

Article

A Transfer Learning Approach for Testing the Adaptive Market Hypothesis: Evidence from BWP/USD to Cryptocurrency Markets

Katlheho Makatjane ^{1,*} , Claris Shoko ²  and Tiisetso Makatjane ³ ¹ Department of Statistics and Population Studies, University of the Western Cape, Bellville 7535, South Africa² Department of Statistics, University of Botswana, Gaborone 0022, Botswana; shokoc@ub.ac.bw³ Department of Statistics and Demography, The National University of Lesotho, Roma 180, Lesotho; tjmakatjane@gmail.com

* Correspondence: kmakatjane@uwc.ac.za

Abstract

The efficient market hypothesis, which holds that prices completely reflect available information, is commonly used in financial market analysis. However, emerging empirical evidence shows that market efficiency develops with time, as posited by the adaptive market hypothesis (AMH), with predictability varying across shifting economic and behavioural regimes. Despite the increasing use of deep learning in financial forecasting, there has been little systematic investigation into whether neural network topologies can successfully identify time-varying efficiency trends across diverse markets. Furthermore, the relevance of transfer learning in studying adaptive behaviour between foreign exchange markets and extremely volatile cryptocurrency markets has received little attention. Using these data, we investigate the AMH by comparing the forecasting performance of various deep learning architectures and determining whether knowledge transfer from a relatively stable fiat currency market, Botswana Pula/US Dollar (BWP/USD), improves the predictive accuracy in a highly volatile cryptocurrency market, Bitcoin/US Dollar (BTC/USD). We use daily data from 1 January 2015 to 11 January 2026 to develop deep neural networks (DNNs) and alpha-recurrent neural networks, and, for generalisation, we benchmark using a recurrent temporal neural network (RTNN), a domain-adversarial neural network (DANN), and KLIEP-based importance-weighted regression. A transfer learning technique is used, in which models are initially trained on BWP/USD and then re-estimated on BTC/USD without freezing any network layers, ensuring complete flexibility and enabling parameters to respond to changing market dynamics. Out-of-sample accuracy measures and rolling long-memory diagnostics are used to evaluate forecast performance in terms of time-varying efficiency. The findings reveal that the RTNN regularly outperforms other forecasting models across marketplaces. Predictive accuracy fluctuates with time, and rolling long-memory measurements show persistent departures from random walk behaviour, which supports the AMH. Transfer learning improves predictive stability in the cryptocurrency market by identifying the existence of transferable informational structures between fiat and digital asset markets. Overall, our results support the idea that market efficiency is dynamic rather than static, and they show that adaptive deep learning systems are an excellent way to test the AMH. The paper suggests that cross-market transfer mechanisms and adaptive modelling methodologies be investigated further in growing foreign exchange and cryptocurrency markets.



Received: 10 January 2026

Revised: 28 February 2026

Accepted: 9 March 2026

Published: 29 June 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article

distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)[Attribution \(CC BY\) license](https://creativecommons.org/licenses/by/4.0/).

Keywords: cross-domain knowledge transfer; data-driven monetary policy; dynamic weighting mechanisms; exchange rate predictability; market regime shifts; time-varying efficiency; volatility adaptation

1. Introduction

Financial forecasting deals with identifying variables that meaningfully explain or predict a financial target, such as asset returns. However, the challenge lies in capturing relationships that are non-linear, time-varying, or depend on rare events. This calls for the selection of market features that hold stable predictive power under changing market dynamics (Fama 1970). The efficient market hypothesis (EMH) posits that markets are always efficient, with asset prices reflecting all available information, making it impossible to consistently "beat" the market through active management. Hence, there are a number of studies on the weak form of the EMH for cryptocurrency markets, particularly the Bitcoin market. For instance, Jakub (2015), Urquhart (2016), Nadarajah and Chu (2017), Bariviera (2017), Tiwari et al. (2018), Aggarwal (2019), and Lade and Yi (2020) have concluded that the Bitcoin market offers strong evidence regarding the non-weak form of efficiency. An alternative to the EMH is the adaptive market hypothesis (AMH) proposed by Lo (2004) as referenced by Mandaci et al. (2019), which is based on an evolutionary approach. The AMH argues that market efficiency is not a static condition but fluctuates, influenced by human behaviour and evolutionary principles. The adaptive market hypothesis is considered valid for exchange rates because studies show that market efficiency is not static but varies over time (Ito et al. 2014; Lim and Brooks 2011), allowing for periods of predictability. This is supported by evidence of episodic return predictability linked to events like central bank interventions or financial crises, which alter market conditions and participant behaviour. Therefore, the AMH offers a more realistic approach than the EMH by acknowledging that market participants adapt to changing environments, creating periods of both efficiency and inefficiency.

Exchange rate markets exhibit phases of high and low efficiency, which the AMH can capture. Events such as coordinated central bank actions, financial crises (for example, the Asian or global financial crises), and changes in economic conditions can temporarily improve or decrease return predictability. These events, therefore, influence how market participants react to and incorporate information into pricing, and research has uncovered evidence of both predictable and unpredictable periods in foreign exchange markets, which supports the AMH. Research by (Noda 2016) supports the adaptive market theory in Japanese stock markets, while (Karaomer and Acaravci 2023) similarly discovered that cryptocurrency market efficiency fluctuates with time. In contrast, markets are more efficient during periods of low volatility, and the AMH explains how market efficiency fluctuates as participants adapt to new conditions, competition, and information systems. This indicates that patterns, anomalies, and predictive correlations in financial markets can emerge, vanish, and resurface at different times and under different market regimes. This adaptive perspective naturally lends support to the transfer learning (TL) paradigm in exchange rate forecasting. Transfer learning, according to (Jallow et al. 2025), entails utilising a machine learning model pre-trained on one currency or market to improve predictions for another, leveraging knowledge from a source domain to boost performance in a target domain. However, exchange rate statistics are influenced by regionally specific monetary policies, macroeconomic considerations, and investor sentiments, posing issues for transfer learning. While the AMH and machine learning have been researched separately, there has been little research that explicitly uses transfer learning to bridge information between a

low-volatility currency market (such as BWP/USD) and a high-volatility cryptocurrency market. Leveraging information from a stable source domain to improve forecasts for a more volatile target domain not only puts the AMH to the test in practice but also addresses the data scarcity issues that are typical in new asset classes, making this method theoretically and practically relevant.

