

A Critical Look At Smart Systems For Hedging Foreign Exchange Risk

Dr. Nevan T. Marcellin

Department of Financial Engineering, Baltic Institute of Economics and Systems, Riga, Latvia

Prof. Eshwari V. Thaloor

Centre for Applied FinTech Research, Amity School of Business, Noida, India

VOLUME02 ISSUE01 (2025)

Published Date: 05 January 2025 // Page no.: - 1-9

ABSTRACT

In today's global economy, fluctuating currency exchange rates are a constant financial headache for businesses, investors, and even entire nations. To protect themselves, companies often "hedge" their currency exposure, much like buying insurance against unfavorable market swings. But deciding when and how to buy this insurance has always been a tough call. This paper explores the shift from relying on gut feelings and basic financial tools to using smart, data-driven "expert systems" to make these decisions. We take a deep dive into how these systems have grown from simple rule-based methods to advanced machine learning models. By reviewing a wide range of research, we'll look at the technologies they use, how well they actually work, and what they mean for the real world of risk management. Our findings show that while no system is a crystal ball, modern machine learning approaches are remarkably good at spotting complex patterns in financial data, often outperforming older methods. However, making them work reliably requires careful setup, rigorous testing, and a healthy respect for their limitations in the unpredictable world of finance. We conclude that these expert systems are a powerful new tool, offering a smarter, more adaptive way to manage the challenges of the foreign exchange market.

Keywords: Foreign Exchange Risk, Hedging, Expert Systems, Decision Support Systems, Machine Learning, Deep Learning, Random Forest, XGBoost, LSTM, Technical Analysis, Time Series Forecasting, Financial Risk Management.

1. INTRODUCTION

1.1 The Modern Business Dilemma: A World of Opportunity and Risk

Imagine you run a business that sells products all over the world. Every sale in a foreign currency, every invoice you pay to an overseas supplier, opens you up to a world of opportunity. But it also exposes you to a subtle, ever-present risk: the value of currencies is constantly changing. The profit you expected from a big European sale could shrink overnight simply because the euro weakened against the dollar. This is the daily reality of foreign exchange (FX) risk, and it's not just a minor accounting issue; it can dramatically impact a company's revenue, costs, and overall financial health [9, 69]. Managing this risk has become a critical strategic challenge for any modern, global enterprise [2, 59].

1.2 Understanding the Three Faces of Currency Risk

Foreign exchange risk isn't a single problem; it shows up in a few different ways [69]:

- Transaction Risk: This is the most straightforward

type. You've agreed to a price in a foreign currency, but you won't get paid for another 90 days. In that time, the exchange rate could move against you, and the payment you receive will be worth less in your home currency than you planned for. It's the risk tied to the day-to-day deals of the business [4, 66].

- Translation Risk: This is more of an accounting challenge. If your company owns a subsidiary in another country, you have to translate its financial statements (assets, liabilities, income) back into your home currency for your consolidated reports. A dip in the exchange rate can make a profitable subsidiary look less valuable on paper [9].

- Economic Risk: This is the big-picture risk. It's about how exchange rate swings can affect your company's entire market position and future profitability. For example, a strong home currency might make your products more expensive for foreign customers, allowing your international competitors to undercut you on price [9, 12].

1.3 The Old Way of Hedging: A Costly Insurance Policy

The traditional way to handle these risks is to hedge, most commonly using a forward contract. This is essentially an agreement with a bank to lock in an exchange rate for a future date [66]. It removes the uncertainty, which is great, but it comes at a cost. Banks charge commissions or build their fees into the rate, and these costs add up, especially for businesses with a lot of international transactions [60, 64].

This creates a difficult choice. Do you hedge everything and accept the high, certain costs? Or do you hedge nothing and leave yourself completely exposed to market volatility? This is the central problem that has pushed financial managers to look for a smarter way.

1.4 A New Approach: Can We Teach a Machine to Be a Better Hedger?

The search for a better way has led to the exciting world of computational finance, where we use the power of artificial intelligence (AI) and machine learning (ML) to tackle financial challenges [53]. Instead of relying on static rules or human intuition alone, we can now build expert decision systems. These are smart systems that learn from historical data to spot patterns and make objective, data-driven recommendations [3, 48].

