Irish Interdisciplinary Journal of Science & Research (IIJSR)
Volume 9, Issue 1, Pages 23-35, January-March 2025

Developing an Intelligent Trading Model for the Ethiopia Commodity Exchange (ECX) Using Deep Reinforcement Learning Algorithms

Yesuf Mussema[1*] & Kuulaa Kakeba (PhD.)[2]

[1,2]*AI and Robotics Center of Excellence, Addis Ababa Science and Technology University, Addis Ababa, Ethiopia.*
Corresponding Author (Yesuf Mussema) Email: yesuf.mussema@aastustudent.edu.et[*]

DOI: https://doi.org/10.46759/IIJSR.2025.9103

Article Received: 09 December 2024 *Article Accepted: 16 February 2025* *Article Published: 23 February 2025*

ABSTRACT

The Ethiopian Commodity Exchange (ECX) faces significant challenges, including manual trading processes, market inefficiencies, and data fragmentation, which hinder its ability to operate effectively in a volatile and dynamic environment. This research develops an intelligent trading model leveraging Deep Reinforcement Learning (DRL) algorithms, specifically Deep Q-Networks (DQN), Double Deep Q-Networks (DDQN), and Advantage Actor-Critic (A2C), to address these issues.

The proposed framework utilizes DRL to enable agents to learn optimal trading policies through interactions with simulated ECX market environments. The model employs historical market data, representing state features such as price trends, trading volumes, and external economic indicators. Actions are defined as buy, sell, or hold decisions, while reward structures are designed to incentivize profit and penalize excessive risk. The research integrates techniques such as experience replay and target networks in DQN, action evaluation in DDQN, and advantage functions in A2C to enhance model performance and stability.

Experimental results demonstrate that the DRL models significantly improve trading efficiency and decision-making accuracy compared to manual processes. DDQN outperforms DQN in managing noisy and volatile market conditions, while A2C excels in handling continuous decision variables, such as dynamic trade volumes. The results highlight the scalability and adaptability of the proposed system in addressing ECX-specific challenges, including risk management and market transparency.

The study concludes that the DRL-based trading model offers transformative potential for the ECX by automating decision-making, optimizing trade execution, and promoting equitable participation among stakeholders. This research provides a foundation for integrating advanced machine learning techniques into emerging commodity markets, ensuring their efficiency and competitiveness in a global context.

Keywords: Ethiopian Commodity Exchange (ECX); Deep Reinforcement Learning (DRL); Double Deep Q-Networks (DDQN); Advantage Actor-Critic (A2C); Weighted Average; Commodity Trading; Intelligent Trading Model; Automated Decision-Making; Risk Management; Performance Metrics; High-Frequency Trading.

1. Introduction

The Ethiopian Commodity Exchange (ECX) was established in 2008 as a privately owned commodity exchange (Ethiopian Commodity Exchange Authority, 2015). It plays a crucial role in Ethiopia's economy by transforming the traditional agricultural market through a fair, transparent, and efficient platform for buying and selling commodities (Gabre-Madhin E., 2008). Guided by the Ethiopia Commodity Exchange Proclamation No. 550/2007, the ECX was established with the primary objective of fostering a just and impartial marketplace for agricultural commodities (Gabre-Madhin E., 2008).

ECX Initially trading in six commodities, it now trades in over 20, including coffee, oilseeds, pulses, cereals, and spices (Ethiopian Commodity Exchange, 2021). Despite its achievements in improving market access and reducing transaction costs, the ECX faces significant challenges such as manual trading processes, price volatility, limited data utilization, and operational inefficiencies.

This research introduces a cutting-edge solution by leveraging Deep Reinforcement Learning (DRL) algorithms, namely Deep Q-Networks (DQN), Double Deep Q-Networks (DDQN), and Advantage Actor-Critic (A2C). These algorithms aim to automate decision-making, optimize trading strategies, and enhance market transparency and efficiency. The study offers a framework for transforming Ethiopia's commodity trading system into a modern, data-driven marketplace while addressing its unique challenges.

1.1. Study Objectives

The main objective of this study is to develop a DRL-based intelligent trading model that addresses ECX's inefficiencies, enhances trading efficiency, and improves decision-making through advanced machine learning. The specific objectives of this study are as follows:

1. Design a DRL framework tailored to ECX, incorporating metrics like Sharpe Ratio, Expected Return, and Maximum Drawdown.