Forecasting in complex, ever-changing financial markets requires systems that can detect changes in market patterns and understand market participants' adaptive behaviour. Deep learning approaches, such as deep neural networks (DNNs) and alpha-recurrent neural networks (Alpha-RNNs), are ideal for capturing the non-linear, time-varying dynamics of financial markets that econometric and statistical models sometimes overlook. Deep neural networks learn hierarchical and complicated dependencies in high-dimensional data without prior assumptions, providing considerable advantages over linear models like exponential smoothing (Boonpan and Sarakorn 2025). Alpha-RNNs are specifically developed for time-varying financial markets and require fewer parameters and training data than other RNN versions, such as long short-term memory (LSTM) or gated recurrent units (GRUs). These architectures reduce overfitting and are able to handle both stationary and non-stationary data; they also provide a direct relationship between the smoothing parameter α and the dynamic behaviour of the model. Adaptive hybrid models, combining signal decomposition methods with deep neural networks and alpha-recurrent neural networks, have shown higher accuracy and resilience in volatile markets, as emphasised by (Liu et al. 2025; Shoko and Sigauke 2023; Wu et al. 2024). By combining the two architectures, these models are able to capture complex hierarchical aspects as well as temporal relationships, yielding a more complete method to predict time-varying exchange rates. Unlike these authors' approaches, the primary goal of our study is to evaluate and compare the forecasting performance of DNN and Alpha-RNN architectures in testing the adaptive market hypothesis, both without and with transfer learning from a low-volatility currency pair, Botswana Pula/US Dollar (BWP/USD), to a high-volatility cryptocurrency pair, Bitcoin/USD (BTC/USD). This enables us to see whether using information from a more stable source market improves the prediction accuracy in a more volatile target market, while also addressing the data scarcity issues typical in new asset classes.

For several reasons, the Botswana Pula/US Dollar exchange rate is preferred over other, more stable currencies like the South African Rand (ZAR/USD) or the Euro (EUR/USD) because BWP/USD historically shows less volatility than ZAR/USD, which is more vulnerable to regular macroeconomic and political shocks (Moyo et al. 2025). Transfer learning is appropriate because it provides a clearer contrast between a highly volatile destination market like BTC/USD and a stable, low-volatility source market. The BWP/USD market, according to (Molefhi 2021), is a smaller developing market whose financial characteristics are more akin to those of emerging, speculative markets like cryptocurrencies. This is in contrast to EUR/USD, which is impacted by broad global macroeconomic variables. Because BWP/USD historical data are often less noisy, deep learning architectures may identify stable patterns and associations that can be modified to predict very turbulent markets. Additionally, the use of BWP/USD adds novelty because most research focuses on major currency pairs; this allows us to investigate how predictions for high-volatility assets can be informed by small, stable, emerging market currencies, offering a unique test of the adaptive market hypothesis in a transfer learning setting. To augment the generalisability and robustness of the forecasting methodology, we broaden our analysis beyond DNNs and Alpha-RNNs by integrating supplementary domain adaptation techniques, including a domain-adversarial neural network (DANN), the Kullback–Leibler Importance Estimation Procedure (KLIEP), and a recurrent temporal neural network (RTNN). This expansive modelling approach enables a thorough comparison of the forecasting efficacy

with and without transfer learning, guaranteeing that any detected enhancements are not limited to a singular design. Ultimately, adaptive interval widths are integrated to address time-varying uncertainty in exchange rate dynamics, in contrast to (Hyndman and Rostami-Tabar 2025), who used prediction interval widths in their study on forecasting interrupted time series, because adaptive intervals better capture dynamic uncertainty, improve risk representation, and remain theoretically consistent with the AMH and our transfer learning design. Finally, explainable artificial intelligence methods, particularly Shapley Additive exPlanations (SHAP), are utilised to elucidate model predictions and ascertain the relative impacts of input features. Collectively, these enhancements bolster the empirical research by augmenting the forecast reliability, refining the interpretability, and offering profound insights into the adaptive behaviour of financial markets. The rest of this paper is structured as follows: Section 2 presents the methodology followed in the study; Section 3 describes its application; and, lastly, Section 4 provides the conclusions and recommendations.

2. Materials and Methods

In this section, we present the methods and procedures followed in the main empirical analysis.

2.1. Designing a Data Preprocessing Structure

This research formulates and assesses deep learning architectures to examine the adaptive market hypothesis regarding the BTC/USD exchange rate using a transfer learning approach. In particular, we use deep neural networks and alpha-recurrent neural networks, which have been trained on BWP/USD, and we examine how well they formulate predictions based on forecasting metrics. We have 4013 daily observations for each currency pair, from 1 January 2015 to 11 January 2026, and utilise daily OHLC (open, high, low, close) values for both BTC/USD and BWP/USD. The data come from Yahoo Finance <https://finance.yahoo.com/quote/BTC-USD/>, which was accessed on 13 January 2026 using the yfinance package in Python 3.12.1. as follows:

```
import yfinance as yf
import pandas as pd
data = yf.download(
    ["BTC-USD", "BWP=X"],
    start="2015-01-01",
    end="2026-01-11"
)

data.to_csv("BTC_BWP_data.csv")
```

We divide the dataset into training, validation, and test sets to retain the chronological sequence that is necessary for generating predictions. The first 80% of observations are utilised for training, the next 10% for validation, and the last 10% for testing. This method preserves the temporal integrity and prevents look-ahead bias by ensuring that future knowledge does not seep into previous periods. Using min-max scaling based on the training set, all input characteristics, such as OHLC (herein referenced Open, High, Low, Close) prices and derived technical indicators, are normalised. The same scaling settings are used for the validation and test sets. Nevertheless, our data processing system is legally responsible for the effective utilisation of the data, and they are submitted to the analytical server for predictive modelling. The system is primarily configured with high-performance computing (HPC) resources through the Amazon machine learning cloud using the Python environment. For training and large-scale data processing, AWS

EC2 instances are employed, providing scalable virtual servers that accommodate the computational demands of complex forecasting tasks. Real-time forecasting deployment is handled via AWS Lambda, which enables a serverless, on-demand execution environment that automatically scales, reduces the operational overhead, and ensures cost efficiency. Daily BTC/USD exchange rate data are transformed into log-returns to stabilise variance, thereby supporting a robust and reliable training and forecasting process. Missing values for BWP/USD occur on weekends and holidays because no trading takes place in this currency market and hence no profits or losses are realised by investors. Therefore, following (Ndlovu and Chikobvu 2023), we fill in these missing values with zeros in the raw price series solely to align the sequences of historical prices across markets. These aligned sequences are then used to compute log-returns, and only valid, non-zero prices contribute to the returns used as model inputs. Figure 1 presents the schematic workflow of our proposed transfer learning approach for the cryptocurrency market.




Figure 1. Workflow of the proposed transfer learning framework. The DNN, Alpha-RNN, RTNN, KLIEP, and DANN are trained and tested on BWP/USD data before applying transfer learning to BTC/USD. The fine-tuned models are deployed on AWS for real-time exchange rate forecasting. **Source:** Authors' own computation.