The earliest versions of these systems were based on technical analysis, which involves looking for patterns in historical price charts [1, 58]. But the real game-changer has been the arrival of modern machine learning. Algorithms like Random Forests [14], XGBoost [18], and deep learning models like LSTMs [34] can analyze vast amounts of data and uncover complex, non-obvious relationships that no human ever could [10, 52, 72].

1.5 What This Paper Is About

This paper is a deep dive into these expert systems, looking specifically at how they can be used as a form of "exchange rate insurance." We want to see if they can help a company make better hedging decisions. Our goal is to bridge the gap between academic research that focuses on using ML to "beat the market" and the practical, real-world need of a business to simply manage its risk effectively.

Here's our roadmap:

- Section 2: We'll explore the big ideas and past research that set the stage for this work.
- Section 3: We'll get into the nitty-gritty of how we built the predictive brain of our expert system.
- Section 4: We'll describe how we designed the complete system and set up a fair test to see how well it performed.

- Section 5: We'll share the results—did the smart system actually make better decisions?
- Section 6: We'll discuss what the results really mean, the challenges of using such a system, and where this technology could go next.
- Section 7: We'll wrap it all up with our main conclusions.

2. STANDING ON THE SHOULDERS OF GIANTS: CORE CONCEPTS

Before we build our own system, it's important to understand the key ideas and debates that have shaped this field.

2.1 Are Markets Predictable? A Tale of Two Theories

At the heart of our quest is a fundamental question: can financial markets be predicted at all? There are two main schools of thought.

- The Efficient Market Hypothesis (EMH): This classic theory, championed by Eugene Fama, argues that asset prices already reflect all available information [23]. If that's true, then future price movements are essentially random and unpredictable. Trying to find patterns is a waste of time, and the best you can do is diversify your investments and hope for the best.
- The Adaptive Markets Hypothesis (AMH): A more modern view, proposed by Andrew Lo, suggests that markets are more like a biological ecosystem than a perfect machine [55]. In this view, investors are like competing animals, constantly adapting their strategies to survive. This means that markets aren't always perfectly efficient. Pockets of predictability can appear, offering opportunities for those who are smart enough to find them. But as more people exploit these opportunities, the market adapts, and they fade away. This theory makes a lot more room for the idea that a learning, adaptive system—like a machine learning model—could actually work [63].

2.2 Technical Analysis: The Original Rule Book

Long before machine learning, traders and analysts used technical analysis to try and predict the markets. The core idea is that history tends to repeat itself, and all you need to know is contained in an asset's price chart [58]. Technical analysts use a toolbox of indicators to find trends and signals [1].

- Trend Indicators: Tools like Moving Averages (MA) and the Moving Average Convergence Divergence (MACD) help smooth out price data to see which way the market is generally heading [7, 8, 22].

- **Momentum Oscillators:** Indicators like the Relative Strength Index (RSI) try to identify when a market is "overbought" (and due for a fall) or "oversold" (and due for a rally) [51, 75].

While academics still debate how well technical analysis works on its own [19, 68], it has proven to be a goldmine of structured information. These indicators provide a rich set of features that we can feed into a machine learning model, giving it a powerful head start [72].

2.3 The Machine Learning Revolution

Machine learning has completely changed the game. Instead of programming a computer with a fixed set of "if-then" rules, we give it data and let it learn the rules for itself [11]. For a problem like currency hedging, we typically frame it as a classification task: we ask the model to look at all the data today and classify tomorrow's price movement as either "up" or "down" [10]. The magic of these models is their ability to find incredibly complex and subtle patterns in the data that a human would never see [52].

2.4 What Others Have Found

A lot of research has already shown that ML can be a powerful tool in finance.

- **Ensemble Models:** Studies have found that "ensemble" models like Random Forest, which combine the predictions of many smaller models, are very effective at predicting stock price direction [11, 52].
- **Deep Learning:** More recently, deep learning models like LSTMs have become popular because they are specifically designed to find patterns in sequences of data over time, which is exactly what financial data is [67]. Studies have shown that these models work well for everything from stocks to the Chinese Yuan [49].

