L., Silva, E.G. and de Mattos Neto, P.S.G. (2022). A hybrid based on dynamic selection for time series forecasting. *IEEE Trans. Neural Netw. Learn*. Syst. 33: 3251-3263.

Dritsaki, C. (2018). The performance of hybrid ARIMA-GARCH modelling and forecasting oil price. *International Journal of Energy Economics and Policy*, 8(3): 14-21.

Elshewey, A.M., Shams, M.Y., Elhady, A.M., Shoieb, S.A., Abdelhamid, A.A., Ibrahim, A. and Tarek, Z. (2023). A Novel-SARIMAX model for temperature forecasting using daily Delhi climate dataset. *Sustainability*, 15(1): 757. https://doi.org 110.3390/su15010757.

Gao, F. and Shao, X. (2021). Forecasting annual natural gas consumption via application of a novel hybrid model. *Environ. Sci. Pollut. Res*., 28: 21411-21424.

Lou, D., Polk, C. and Skouras, S. (2019). A tug of war: Overnight versus intraday expected returns. *Journal of Financial Economics*, 134(1): 192-213, https://doi.org/10.1016/j.jfineco.2019.03.011.

Li, J., Zhu, S., Wu, Q., and Zhang, P. (2018). A Hybrid Forecasting Model Based on EMD-GASVM-RBFNN for Power Grid Investment Demand. *Mathematical Problems in Engineering*, 17 pages. https://doi.org/10.1155/2018/7416037.

Li, J., Wu, Y. and Fan, L. (2021). Monthly Henry Hub natural gas spot prices forecasting using variational mode decomposition and deep belief network. *Energy*, 227: 120478.

Lin, C., Chang, H. and Chou, R.K. (2023). Overnight versus intraday returns of anomalies in China. *Pacific-Basin Finance Journal*, 102007 https://doi.org/10.1016/j.pacfin.2023.102007.

Liu, C., Sun, B., Zhang, C. and Li, F. (2020). A Hybrid Prediction Model for Residential Electricity Consumption Using Holt-Winters and Extreme Learning Machine. *Applied Energy*, 275.

Ma, T., Antoniou, C. and Toledo, T. (2020). Hybrid machine learning algorithm and statistical time series model for network-wide traffic forecast. *Transp. Res. Part C Emerg. Technol*. 111: 352-372.

Pala, Z. and Unluk I.H. (2022). Comparison of hybrid and non-hybrid models in short-term predictions on time series in the R development environment. *Dicle University Journal of Engineering*, 13(2): 199-204.

Sobti, N. (2025). What triggers intraday price jumps and co-jumps in gold? *International Review and Financial Analysis*, 105, 104380. https://doi.org/10.106/j.i.irfa.104380

Xekalaki, E. and Degiannakis, S. (2010). *ARCH Models for Financial Applications*. Wiley: Chichester.

Ye, Z., Jiang, D. and Luo, Y (2023). Factor beta, overnight and intraday expected returns in China. *Global Finance Journal*, 56: 100827,https://doi.org/10.1016/j.gfj.2023.100827.

Zaini, N., Malek, M.A., Yusoff, N., Mardi, N.H. and Norhisham, S. (2018). Daily River Flow Forecasting with Hybrid Support Vector Machine – Particle Swarm Optimization. *IOP Conference Series: Earth and Environmental Science, Volume 140, 4th International Conference on Civil and Environmental Engineering for Sustainability (IConCEES 2017)* 4–5December2017*,* Langkawi*,* Malaysia.