2.2. Deep Neural Networks

We use a feedforward deep neural network (DNN) to forecast the following day's price based on the preceding 30 days of historical data. The network learns a complex mapping function $f(\cdot)$ from input features to output forecasts. The input consists of a feature vector of n lagged returns, while the output is a single predicted price. The network comprises three hidden layers with diminishing neuron counts, namely 128 neurons in the first layer, 64 in the second, and 32 in the third, as seen in Figure 2. Each neuron computes the weighted sum of its inputs and applies a ReLU activation function to introduce non-linearity and enable the network to capture intricate patterns. Mathematically, the output of the l -th hidden layer is expressed as follows:

$$\mathbf{h}^{(l)} = \phi^{(l)} \left(\mathbf{W}^{(l)} \mathbf{h}^{(l-1)} + \mathbf{b}^{(l)} \right), \quad (1)$$

where $\phi^{(l)}$ denotes the activation function (i.e., ReLU), $\mathbf{W}^{(l)}$ is the weight matrix, $\mathbf{b}^{(l)}$ is the bias vector, and $\mathbf{h}^{(l-1)}$ is the output from the previous layer ($\mathbf{h}^{(0)}$ which corresponds to the input vector). The output layer consists of a single linear neuron that aggregates the transformed signals from the final hidden layer to produce the predicted price as,

$$\hat{y} = \mathbf{h}^{(L+1)}, \quad (2)$$

where $\hat{y}$ represents the forecasted value. The network is trained using the Adam optimiser, with the mean squared error as the loss function and the mean absolute error as the evaluation metric. This architecture efficiently models short-term dependencies in price sequences while maintaining computational simplicity.




Figure 2. Proposed DNN architecture. **Source:** authors' own computation.

2.3. Alpha-Recurrent Neural Networks

This research builds on the work of (Makatjane et al. 2025), who compared several recurrent neural network variations and coupled the best-performing Alpha-RNN with the TBATS model. We apply transfer learning to the BTC/USD exchange rate using their suggested Alpha-RNN and adjust it for univariate forecasting in accordance with the adaptive market hypothesis. The Alpha-RNN is an extension of the basic RNN with an extra smoothing parameter, α , and it regulates the contribution of the current hidden state to the exponentially smoothed hidden state. This architecture, unlike LSTM, does not have long-term or extracellular memory (Dixon 2022); instead, it regulates its effective memory dynamically using the smoothing parameter. In this study, we build a time-varying α_t , which allows the network to change its memory length according to changing market circumstances. At each time step, let s range from $t - p + 2$ to t and define the Alpha-RNN as

$$\hat{y}_{t+1} = \omega_y \hat{h}_t + b_y, \quad \hat{h}_s = g(U_h \tilde{h}_{s-1} + \omega_h y_s + b_h), \quad \tilde{h}_s = \hat{\alpha}_s \hat{h}_s + (1 - \hat{\alpha}_s) \tilde{h}_{s-1}. \quad (3)$$

where $g(\cdot)$ is the activation function, and $\hat{\alpha}_s \in [0, 1]$ is a dynamic smoothing parameter. Equation (3) thus comprises the output function, the hidden state update, and the smoothed hidden state update. Unrolling the recursion, this yields

$$\tilde{h}_t = \sum_{k=0}^{t-1} \hat{\alpha}_{t-k} \prod_{j=0}^{k-1} (1 - \hat{\alpha}_{t-j}) \hat{h}_{t-k} + \prod_{j=0}^{t-1} (1 - \hat{\alpha}_{t-j}) \tilde{h}_0, \quad (4)$$

which illustrates that the influence of past hidden states decays according to the learned $\hat{\alpha}_t$ values. It is therefore inaccurate to assert that $\alpha = 1$ produces infinite memory; in fact, when $\alpha_t = 1$, all contributions from previous smoothed states vanish, leaving only the current hidden state. Nonetheless, effective long-term memory arises only when $\alpha_t < 1$, with the network retaining past information recursively over multiple steps.

For illustration, consider a simplified autoregressive form obtained by setting $\phi = U_h = \omega_y = 1$ and $b_y = b_h = 0$, which gives a one-step-ahead forecast computed by

$$\hat{y}_{t+1} = \phi_1 y_t + \phi_2 y_{t-1} + \phi(1 - \hat{\alpha}_t) \tilde{h}_{t-2}, \quad (5)$$

where $\phi_1 = \phi$ and $\phi_2 = \hat{\alpha}_t \phi^2$. This form explicitly demonstrates that the contribution of earlier lags diminishes as $\hat{\alpha}_t$ approaches one, in full agreement with Equation (4).

The concept of stability is critical in time-series modelling, ensuring that the influence of perturbations diminishes over time. For the Alpha-RNN, stability is governed by the properties of the activation function $g(\cdot)$, which must satisfy $|g(x)| < 1$ for all x . Given an invertible non-linear lag operator $\Phi(B)$, the model may be expressed as

$$y_t = \Phi^{-1}(B) \epsilon_t, \quad (6)$$

where ϵ_t represents a white noise innovation. The smoothed hidden state $\tilde{h}_t$ satisfies

$$\tilde{h}_t = \hat{\alpha}_t \cdot g(U_h \tilde{h}_{t-1} + \omega_h x_t + b_h) + (1 - \hat{\alpha}_t) \cdot \tilde{h}_{t-1}, \quad (7)$$

and the network remains stable provided that the condition $|g(x)| < 1$ holds universally.

The Alpha-RNN is trained via stochastic gradient descent, with the time-varying smoothing parameter $\hat{\alpha}_t$ considered a learnable parameter. To increase generalisation and reduce overfitting, L2 regularisation is applied to the network weights, while suitable learning rates and momentum are utilised to expedite convergence and preserve stability. Early stopping is used with a defined patience parameter, and training is stopped if the performance on a validation set does not increase after a certain number of successive epochs. The model is trained using mini-batches of a fixed size across numerous epochs, enabling the network to successfully learn both short-term and long-term data dependencies. This training technique not only improves the network's capacity to adaptively capture memory via $\hat{\alpha}_t$ but also addresses the vanishing gradient issue common in basic RNNs. The dynamic smoothing parameter allows the Alpha-RNN to simulate non-stationary time series while retaining partial autocorrelation patterns, capturing complicated temporal correlations without the need for more complex designs like LSTM or GRUs (Almosova and Andresen 2023). The time-varying $\hat{\alpha}_t$ balances short-term responsiveness with long-term memory retention, creating a flexible and robust forecasting framework for the adaptive market hypothesis. Importantly, all architectures considered in this study use the same training approach, which includes L2 regularisation, learning rate scheduling, momentum, mini-batch training, patience-based early stopping, and epoch control. This ensures a fair and consistent comparison of the predictive performance.

3. Application to BTC/USD Exchange Rate

This section describes the application of the methods in Section 2 and the analysis of the results. This study seeks to provide a substantial test of the adaptive market hypothesis. We delineate the outcomes of using the deep neural networks and alpha-recurrent neural networks, along with the ultimate transfer learning methodology, in succession. The section starts with an examination of the deep neural network architecture, moves on to the exponentially smoothed recurrent neural network architecture, and concludes with the results of the combined time-varying dynamic weighted ensemble approach. The BWP/USD closing prices in Figure 3 demonstrate a consistent upward trend over time, especially from 2020 onwards, with the values fluctuating. The kernel density plot, as shown in Figure 3c, shows two main peaks at about 10.5 and 13.5, indicating that prices tend to cluster around these two levels. The Q-Q plot, as given in Figure 3b, indicates that

the data are approximately normally distributed; however, deviations at the tails suggest heavier tails and the presence of extreme values. The box plot in Figure 3d indicates a median of approximately 10.5, with the interquartile range between 10 and 13.5. The whiskers extend from 9 to 14, demonstrating that most prices are concentrated within this range, with no notable outliers present. These visualisations demonstrate moderate volatility, accompanied by sporadic extreme fluctuations. The same results are also reported by Mosanawe and Makatjane (2025) in their study of dynamic volatility and tail risk in BTC, BWP, and ZAR exchange rates.