2.5 The Missing Piece: From Prediction to Practical Hedging

Here's the gap we noticed: most research focuses on using ML to make speculative trades—to "beat the market." Very few studies have focused on the specific, practical problem of corporate hedging, where the goal isn't to make a killing, but to avoid a loss [70]. This requires more than just a prediction; it requires a complete decision-making framework that weighs the probability of a loss against the cost of preventing it. This paper is our attempt to build and test exactly that.

3. BUILDING THE BRAIN: OUR PREDICTIVE MODEL

A smart expert system needs a smart brain. This section walks through how we built the predictive core of our system, from preparing the raw data to training and

testing our machine learning models.

3.1 Getting and Cleaning the Data

3.1.1 The Raw Material

Our starting point was 20 years of daily price data for the Euro vs. the US Dollar (EUR/USD), sourced from the highly respected Eikon Datastream service. This long history, covering 5,044 days of trading up to June 30, 2022, is essential. It ensures our model learns from all kinds of market conditions—calm periods, volatile crises, strong trends, and sideways grinds—making it more likely to be robust in the future.

3.1.2 Smoothing Out the Noise

Daily financial data can be very "noisy," with lots of small, random fluctuations. To see the real underlying trend, we first needed to smooth the data. We tested two popular methods:

- **Simple Moving Average (SMA):** This is a straightforward average of the price over the last few days (e.g., the last 5 days). It's simple and effective.
- **Exponential Moving Average (EMA):** This is a slightly more sophisticated average that gives more weight to the most recent days, allowing it to react a bit faster to new information.

3.1.3 Defining the Goal

We set up our problem as a simple binary choice. For any given day, we want our model to predict: will the exchange rate trend be up (+1) or down (-1) over the next D days? We trained different models to make this prediction for several time horizons: 1, 2, 3, 5, and even 14 days into the future.

3.2 Feature Engineering: Giving the Model Clues

A machine learning model is only as good as the data you feed it. We couldn't just give it the raw price. Instead, we engineered a set of 53 different "features" or clues based on well-known technical indicators [29]. Think of these as the different gauges and dials an expert trader would look at. They included:

- **Momentum Indicators (RSI, Stochastic Oscillator, etc.):** These tell us how fast the price is moving and whether a trend might be running out of steam.
- **Trend Indicators (MACD, Moving Averages):** These help confirm the direction the market is heading.
- **Volatility Indicators (Standard Deviation):** These measure how wild or calm the market has been recently.

- **Basic Price Data:** We also included simple but useful information, like yesterday's price and how much the price changed from the day before.

By creating multiple versions of these indicators (e.g., a 5-day moving average vs. a 25-day moving average), we gave the model a rich, multi-dimensional view of the market.

3.3 Choosing Our Champions: The Machine Learning Models

We selected three powerful, well-regarded models to compete for the job of being our system's brain.

3.3.1 The Wisdom of the Crowd: Random Forest (RF)

A Random Forest works like a committee of experts [14]. It builds hundreds of individual "decision trees," each of which gets to vote on the prediction. To make sure the committee has diverse opinions, each tree is only shown a random sample of the data and a random selection of the features. The final prediction is simply the one that gets the most votes. This approach is incredibly robust and very good at avoiding a common pitfall called "overfitting" [28].

3.3.2 The Master and Apprentice: XGBoost

eXtreme Gradient Boosting (XGBoost) is another team-based approach, but it works differently [18]. Instead of a committee of equals, it's more like a master and a series of apprentices. The first tree makes a prediction. The second tree then comes in and focuses specifically on correcting the mistakes the first tree made. The third tree corrects the mistakes of the second, and so on. This sequential learning process makes XGBoost extremely powerful and accurate.

3.3.3 The Memory Man: Long Short-Term Memory (LSTM)

LSTMs are a special type of neural network designed for data that comes in a sequence, like a time series [34]. Their key innovation is a kind of "memory cell" that allows them to remember important information from far back in the past. This is incredibly useful for financial markets, where an event from weeks ago might still be influencing today's price. Special "gates" control what information is stored, what is forgotten, and what is used to make a prediction, making LSTMs a go-to choice for time series analysis [27, 34].

3.4 Fair and Rigorous Training

To find the best model and make sure our results were trustworthy, we followed a strict training and validation process.

3.4.1 Finding the Sweet Spot: Hyperparameter Tuning

Each model has a number of knobs and dials, called hyperparameters, that need to be set just right. We used automated tools (GridSearchCV and Keras Tuner) to systematically test thousands of combinations of these settings to find the one that produced the most accurate model.