2. Implement and train DQN, DDQN, and A2C algorithms for optimal trading strategy development.

3. Evaluate the performance of DRL algorithms in a simulated ECX environment using quantitative financial metrics.

4. Incorporate risk management and regulatory compliance mechanisms to ensure system stability and sustainability.

2. Literature Review

2.1. Overview of Algorithmic Trading Systems and Platforms

Algorithmic trading has revolutionized commodity market trading by automating trading decisions based on predefined rules and data analysis (Moody & Saffell, 1998). These systems leverage advanced machine learning algorithms and quantitative methods to enhance efficiency and reduce human error in decision-making processes (Moody & Saffell, 1998). Platforms such as X_TRADER, CQG Integrated Client, and MetaTrader 5 exemplify the diversity of tools available for automating trade execution and strategy development. X_TRADER, developed by Trading Technologies International, is widely recognized for its Autospreader® feature, which enables real-time creation and execution of complex spread strategies (Trading Technologies International Inc., n.d.). Similarly, CQG's Integrated Client provides robust analytics and an algorithmic order-routing system tailored for commodity and futures trading (CQG Inc., [Accessed: Dec. 16, 2024]). MetaTrader 5 further enhances algorithmic trading by enabling the development and deployment of custom-built Expert Advisors (EAs), particularly in commodities such as oil and gold (MetaQuotes Software Corp., n.d.).

This thesis examines how DRL-based algorithms, including DQN, DDQN, and Actor-Critic methods, can optimize trading strategies for the ECX by improving metrics like Sharpe Ratio, Expected Return, and Maximum Drawdown. By analyzing these algorithms' adaptability and scalability, the research highlights their potential to transform ECX's trading system.

2.2. Evolution of Algorithmic Trading

The transition from manual trading to electronic and algorithmic trading has been marked by significant technological advancements (Gibbons, 2021). Early innovations, such as the adoption of the Financial Information eXchange (FIX) protocol in 1992, laid the groundwork for real-time data sharing and automated trade execution

(FIX Trading Community, 2024). By the 2000s, the development of high-frequency trading (HFT) algorithms and the proliferation of alternative trading systems (ATS) further reduced reliance on traditional floor traders (Biais et al., 2015).

2.3. Algorithmic Trading and DRL

While traditional algorithmic trading systems relied on statistical and econometric models, their inability to adapt to the dynamic and non-linear nature of financial markets has led to the integration of Deep Reinforcement Learning (DRL) (Moody & Saffell, 1998). DRL enhances trading systems by enabling agents to learn optimal policies directly from market data (Yang et al., 2020). Unlike static models, DRL-based systems adapt to changing market conditions and continuously improve their decision-making capabilities, offering significant advantages in highly volatile commodity markets (Moody & Saffell, 1998; Sutton & Barto, 2018).

2.4. Factors Influencing Commodity Trading

Commodity trading is influenced by global supply-demand imbalances, geopolitical tensions, and macroeconomic indicators (Miller & Black, 2018). For instance, disruptions in East Africa’s coffee supply caused by climate change have led to significant price fluctuations, demonstrating the market's vulnerability to external factors (Worku, 2014). Additionally, poor storage and transportation networks exacerbate inefficiencies, while behavioral challenges, such as limited trust among participants, hinder transparency and reliability (Worku, 2014; Sexton, 1988). The research addresses these challenges by investigating how DRL models can dynamically adapt to external factors, optimize risk management strategies, and enhance performance metrics such as Sortino and Calmar Ratios.

2.5. Applications of Deep Reinforcement Learning (DRL)

Deep Reinforcement Learning (DRL) combines the representational power of deep learning with the decision-making capabilities of reinforcement learning, making it an ideal approach for solving complex financial market challenges. DRL has been successfully applied to optimize trading execution, reduce risk exposure, and improve profitability (Miller & Black, 2018). For example, studies have demonstrated DRL’s effectiveness in navigating volatile commodity markets, leveraging sentiment analysis, and outperforming traditional methods in metrics such as the Sharpe Ratio and Maximum Drawdown (Moody & Saffell, 1998; Zhang et al, 2020).

This research focuses on how DRL can be tailored to the ECX's unique trading environment, exploring how algorithms like Weighted Average and A2C can be fine-tuned to maximize profitability while ensuring compliance and stability (Yang et al., 2020).

2.6. Key DRL Techniques in Commodity and Financial Markets

2.6.1. Deep Q-Network (DQN)

DQN, introduced by Mnih et al. (2015), is a breakthrough in deep reinforcement learning (DRL) that combines Q-learning with deep neural networks. The algorithm approximates the Q-value function, estimating expected future rewards for each state-action pair (Sutton & Barto, 2018). In trading, DQN helps identify optimal actions (e.g., buy, sell, hold) to maximize cumulative rewards under various market conditions (Yang et al., 2020).

2.6.1.1. Key enhancements in DQN include:

- **Experience Replay**: Stores agent experiences and samples them randomly to stabilize training.
- **Target Network**: Reduces oscillations by providing fixed targets for Q-value updates.

Applications in algorithmic trading, such as portfolio management (Ding et al., 2018), have demonstrated its effectiveness. However, challenges like overestimation bias can limit its reliability in volatile markets (Lillicrap et al., 2015).

2.6.2. Double Deep Q-Network (DDQN)

Van Hasselt et al. (2016) proposed DDQN to address overestimation bias inherent in DQN by decoupling the action selection and evaluation processes. This modification significantly improves the stability and accuracy of Q-value estimation, making DDQN a preferred choice in complex and noisy environments like commodity trading (Zhang et al, 2020).