Figure 3. BWP/USD closing price. Figure (a) represent a time series from 1 January 2015 to 11 January 2026; while Figure (b) is the kernel density plot. Moreover, Figure (c) represents the normal quantile-quantile plot with the red line representing a 45 line degree. Finally, Figure (d) is the normal Whisker's plot.

3.1. Forecasting Performance Without Transfer Learning

The analysis of the forecasting performance of the deep neural network, alpha-recurrent neural network, domain-adversarial neural network, KLIEP-based importance-weighted regression, and recurrent temporal neural network provides a comprehensive evaluation of the use of deep learning in forecasting the BWP/USD exchange rate. Table 1 summarises the comparative forecasting performance of all architectures in terms of bias, uncertainty calibration, and computational efficiency. Among the individual deep learning architectures, the DNN demonstrates strong and stable predictive performance, achieving a test mean absolute error of 0.042098 and a mean squared error of 0.003728, with a root mean squared error of 0.061057. This architecture exhibits a modest positive bias; we found that the MFE is 0.022016 and it maintains test coverage of 92.3% with a relatively narrow adaptive interval width of 0.1002. These results indicate that the DNN effectively captures the global and long-term dynamics of the BWP/USD exchange rate, producing reliable forecasts at a moderate computational cost of 45.01 s. The Alpha-RNN, which is designed to exploit temporal dependencies in sequential data, requires a substantially longer training time of 198.67 s. In terms of forecasting performance, it records a test MAE of 0.055856 and an MSE of 0.008288, with an RMSE of 0.091038. While its point forecast accuracy is weaker than that of the DNN, the Alpha-RNN achieves higher test coverage of 95.1% at the expense of wider predictive intervals, reported as 0.1484. This trade-off suggests that the Alpha-RNN captures short-term temporal variability more aggressively, resulting in increased uncertainty but improved test coverage. The RTNN emerges as the strongest individual performer across most forecasting metrics. It achieves the lowest MAE, which is recorded at 0.034632, and a competitive MSE of 0.003207, corresponding to an RMSE of 0.056630. Additionally, the RTNN maintains high test coverage of 95.3% with a narrow

AIW of 0.1016, while requiring only 36.92 s of training time. These results indicate that the RTNN offers an effective balance between predictive accuracy, uncertainty calibration, and computational efficiency.

In contrast, the domain-adversarial neural network performs poorly in this forecasting task. It records a very large MAE of 4.574795 and an MSE of 21.180155, accompanied by a substantial negative bias (MFE = -4.574795). Furthermore, the architecture exhibits zero test coverage and an excessively wide predictive interval of 2.6091. These results indicate severe negative transfer, where the adversarial domain alignment objective dominates the regression task, rendering the DANN unsuitable for forecasting the BWP/USD exchange rate in this setting. The KLIEP-based importance-weighted regression model offers a negligible computational cost, with a training time of only 0.16 s; however, this efficiency comes at the expense of predictive accuracy. The architecture produces a test MAE of 0.213480 and an MSE of 0.067188 with an RMSE of 0.25921. Although it achieves test coverage of 93.5%, the large AIW of 1.7282 indicates weak uncertainty calibration. These findings suggest that distribution reweighting alone is insufficient to capture the non-linear and temporal structure of the exchange rate series.

Table 1. Performance metrics for DNN, Alpha-RNN, and dynamic ensemble on BWP/USD.

Model	Training Time (s)	Test MAE	Test MSE	Test MFE	Test MAPE	Test RMSE	Test Coverage	AIW
DNN	45.01	0.042098	0.003728	0.022016	0.313552	0.061057	92.3%	0.1002
Alpha-RNN	198.67	0.055856	0.008288	0.003675	0.414716	0.091038	95.1%	0.1484
DANN	41.75	4.574795	21.180155	-4.574795	34.010705	4.602261	0.000%	2.6091
KLIEP	0.16	0.213480	0.067188	-0.183588	1.575491	0.25921	93.5 %	1.7282
RTNN	36.92	0.034632	0.003207	-0.000370	0.257361	0.056630	95.3%	0.1016

Finally, we evaluate the stability and generalisation of our RTNN architecture, which is an essential task to determine whether our forecasting approach can accurately predict price movement and prevent the underestimation of potential losses. The adaptive interval width of 0.1016, in conjunction with the 99% prediction interval depicted in Figure 4, offers valuable insights into the forecasting performance of the RTNN architecture and the uncertainty surrounding its forecasts. The relatively narrow AIW indicates that the architecture is capable of generating precise predictions while accounting for inherent uncertainties. The true values, represented by the red line, and the predicted ones, denoted by the blue line, remain predominantly within the 99% prediction interval (shaded black region), demonstrating that our RTNN architecture effectively captures the underlying BWP/USD price movement. This suggests that the RTNN architecture is robust in its ability to provide reliable out-of-sample forecasts—a crucial feature for economic and market applications. These findings are also in accordance with those found by (Makatjane et al. 2025) in their study of the downside risk of the South African mining index and those of (Mathonsi and Zyl 2025) in their study of multivariate anomaly detection based on prediction intervals constructed using deep learning.

The average SHAP values, as illustrated in Figure 5, further show that the prediction of the BWP/USD exchange rate is mostly based on information about the time structure and the opening price and not on extreme prices throughout the day. The monthly seasonality has the greatest effect, with a mean absolute SHAP value of 0.0126 and an opening price of 0.0108, while the opening price of 0.0108 and the weekly pattern of 0.0108 are close behind. This shows how important recurrent calendar patterns and opening price discovery are in determining how exchange rates move. These results are in line with those of previous studies that show how interpretable machine learning methods help in locating the most important time-based factors in financial forecasting models; see, for

instance, (Ergenç and Aktaş 2025). The findings indicate that periodic trade settlements, institutional foreign currency flows, and regular fiscal or commercial cycles are pivotal in BWP/USD dynamics (Honde 2024), especially in less liquid foreign exchange markets, where information is integrated at the commencement of the trading period. On the other hand, intraday and finer-scale temporal factors only slightly improve the forecast accuracy. This is seen in the low SHAP value for the daily calendar effect of 0.0011, the intraday low price given by 0.0012, and the quarterly indication of 0.0013. The year variable shows how the structure changes over time, with a mean absolute SHAP value of 0.0027. This means that short- to medium-term seasonal and market-opening influences have a larger impact on exchange rate behaviour, and this trend is consistent with previous research highlighting that accurately representing temporal hierarchies improves the interpretability and resilience in financial and economic forecasting, especially for volatility and tail risk applications (Ben Mrad and Alsowayan 2026; Chan et al. 2026).