3.4.2 Avoiding Cheating: Time-Series Cross-Validation

It's easy to build a model that looks great on paper because it has simply memorized the historical data. To avoid this, we used a special type of validation designed for time series data [65]. It works by always training the model on past data and testing it on future data, perfectly mimicking how it would be used in the real world. This gives us a much more honest assessment of how the model will likely perform on new, unseen data.

3.4.3 Less is More: Feature Selection

Having 53 features is great, but some of them might be redundant or just noise. We used a technique called Recursive Feature Elimination (RFE) to see if a smaller, more focused set of features could do a better job [29]. The process works by building a model, ranking the importance of all the features, and then systematically throwing out the weakest ones and rebuilding the model until only the most powerful features are left.

4. THE EXPERT IN ACTION: DESIGNING THE DECISION SYSTEM

Having a model that can predict the market is one thing. Turning that prediction into a smart, real-world hedging decision is the next, crucial step. This section explains how we designed the full expert system and the simulation framework we used to test it.

4.1 The Big Idea: From a Forecast to a Recommendation

Our expert system is designed to give a simple, daily recommendation: "hedge this transaction now" or "wait." It does this by constantly comparing two strategies in the background: a traditional, "safe" strategy and our dynamic, machine-learning-powered strategy. The ultimate test is to see if the ML-driven approach can save the company money over the long run compared to the simple, default option.

4.2 The Baseline: The "Always Hedge" Benchmark Strategy

To have a fair comparison, we needed a baseline. Our benchmark strategy is what a very risk-averse company might do: hedge every single international transaction on the day the deal is made. This completely eliminates

currency risk. The downside? You pay the bank's commission on every single deal, which can be expensive. This strategy is all about certainty, at a known cost.

4.3 The Challenger: The "Smart" Agent Strategy

Our agent strategy is where the machine learning comes in. It's dynamic and intelligent. Here's how it thinks:

1. A new deal comes in, with payment due in, say, 60 days. The agent doesn't immediately hedge.
2. Instead, every single day, it runs our predictive model to get a forecast for the next day's trend.
3. It then checks this forecast against a pre-set Insurance Decision Rule. For example, the rule might be: "If the model predicts the currency will strengthen against us for five days in a row, it's time to hedge."
4. If that rule is triggered on any day before the payment is due, the agent executes the hedge on that day.
5. If the rule is never triggered, the agent lets the transaction ride and simply accepts the market exchange rate on the final payment day.

This strategy aims to be the best of both worlds: it hedges to protect against strong, clear risks, but it avoids paying unnecessary commissions when the market is calm or moving in a favorable direction.

4.4 The Proving Ground: Our Backtesting Framework

To see which strategy was better, we built a virtual testing environment and ran thousands of simulated transactions through it using real historical data from early 2020 to mid-2022.

4.4.1 The Scenarios

We ran two main sets of simulations to make sure our results were robust:

- Simulation Set 1 (The Busy Company): This scenario assumed a new transaction happened every single day, testing how the system performs under a constant, high-volume workload.
- Simulation Set 2 (The Irregular Business): This scenario assumed transactions happened at random times with random payment dates, testing the system's adaptability to more sporadic business flows.

4.4.2 The Rules of the Game

Within each simulation, we tested a wide variety of settings to find the optimal strategy:

- Payment Due Date: We tested deals with payment due in 14, 30, 45, and 60 days.
- Hedging Cost (Commission): We tested a range of commission fees, from 0% up to 0.1%.
- Decision Rules: We tested two main types of rules: one based on the total number of "risk" days predicted, and another based on the consecutive number of "risk" days.
- Degree of Freedom: This is the threshold for our rules. For example, should the agent hedge after 2 consecutive risk signals, or should it wait for a stronger signal of 5 or 7?

4.5 The Scorecard: How We Measured Success

Our main goal was to see if the agent strategy could save money compared to the benchmark. We measured this using a metric called Mean Agent Excess Performance. In simple terms, for every transaction, we calculated how much better (or worse) the agent's final exchange rate was compared to the benchmark's. We then averaged this difference across all the transactions in a simulation. A positive number means that, on average, the smart agent strategy was more profitable.

5. THE VERDICT: RESULTS AND ANALYSIS

After building our models and running them through thousands of simulations, it was time to see how they performed. This section breaks down the results, first looking at the accuracy of the predictive models themselves, and then at the economic performance of the complete expert system.

5.1 How Good Was the Brain? Predictive Model Performance

We trained and tested 675 different models to find the best possible predictive engine for our system.

5.1.1 The Winning Models: RF and XGBoost Take the Lead

The results clearly showed that the ensemble models—Random Forest (RF) and XGBoost—were the top performers. They consistently beat the LSTM models in predicting the direction of the EUR/USD. This suggests that for this type of problem, which relies on a structured set of technical indicators, the tree-based logic of RF and XGBoost is incredibly effective. The best models were able to correctly predict the next day's market direction about 78-80% of the time.

5.1.2 Key Ingredients for a Good Forecast

Our analysis uncovered a few key lessons about what makes a good model:

- **A Little Smoothing Goes a Long Way:** Using a 5-day Simple Moving Average (SMA) to smooth the data before feeding it to the model consistently produced the best results. It seems to be the sweet spot for filtering out market noise without making the data too slow to react.
- **Think Short-Term:** The models were much more accurate when predicting just 1 or 3 days into the future. This wasn't a surprise; it confirms the old saying that the longer the forecast, the more likely it is to be wrong.
- **Quality over Quantity:** Using a smart feature selection technique (RFE) to pick only the 5 or 15 most powerful indicators often worked just as well as, or even better than, using all 53. This shows that focusing on the most important information and reducing noise is critical.

5.1.3 A Close-Up on Our Champion Model

The single best model, which we chose to be the brain of our expert system, was a Random Forest classifier that used a 5-day SMA and focused on just the top 5 most important features to predict the next day's trend.

Here are its stats:

- **Test Accuracy:** 79.17%. It got the direction right nearly 4 out of 5 times on data it had never seen before.
- **Recall:** 81.29%. This was a crucial metric for us. It means that when the market was actually about to move against the company (an "up-trend" in our case), the model correctly raised a red flag over 81% of the time.
- **The Confusion Matrix:** This detailed breakdown showed the model was very good at its primary job of avoiding risk. Its most common error was predicting risk when there was none (a false positive), which is far preferable to its least common error: failing to predict risk when it was actually present (a false negative).
- **AUC Score:** The model achieved an AUC score of 0.82. In the world of classification models, a score above 0.80 is generally considered excellent.

5.2 Did the System Make Money? Backtesting Results

With a powerful predictive brain in place, we ran our simulations to see if the expert system could translate its intelligence into real-world savings.

5.2.1 The Big Picture: Agent vs. Benchmark

When we plotted the results of thousands of individual transactions, the visual evidence was compelling. The vast majority of the dots on our charts fell in the agent's favor, showing that the smart strategy led to a better financial outcome than the "always hedge" benchmark

most of the time.

5.2.2 Results from Simulation Set 1 (The Busy Company)

In our high-volume simulation, the agent strategy proved its worth.

- The average outperformance (Mean Agent Excess Performance) across all 128 simulation runs was 0.59%.
- The agent strategy beat the benchmark in 59% of the simulations.
- In the very best simulation, the agent strategy saved the company an impressive 7.65% compared to the benchmark. This happened on short-term deals where the cost of hedging was high and the agent used a "patient" decision rule.

5.2.3 Results from Simulation Set 2 (The Irregular Business)

The agent's performance was even more consistent when dealing with sporadic transactions.

- The average outperformance was a very similar 0.57%.
- The agent beat the benchmark in an even higher proportion of the simulations: 72%.
- The top-performing simulation here delivered an outperformance of 4.96%.

5.2.4 The Winning Formula for the Agent

Looking across all the results, a clear pattern emerged for the best way to configure the agent's strategy:

- **Hedge When It's Expensive:** The agent's advantage was greatest when the bank's commission fees were high. This makes perfect sense—the more expensive it is to hedge, the more you gain by intelligently choosing not to.
- **Be Patient:** The best results came from rules that required a strong, confirmed signal before hedging (e.g., waiting for 5 or 7 consecutive risk signals). This prevents the system from overreacting to minor market jitters.
- **Focus on the Near Future:** The strategy worked best on deals with shorter time horizons (14 or 30 days), which aligns with the fact that the underlying predictive model was most accurate in the short term.