Figure 4. Prediction interval widths without transfer learning.




Figure 5. SHAP values for feature importance for BWP/USD.

3.2. Testing the Adaptive Market Hypothesis Without Transfer Learning

Because the RTNN demonstrates high forecasting performance, we use this architecture to test the AHM and follow the work by (Hurs et al. 1965), as cited by (Henriques et al. 2020), and use the rolling Hurst exponent estimated over a 100-day timeframe. Our results, as tabulated in Table 2, show that the BWP/USD market has significant and long-lasting positive trending behaviour. The rolling mean Hurst exponent of 0.7284 indicates a strong

tendency toward trending behaviour; this is far higher than the random walk standard of $H = 50\%$. The minimum and highest H values of 0.5268 and 0.8176, respectively, reveal that the market always shows trending behaviour, with no windows indicating near-random or mean-reverting behaviour. In particular, 100% of the rolling windows show trending behaviour and less than 50% show mean reversion, while nearly 50% are close to random. These findings corroborate the idea that the Botswanan foreign currency market is mostly trend-following rather than consistently efficient, aligning with the AMH. The results also show that the adaptive RTNN is successful in capturing how the market changes over time, because it is able to adjust to both long-term trends and short-term changes, which makes its predictions more reliable in a complex and changing environment.

Table 2. Summary statistics of rolling Hurst exponent for BWP/USD.

Statistic	Value	Interpretation
Mean H	0.7284	Slightly trending
Min H	0.5268	Strong mean-reverting behaviour
Max H	0.8176	Strong trending behaviour
Proportion of trending windows ($H > 0.5$)	100.00%	Market exhibits trending behaviour
Proportion of near-random windows ($H \approx 0.5$)	0.00%	Market approximately random
Proportion of mean-reverting windows ($H < 0.5$)	0.00%	Market exhibits mean reversion

To support the results reported in Table 2, we also plot the rolling Hurst values, as presented in Figure 6. None of the rolling Hurst values in this figure come close to the market efficiency criterion of $H = 50\%$. Our findings corroborate those of (Gomez et al. 2022), who demonstrated that their enhanced Hurst exponent encapsulated the dynamics anticipated by the adaptive market hypothesis. The main difference is that they employed a considerably longer window of 1024 trading days (about four years), using a technical analysis (TA) approach, whereas, in this study, we simply use 100 trading days (about one year). Although the timeframe is shorter, our results still show the time-varying efficiency patterns that the AMH describes.




Figure 6. Rolling Hurst exponent for BWP/USD closing prices with 100-day window. The red dashed line at $H = 0.5$ indicates the benchmark for an efficient market.

Moreover, these results provide a clear picture of the BWP/USD exchange rate movement. The high prediction performance and low bias suggest that the market has a structure that changes with time, which is not reflective of rigid efficiency but is consistent with the adaptive market hypothesis of Lo (2004). Our established RTNN accurately captures both short-term variations and longer-term trends, giving policymakers and business organisations solid predictions that could help with interventions, reserve management, and hedging strategies. The rolling Hurst exponent results support these findings because the market shows significant trend-following behaviour with sustained predictability, as seen

by the mean of 0.7284 and 100% of rolling windows moving. These changes show how important it is to have an adaptable architecture that changes with the market in real time.

3.3. Forecasting Performance with Transfer Learning

Transfer learning enables models trained in one domain to adapt their learnt representations to a similar but separate task. In this study, we compare the performance of five forecasting architectures, the DNN, Alpha-RNN, DANN, KLIEP, and RTNN, when moving from the relatively stable BWP/USD exchange rate to the volatile BTC/USD market. Table 3 presents the performance results of these architectures in transfer learning, where the RTNN has the highest overall forecasting performance, with the lowest MSE recorded at 0.000042, while the MAE is found to be 0.005083, the RMSE is 0.006486, and the MAPE is 0.55%. Moreover, it is discovered that this architecture produces good test coverage of 99.0% and a short AIW of 0.0176. This shows that the RTNN successfully responds to BTC/USD volatility while maintaining low bias at 0.001064. The Alpha-RNN achieves high forecasting performance in point predictions, with an MSE of 0.004757, an MAE of 0.53187, and decent test coverage of 97.1%, but its lengthy training period of 10,129 s limits its usefulness. The pre-trained DNN for BWP/USD achieves a moderate performance MSE of 0.008288 and an MAE of 0.055856 with good coverage of 95.1%, but it has a comparatively broader AIW of 0.1484, suggesting lesser accuracy in short-term BTC/USD swings. The DANN, on the other hand, performs rather well, with an MSE of 0.030277 and MAE of 0.170335, but with lesser test coverage at 96.6% and a greater AIW of 0.240, indicating difficulties in modelling fast domain transitions despite adversarial training. Finally, KLIEP with ridge regression achieves an exceptionally short training time of 0.01 s and perfect coverage at 100%, but its MSE of 7.297906 together with an MAE of 2.665 indicate low pointwise accuracy, rendering it unsuitable for accurate short-term forecasting in volatile markets. Overall, the RTNN is the superior transfer learning model, demonstrating a harmonious balance among predicted accuracy, interval dependability, and flexibility to BTC/USD volatility. The Alpha-RNN and DNN successfully capture overarching patterns, but KLIEP and DANN demonstrate difficulties in accurate pointwise forecasting. These results correspond with prior studies suggesting that deep networks may struggle in volatile markets owing to gradient instability and transfer domain discrepancies (Rossolini et al. 2022).

Table 3. Performance comparison of forecasting models.

Model	Train Time (s)	MSE	MAE	RMSE	MFE	MAPE (%)	AIW (Δ)	Test Coverage
DNN	70.02	0.015709	0.116258	0.125334	−0.116258	12.29	0.208035	0.988
Alpha-RNN	10,129.02	0.004757	0.053187	0.218847	−0.021713	0.47	0.015551	0.971
DANN	84.41	0.030277	0.170335	0.174004	−0.170335	18.32	0.240003	0.966
KLIEP	0.01	7.297906	2.665002	2.701464	−2.665002	23.60	3.532006	1.000
RTNN	44.97	0.000042	0.005083	0.006486	0.001064	0.55	0.017624	0.990

Figure 7 presents the fitted conditional mean for the RTNN together with the 99% AIW. We use the adaptive interval widths because they are dynamic, data-driven, and context-aware boundaries for predictions or confidence estimates, optimising accuracy by widening in high-uncertainty zones and narrowing in low-uncertainty ones. Unlike the prediction interval (PI), which uses fixed intervals, as emphasised by (Hajibabae et al. 2024), the AIW improves performance in tasks like wind speed forecasting, heteroscedastic regression, and risk assessment. According to visual examination, the architecture fits the BTC/USD closing prices well. The narrow interval of 0.017624 with coverage of 99%

indicates that the architecture achieves high forecasting performance with low uncertainty in out-of-sample forecasting.



Figure 7. Fitted conditional mean for BTC/USD closing prices coupled with AIW.