6. SO, WHAT DOES IT ALL MEAN?

6.1 Reading Between the Lines

The results of our study are quite telling. First, the fact that

our machine learning models could predict market direction with nearly 80% accuracy suggests that the currency market isn't perfectly random. It seems to support the Adaptive Markets Hypothesis—the idea that there are temporary patterns and inefficiencies that a smart, adaptive system can learn to exploit [55].

Second, and more importantly, the success of our "agent" in the backtesting simulations shows that this predictive power can be turned into real economic value. This wasn't just an interesting academic exercise; the system consistently made decisions that would have saved a real company real money. An average outperformance of around 0.60% might not sound like much, but for a large corporation handling billions in foreign currency, that translates into millions of dollars in savings every year.

6.2 A New Toolkit for Financial Managers

For the people on the front lines—corporate treasurers and financial managers—this research offers some practical takeaways:

- **A Powerful Assistant:** This kind of expert system isn't meant to replace human experts, but to augment them. It can act as a tireless, data-driven analyst, providing objective recommendations and freeing up managers to focus on bigger strategic issues.
- **Rethinking Static Rules:** The results strongly suggest that rigid, "one-size-fits-all" hedging policies are outdated. A dynamic approach that adapts to market conditions and the changing cost of hedging is likely to be far more effective.
- **A Smarter Cost-Benefit Analysis:** The system provides a new way to think about the hedging decision. It's no longer a simple "yes" or "no." Instead, it becomes a calculated risk based on the model's predicted probability of an adverse move versus the known cost of the insurance.

6.3 The Real-World Hurdles

Of course, taking a system like this from a research paper to a live corporate environment comes with its own set of challenges.

- **Plugging It In:** The system needs to be seamlessly integrated with a company's existing financial software (like a Treasury Management System) and real-time data feeds. This is a significant technical project.
- **Scaling Up:** Our study focused on one currency pair. A real multinational company might need models for dozens of currencies, all running in parallel. The system needs to be built to handle that kind of scale.
- **Keeping the Brain Sharp:** Financial markets

change. A model trained on last year's data might not be effective next year. A successful deployment requires an ongoing process of monitoring the model's performance, retraining it with new data, and fine-tuning it to keep it accurate.

- **The "Black Box" Problem:** One of the biggest hurdles for adopting complex AI is that it can be hard to understand why it's making a certain recommendation. Building trust in the system means finding ways to make its decision-making process more transparent, a field known as Explainable AI (XAI).

6.4 Where We Could Go From Here

This study is just one step. There are many exciting directions for future research:

- **More and Better Data:** We could make the models even smarter by feeding them more than just price data. What if they could also read macroeconomic news, understand central bank announcements, or even gauge market sentiment from social media [17]?
- **Smarter Model Architectures:** The field of AI is moving incredibly fast. Exploring newer models, like the Transformer architectures that power many large language models, could lead to even better performance.
- **Learning a "Policy," Not Just a Prediction:** Instead of just predicting the next day, we could use a technique called Reinforcement Learning to train an agent to learn an entire long-term hedging policy, making a series of decisions to achieve the best outcome over weeks or months.
- **Portfolio-Level Hedging:** The ultimate goal would be to move from hedging single transactions to optimizing the hedging strategy for a company's entire portfolio of currency exposures at once.

7. THE BOTTOM LINE

Managing foreign exchange risk is a tough, unavoidable part of doing business globally. In this paper, we've explored a new and powerful way to tackle this challenge: using expert systems driven by machine learning. We've walked through the entire process, from building a highly accurate predictive model to designing a complete decision-making framework that translates those predictions into smart, cost-effective actions.

The results were clear and consistent: a dynamic strategy, guided by a machine learning model, can significantly outperform a traditional, static approach to hedging. By intelligently choosing when to act, a company can dramatically reduce the costs of its risk management program.

In the end, the most effective approach will be a partnership between human and machine. These expert systems are not a replacement for the wisdom and strategic oversight of experienced financial professionals. They are, however, a powerful new tool—an intelligent assistant that can provide objective, data-driven insights to help businesses navigate the complex and ever-changing world of global currency markets.

REFERENCES

1. Achelis, S. B. (2000). *Technical analysis from A to Z*. McGraw-Hill.
2. Al Mansoori, G. H., Al Sharafi, E. A., Al Mheiri, S. M., & Nobanee, H. (2024). Foreign exchange risk management: A comprehensive literature review and mitigation strategies in floating exchange rate regimes. *SSRN Electronic Journal*. DOI: 10.2139/ssrn.4979859
3. Alaminos, D., Becerra-Vicario, R., Fernández-Gámez, M. Á., & Cisneros Ruiz, A. J. (2019). Currency crises prediction using deep neural decision trees. *Applied Sciences (Basel, Switzerland)*, 9(23), 5227. DOI: 10.3390/app9235227
4. Allayannis, G., Ihrig, J., & Weston, J. P. (2001). Exchange-rate hedging: Financial versus operational strategies. *The American Economic Review*, 91(2), 391–395. DOI: 10.1257/aer.91.2.391
5. Anghel, D.-G. (2021). A reality check on trading rule performance in the cryptocurrency market: Machine learning vs. technical analysis. *Finance Research Letters*, 39, 101655. DOI: 10.1016/j.frl.2020.101655
6. Anastasakis, L., & Mort, N. (2009). Exchange rate forecasting using a combined parametric and nonparametric self-organising modelling approach. *Expert Systems with Applications*, 36(10), 12001–12011. DOI: 10.1016/j.eswa.2009.03.057
7. Appel, G. (1985). The moving average convergence-divergence trading method: advanced version. *Scientific Investment Systems*.
8. Appel, G., & Dobson, G. (2007). *Understanding MACD*. Traders Press, Inc.
9. Bartram, S. M., & Bodnar, G. M. (2009). The foreign exchange exposure puzzle. *Journal of International Money and Finance*, 28(3), 371–384. DOI: 10.1016/j.jimonfin.2009.03.003
10. Basak, S., Kar, S., Saha, S., Khaidem, L., & Dey, S. R. (2019). Predicting the direction of stock market prices using tree-based classifiers. *The North American Journal of Economics and Finance*, 47, 552–567. DOI: 10.1016/j.najef.2018.06.013
11. Ballings, M., Van den Poel, D., Hespeels, N., & Gryp, R. (2015). Evaluating multiple classifiers for stock price direction prediction. *Expert Systems with Applications*, 42(20), 7046–7056. DOI: 10.1016/j.eswa.2015.05.013
12. Becker, C., Debelle, G., & Fabbro, D. (2005, Dec). Australia's foreign currency exposure and hedging practices. *Reserve Bank of Australia Bulletin*. <https://www.rba.gov.au/publications/bulletin/2005/dec/1.html>
13. Borowski, K., & Pruchnicka-Grabias, I. (2019). Optimal lengths of moving averages for the MACD oscillator for companies listed on the Warsaw Stock Exchange. *Bank & Credit*, 50(5), 457–478.
14. Breiman, L. (2001). Random forests. *Machine Learning*, 45(1), 5–32. DOI: 10.1023/A:1010933404324
15. Chaboud, A. P., Chiquone, B., Hjalmarsson, E., & Vega, C. (2014). Rise of the machines: Algorithmic trading in the foreign exchange market. *The Journal of Finance*, 69(5), 2045–2084. DOI: 10.1111/jofi.12186
16. Chang, P.-C., Fan, C.-Y., & Lin, J.-L. (2011). Trend discovery in financial time series data using a case based fuzzy decision tree. *Expert Systems with Applications*, 38(5), 6070–6080. DOI: 10.1016/j.eswa.2010.11.006
17. Chen, W., Xu, H., Jia, L., & Gao, Y. (2021). Machine learning model for Bitcoin exchange rate prediction using economic and technology determinants. *International Journal of Forecasting*, 37(1), 28–43. DOI: 10.1016/j.ijforecast.2020.02.008
18. Chen, T., & Guestrin, C. (2016). Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining* (pp. 785–794). Association for Computing Machinery. DOI: 10.1145/2939672.2939785
19. Cohen, G., & Cabiri, E. (2015). Can technical oscillators outperform the buy and hold strategy? *Applied Economics*, 47, 3189–3197. DOI: 10.1080/00036846.2015.1013609