While the Bitcoin market is highly unstable, as emphasised by (Aizenman 2019), our training process for the RTNN includes temporal hierarchies at the daily, weekly, monthly, quarterly, and annual levels, as well as basic price attributes like open, high, and low. These are included to determine which one has the greatest effect on transfer learning models when applied to the Bitcoin market. To measure their impacts, we use the methods described in (Seropola et al. 2025) and emphasise SHAP analysis, which is an explainable artificial intelligence tool. Figure 8 shows the findings, which highlight the feature significance that comes from the RTNN design. It is clear that the open price feature, which is the opening price at which BTC is traded, has the greatest effect on how well the model predicts prices. The high and month features come next. The temporal hierarchy elements have a considerable effect on the month and quarter, whereas the day, week, and year have a smaller effect. The typical absolute SHAP values are as follows: open = 0.006664, high = 0.010975, low = 0.006515, day = 0.001822, week = 0.002780, month = 0.004911, quarter = 0.003648, and year = 0.000758. The findings show that both price levels and time patterns are very important for the RTNN's success, with the highest prices and long-term time impacts being the most important.




Figure 8. SHAP values for feature importance.

3.4. Transfer Learning to Test the Adaptive Market Hypothesis

In order to test the adaptive market hypothesis in transfer learning, we use the rolling Hurst exponent, which is in line with the technique presented in Section 3.2 and follows the empirical analysis of (Kroha and Skoula 2018). The rolling Hurst exponent measures how efficient the market is over time by considering long-range dependencies and persistence in asset returns. The statistics for the BTC/USD exchange rate show that the market is quite active and stable. Table 4 shows that the mean rolling Hurst exponent is 0.7833, with values ranging from 0.6330 to 0.8396. All values are much higher than the efficient market benchmark of $H = 0.5$, which means that there are strong long-memory effects and trend persistence in BTC/USD price movement, as seen in the BWP/USD market as well. This suggests that the persistence characteristics are preserved even when knowledge is shared across markets. It is important to further note that 100% of the rolling windows show $H > 0.5$, and none of the windows show near-random behaviour with $H \approx 0.5$ or mean-reverting behaviour with $H < 0.5$. The prevalence of persistent regimes indicates a consistent departure from weak-form efficiency across the entire data period. This tendency is completely in line with the AMH, which states that market efficiency changes over time instead of being static. These findings are achieved using a small rolling window of 100 trading days, which is important since it shows that adaptable and enduring market structures can be found even over short time periods.

Table 4. Summary statistics of the rolling Hurst exponent for BTC/USD.

Statistic	Value
Mean Hurst exponent	0.7833
Minimum Hurst exponent	0.6330
Maximum Hurst exponent	0.8396
Trending regimes ($H > 0.5$)	100.00%
Near-random regimes ($H \approx 0.5$)	0.00%
Mean-reverting regimes ($H < 0.5$)	0.00%

Figure 9 graphically corroborates these results by showing persistent departures of the rolling Hurst exponent over the efficiency threshold. The constantly high rolling Hurst exponent seen in the BTC/USD exchange rate has profound effects on the economy and the market. A Hurst exponent that is always higher than the efficient market benchmark of $H = 0.5$ means that price dynamics are marked by long-range dependency and trend persistence. This means that previous price movements include economically useful information about future returns. This consistent departure from weak-form efficiency supports the adaptive market hypothesis, which asserts that market efficiency changes in response to changing market circumstances, rather than being static. Our results corroborate those presented by (Raju et al. 2025), who demonstrated that their strategies yielded superior returns per unit of risk, hence verifying the adaptive market hypothesis via the emphasis on cyclical fluctuations in market profitability. Moreover, they validate the efficacy of Sharpe ratio-based strategies compared to the buy-and-hold method, supporting the concept of changing market efficiency in accordance with the adaptive market hypothesis.

The resemblance between the persistence characteristics of BTC/USD and those previously identified in BWP/USD indicates that long-memory effects constitute a structural attribute that endures across markets, even when transitioning knowledge from a comparatively stable fiat currency to a highly volatile cryptocurrency. This discovery has significant ramifications for cross-market forecasting and policy evaluation, suggesting that models developed for established markets may preserve economically relevant information when used in nascent digital asset markets. From the point of view of a market participant(s),

the fact that persistent regimes are in charge means that tactics based on following trends and momentum work better than static trading methods. (Erwanto 2026) further indicates that a trend-following approach seeks to benefit from long-term market movement by responding to pricing behaviour rather than anticipating outcomes. Trend trading, therefore, works by immediately reducing losses, allowing winnings to run. Moreover, trends do not materialise every time, as (Tsoku and Makatjane 2026) emphasises by demonstrating that the computed dynamic spread for pairs trading strategy initiates at a value close to zero and progressively decreases until the conclusion of 2021, indicating a substantial negative divergence from equilibrium. The spread, as reported in the mentioned work, demonstrates significant peaks during major market rallies, accompanied by abrupt declines and heightened volatility periods. From 2022 to late 2025, the spread exhibits fluctuations, characterised by both upward and downward trends, along with periodic deviations exceeding the ± 1 standard deviation. The observed dynamics indicate periods when the market might be overvalued or undervalued relative to its long-term equilibrium, highlighting potential opportunities for mean reversion. According to (Monge et al. 2025), these patterns can be used in a data-driven manner to produce actionable trading signals for cryptocurrency pairing. For investors and risk managers, the existence of long-memory dynamics indicates that disturbances in the BTC/USD market are likely to persist over time, underscoring the need for flexible risk management frameworks that can react to changing market conditions. Furthermore, the capacity to identify persistent behaviour within a short rolling timeframe underscores the practical applicability of the methodology for real-time monitoring, trading, and regulatory supervision in rapidly evolving cryptocurrency markets.




Figure 9. Rolling Hurst exponent for BTC/USD closing prices with 100-day window. The red dashed line at $H = 0.5$ indicates the benchmark for an efficient market.

4. Conclusions and Recommendations

This study examined the validity of the adaptive market hypothesis in both fiat and cryptocurrency markets through the application of deep learning architectures, namely deep neural networks, alpha-recurrent neural networks, domain-adversarial neural networks, KLIEP-based importance-weighted regression, and recurrent temporal neural networks. Utilising the BWP/USD exchange rate as the baseline market and BTC/USD as the transfer learning objective, the empirical findings consistently corroborate the fundamental assertion of the AMH: that market efficiency is dynamic rather than fixed. Without transfer learning, the RTNN was the best model for the BWP/USD market. It had better predictive accuracy, less bias, and better-calibrated uncertainty, as shown by its small adaptive interval widths and high coverage rates. These findings showed that the RTNN is successful in capturing both long-term structural trends and short-term changes in the exchange rate. The rolling Hurst exponent analysis confirmed these results by showing that all rolling windows had persistent trending behaviour. This disproved the idea of weak-form efficiency and confirmed that the market is adaptable.

With respect to (Feng and Zhang 2023), transferring information from a stable, controlled market like the Botswana Pula foreign market to a highly volatile, sentiment-driven market like Bitcoin necessitates the adaptation of technical abilities, as well as a significant adjustment in risk management measures. While BWP/USD is reasonably stable and has a predictable monetary policy, BTC/USD is volatile, trades around the clock, and, in certain situations, gives predictive insights for other currency pairings. Hence, our transfer of knowledge from the more stable BWP/USD market to the very unstable BTC/USD market gives further proof that the AMH is correct. Specifically, our established RTNN consistently showed exceptional predictive performance in the cryptocurrency sector, having the lowest forecasting errors, minimal bias, and robust uncertainty calibration. The maintenance of long-memory traits across markets, shown by consistently elevated rolling Hurst exponent values, implies that persistence is a fundamental attribute that transcends market classification. This conclusion emphasises the economic significance of moving acquired representations from established fiat markets to nascent digital asset markets.

The SHAP explainability analysis adds to the empirical value of this investigation. The findings indicated that price-based attributes, especially the high price, significantly impacted the RTNN's forecasting performance, followed by open and low prices. Temporal hierarchy features such as month and quarter also played a meaningful role, while day, week, and year exhibited comparatively lower impacts. These results suggest that both price formation processes and temporal patterns are important in understanding how markets adapt, particularly when using transfer learning.

From an economic and policy perspective, the results carry important implications. For policymakers, particularly central banks such as the Bank of Botswana, adaptive deep learning models can enhance exchange rate monitoring, reserve management, and the timing of foreign exchange interventions by providing more accurate and uncertainty-aware forecasts. For investors and risk managers, the dominance of persistent regimes implies that trend-following and momentum-based strategies may outperform static trading rules, particularly in markets characterised by long-memory dynamics. However, reliance on algorithmic forecasts should be addressed by continuous validation and prudent interpretation to mitigate systemic risks arising from homogeneous model-driven behaviour.

Recommendations

Future research should expand this paradigm to include real-time forecasting and decision support systems, especially in emerging and frontier markets. Further investigation of Transformer-based architectures and reinforcement learning models may improve the flexibility and interpretability during regime transitions. Furthermore, broadening the research to include other asset classes, macroeconomic factors, and alternative market efficiency metrics would offer a more complete picture of adaptive efficiency in a variety of economic circumstances. Overall, this work reveals that deep learning-based adaptive models provide a solid and empirically grounded method for assessing developing market efficiency in both conventional and digital financial systems.

Author Contributions: Conceptualisation, C.S.; methodology, K.M.; Validation, C.S. and T.M.; Formal analysis, K.M.; Investigation, K.M.; Data curation, K.M.; Writing—original draft, K.M.; Writing—review and editing, C.S. and T.M.; Visualization, K.M. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The data presented in this study are openly available in <https://github.com/Makatjane/Data> (accessed on 13 January 2026).

Conflicts of Interest: The authors declare no conflict of interest.

References

- Aggarwal, Divya. 2019. Do Bitcoins follow a random walk model? *Research in Economics* 73: 15–22. [\[CrossRef\]](#)
- Aizenman, Joshua. 2019. On the Built-In Instability of Cryptocurrencies. *VoxEU Column*, February 12. Available online: <https://cepr.org/voxeu/columns/built-instability-cryptocurrencies> (accessed on 22 February 2026).
- Almosova, Anna, and Niek Andresen. 2023. Nonlinear inflation forecasting with recurrent neural networks. *Journal of Forecasting* 42: 240–59. [\[CrossRef\]](#)
- Bariviera, Aurelio F. 2017. The Inefficiency of Bitcoin Revisited: A Dynamic Approach. *Economics Letters* 161: 1–4. [\[CrossRef\]](#)
- Bartos, Jakub. 2015. Does Bitcoin Follow the Hypothesis of Efficient Market? *International Journal of Economic Sciences* 4: 10–23. [\[CrossRef\]](#)
- Ben Mrad, Ali, and Hamad Mohammad Alsowayyan. 2026. Interpretable machine learning for price index forecasting: A case study with rolling windows and SHAP. *Engineering, Technology & Applied Science Research* 16: 30954–962. [\[CrossRef\]](#)
- Boonpan, Samak, and Weerachai Sarakorn. 2025. Deep Neural Network Model Enhanced with Data Preparation. *Journal of Open Innovation* 11: 100438. [\[CrossRef\]](#)
- Chan, Chia-Pang, Fang-Kai Tang, and Jun-He Yang. 2026. Financial Performance Prediction and Stability Analysis Using SHAP-Enhanced Machine Learning Models. *Computational Economics*. [\[CrossRef\]](#)
- Dixon, Matthew. 2022. Industrial Forecasting with Exponentially Smoothed Recurrent Neural Networks. *Technometrics* 64: 114–24. [\[CrossRef\]](#)
- Ergenç, Cansu, and Rafet Aktaş. 2025. Sector-specific financial forecasting with a machine learning algorithm and SHAP interaction values. *Financial Internet Quarterly* 21: 42–66. [\[CrossRef\]](#)
- Erwanto. 2026. Mastering Trend Following Strategy: How It Works and Misconceptions. *Go Trade*, January 12. Reviewed by GoTrade Internal Analyst. Available online: <https://www.heygotrade.com/en/blog/mastering-trend-following-strategy/> (accessed on 15 February 2026).
- Fama, Eugene F. 1970. Efficient Capital Markets: A Review of Theory and Empirical Work. *The Journal of Finance* 25: 383–417. [\[CrossRef\]](#)
- Feng, Wenjun, and Zhengjun Zhang. 2023. Currency exchange rate predictability: The new power of Bitcoin prices. *Journal of International Money and Finance* 132: 102811. [\[CrossRef\]](#)
- Gómez-Águila, A., J. E. Trinidad-Segovia, and M. A. Sánchez-Granero. 2022. Improvement in Hurst Exponent Estimation and Its Application to Financial Markets. *Financial Innovation* 8: 86. [\[CrossRef\]](#)
- Hajibabaei, Parisa, Farhad Pourkamali-Anaraki, and Mohammad Amin Hariri-Ardebili. 2024. Adaptive Conformal Prediction Intervals Using Data-Dependent Weights With Application to Seismic Response Prediction. *IEEE Access* 12: 53579–97. [\[CrossRef\]](#)
- Henriques, Teresa, Maria Ribeiro, Andreia Teixeira, Luísa Castro, Luís Antunes, and Cristina Costa-Santos. 2020. Nonlinear Methods Most Applied to Heart-Rate Time Series: A Review. *Entropy* 22: 309. [\[CrossRef\]](#)
- Honde, George. 2024. Botswana: Challenges and Policy Responses to Avoid the Trap. In *Avoiding the Middle-Income Trap in Africa: Economic Challenges and Policy Responses*. Cham: Springer Nature Switzerland, pp. 207–58. [\[CrossRef\]](#)
- Hurst, H. E., R. P. Black, and Y. M. Simaika. 1965. *Long-Term Storage: An Experimental Study*. London: Constable and Co. Ltd.
- Hyndman, Rob J., and Bahman Rostami-Tabar. 2025. Forecasting interrupted time series. *Journal of the Operational Research Society* 76: 790–803. [\[CrossRef\]](#)
- Ito, Mikio, Akihiko Noda, and Tatsuma Wada. 2014. International Stock Market Efficiency: A Non-Bayesian Time-Varying Model Approach. *Applied Economics* 46: 2744–54. [\[CrossRef\]](#)
- Jallow, Haruna, Ronald Waweru Mwangi, Alieu Gibba, and Herbert Imboga. 2025. Transfer learning for predicting of gross domestic product growth based on remittance inflows using RNN-LSTM hybrid model: A case study of The Gambia. *Frontiers in Artificial Intelligence* 8: 1510341. [\[CrossRef\]](#)
- Karaomer, Yunus, and Songül Kakilli Acaravci. 2023. Adaptive Market Hypothesis: Evidence from the Cryptocurrency Market. *Iranian Journal of Management Studies* 16: 125–38.
- Kroha, Petr, and Miroslav Skoula. 2018. Hurst Exponent and Trading Signals Derived from Market Time Series. Paper presented at 20th International Conference on Enterprise Information Systems (ICEIS 2018), Funchal, Portugal, March 21–24; pp. 371–78. Available online: <https://www.scitepress.org/PublishedPapers/2018/66670/66670.pdf> (accessed on 5 January 2026).
- Lade, Suvarna, and Junesuh Yi. 2020. Does the South Korea Bitcoin Market Is Efficient? *International Journal of Management* 11: 1592–97.
- Lim, Kian-Ping, and Robert Brooks. 2011. The Evolution of Stock Market Efficiency Over Time: A Survey of the Empirical Literature. *Journal of Economic Surveys* 25: 69–108. [\[CrossRef\]](#)
- Liu, Zihan, Zijia Zhang, and Weizhe Zhang. 2025. A Hybrid Framework Integrating Traditional Models and Deep Learning for Multi-Scale Time Series Forecasting. *Entropy* 27: 695. [\[CrossRef\]](#)
- Lo, Andrew W. 2004. The Adaptive Markets Hypothesis: Market Efficiency from an Evolutionary Perspective. *Journal of Portfolio Management* 30: 15–29. [\[CrossRef\]](#)
- Makatjane, Katleho, Claris Shoko, and Caston Sigauke. 2025. An End-to-End Combined Forecasting Architecture. *Journal of Statistics Applications and Probability* 14: 59–75. [\[CrossRef\]](#)

- Mandaci, Pinar Evrim, Dilvin Taşkın, and Zeliha Can Ergün. 2019. Adaptive Market Hypothesis. *International Journal of Economics and Business Administration* 7: 84–101. [\[CrossRef\]](#)
- Mathonsi, Thabang, and Terence L. van Zyl. 2025. Multivariate anomaly detection based on prediction intervals constructed using deep learning. *Neural Computing and Applications* 37: 707–21. [\[CrossRef\]](#)
- Molefhi, Koketso. 2021. The impact of macroeconomic variables on capital market development in Botswana's economy. *African Journal of Economic Review* 9: 204–22. [\[CrossRef\]](#)
- Monge, Manuel, Rafael Hurtado, and Juan Infante. 2025. Time trends and persistence of the return difference between growth and value investment strategies. *PLoS ONE* 20: e0332690. [\[CrossRef\]](#) [\[PubMed\]](#)
- Mosanawe, Letlhogonolo, and Katleho Makatjane. 2025. Dynamic Volatility and Tail Risk in BTC, BWP, and ZAR Exchange Rates: Bayesian SARMA-GARCH with Skewed Error Distributions. In *Statistics, Optimisation and Information Computing*. Melbourne: International Academic Press. Available online: <http://www.iapress.org/index.php/soic/article/view/2883> (accessed on 5 January 2026).
- Moyo, Edwin, Katleho Makatjane, and Ramasamy Sivasamy. 2025. An Improved Volatility Forecasting with GMLE-Based MS-GARCH Models: Evidence from BWP/USD, ZAR/USD, and BTC/USD. *Journal of Statistics Applications and Probability* 14: 343–60. [\[CrossRef\]](#)
- Nadarajah, Saralees, and Jeffrey Chu. 2017. On the Inefficiency of Bitcoin. *Economics Letters* 150: 6–9. [\[CrossRef\]](#)
- Ndlovu, Thabani, and Delson Chikobvu. 2023. Extreme Risk Comparison of Bitcoin/USD and ZAR/USD. *Risks* 11: 100. [\[CrossRef\]](#)
- Noda, Akihiko. 2016. A Test of the Adaptive Market Hypothesis Using a Time-Varying AR Model in Japan. *Finance Research Letters* 17: 66–71. [\[CrossRef\]](#)
- Raju, Ajay, Uday Kumar Jagannathan, and Triveni Pundalika. 2025. Empirical analysis of the adaptive market hypothesis using the Sharpe Ratio for buy and sell opportunities in global indices. *The Economics and Finance Letters* 12: 120–42. [\[CrossRef\]](#)
- Rossolini, Giulio, Alessandro Biondi, and Giorgio Buttazzo. 2022. Increasing the Confidence of Deep Neural Networks by Coverage Analysis. *IEEE Transactions on Software Engineering* 49: 802–15. [\[CrossRef\]](#)
- Seropola, Boago, George Anderson, and Katleho Makatjane. 2025. Using Machine Learning and Explainable Artificial Intelligence for Election Voting Prediction. Paper presented at 2025 IEEE/ACIS 23rd International Conference on Software Engineering Research, Management and Applications (SERA), Las Vegas, NV, USA, May 29–31; pp. 1–8. [\[CrossRef\]](#)
- Shoko, Claris, and Caston Sigauke. 2023. Short-Term Forecasting of COVID-19 Using Support Vector Regression. *American Journal of Infection Control* 51: 1095–107. [\[CrossRef\]](#)
- Tiwari, Aviral Kumar, Rabin K. Jana, Debojyoti Das, and David Roubaud. 2018. Informational Efficiency of Bitcoin—An Extension. *Economics Letters* 163: 106–9. [\[CrossRef\]](#)
- Tsoku, Johannes Tshepiso, and Katleho Makatjane. 2026 Deep Learning-Based Pairs Trading: Real-Time Forecasting of Cointegrated Cryptocurrency Pairs. *Frontiers in Applied Mathematics and Statistics* 12: 1749337. [\[CrossRef\]](#)
- Urquhart, Andrew. 2016. The Inefficiency of Bitcoin. *Economics Letters* 148: 80–82. [\[CrossRef\]](#)
- Wu, Shujun, Zengchuan Dong, Sandra M. Guzmán, Gregory Conde, Wenzhuo Wang, Shengnan Zhu, Yiqing Shao, and Jinyu Meng. 2024. Two-Step Hybrid Model for Monthly Runoff Prediction. *Ecological Informatics* 84: 102914. [\[CrossRef\]](#)

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.