In the context of financial markets, DDQN has been employed for tasks such as high-frequency trading and risk management. By reducing overestimation, it offers more reliable trading decisions, especially in scenarios where DQN might overfit to short-term market anomalies (Zhang et al, 2020).

2.6.3. Actor-Critic (A2C)

Actor-Critic methods, introduced by Konda and Tsitsiklis (2000), combine the advantages of value-based and policy-based approaches. The actor learns the policy (mapping states to actions), while the critic evaluates the chosen actions using a value function. This dual architecture facilitates stable and efficient learning, especially in continuous action spaces (Konda & Tsitsiklis, 2000).

Variants like Advantage Actor-Critic (A2C) and Proximal Policy Optimization (PPO) have been widely used in trading environments. For instance, Jiang et al. (2017) applied Actor-Critic models to develop a stock trading system that dynamically adjusted trading volumes based on market volatility. The actor-critic framework is particularly advantageous in environments with frequent and subtle market changes, such as agricultural commodity exchanges.

2.6.4. Application of Weighted Average in Algorithmic Trading

The implementation applies the concept of the weighted average to analyze historical trading data, primarily focusing on the computation of a weighted average price for commodities and its influence on determining trading actions: Buy, Sell, and Hold. This process aids in creating an actionable strategy that incorporates price trends and performance metrics to optimize trading decisions.

2.6.4.1. Action Signal Generation

- **Buy Signal**: Triggered when the Closing Price surpasses the 50-period Moving Average (50MA), indicating a positive price momentum.
- **Sell Signal**: Triggered when the Closing Price falls below the 50MA, suggesting a potential decline in price.

Irish Interdisciplinary Journal of Science & Research (IIJSR)
Volume 9, Issue 1, Pages 23-35, January-March 2025

- **Hold Signal**: Default action when neither a Buy nor a Sell signal is generated, implying market stability or uncertainty (Smith, 2018; Johnson A. , 2020).

2.7. Insights and Challenges in ECX

The ECX has made significant strides in modernizing Ethiopia's commodity trading landscape. However, persistent challenges, such as infrastructural limitations, inefficiencies in manual systems, and lack of advanced technological integration, continue to hinder its growth (Amha & Gebremeskel, 2019; Worku, 2014). Algorithmic trading powered by DRL offers a pathway to address these issues by enhancing scalability, democratizing access to market data, and improving decision-making processes (Zhang et al, 2020; Lee, 2021).

3. Methodology

This methodology section provides a structured and comprehensive outline of the research approach, data collection, algorithm implementation, evaluation metrics, and theoretical foundations essential for understanding and conducting the study on deep reinforcement learning in algorithmic trading within the ECX.

3.1. Data Collection and Preprocessing

The research utilizes seven years of historical trading data from the ECX, encompassing commodity prices, trading volumes, and other relevant variables, such as opening and closing prices, maximum and minimum prices, and total trade value. External factors, including economic indicators, are incorporated to enrich the dataset. Data preprocessing includes cleaning to address missing and duplicate values, normalization to scale data uniformly, and feature engineering to capture market dynamics effectively. The dataset is divided into training, validation, and testing subsets, enabling model training, hyperparameter tuning, and performance evaluation.

Table 1. Commodity Exchange Data Attributes

Column	Description	Type	Example
Date	Exact day a commodity transaction occurred; enables temporal trend analysis.	Date	2024-12-08
Commodity	Primary good being traded at ECX.	String	Coffee
Commodity Type	Processing method or classification of the commodity.	String	Washed, Unwashed
Commodity Name	Specific categorization or extended description of the commodity.	String	Export Coffee
Description	Detailed classification, including type, grade, and origin.	String	Unwashed Kelem Welega 7
Origin	Geographical source of the commodity, indicating regional production.	String	Jimma, Goji
Opening	Initial price at session start; reflects market sentiment.	Float	1698.00 (ETB/kg)

Price			
Closing Price	Final price at session close, showing market performance.	Float	1720.00 (ETB/kg)
Max Price	Highest recorded price during a session; indicates peak demand.	Float	1720.00 (ETB/kg)
Min Price	Lowest recorded price during a session; shows market support level.	Float	1698.00 (ETB/kg)
Total Volume	Total quantity of the commodity traded during a session.	Integer	480,000
Kilogram	Standardized unit for measuring trade quantities in kilograms.	Float	10,200.00
Ton	Aggregate trade weight converted to metric tons.	Float	10.2
Total Value	Total monetary value of all trades, calculated as price × volume.	Float	480,000.00 (ETB)
GDP	Gross domestic product during the observed period, providing economic context.	Float	12,044,500.00
Exchange Rate	Value of the Ethiopian Birr (ETB) relative to the US Dollar.	Float	29.07
Inflation Rate	General price increase rate, reflecting macroeconomic trends.	Float	12.86

3.1.1. Data Cleaning

The data cleaning process ensures that the dataset is free from missing values, inconsistencies, and erroneous entries, enabling effective analysis.

Procedure:

- **Missing Data Handling**: The dataset was initially inspected for missing or null values. If any columns contained missing values, they were filled with appropriate replacements (e.g., "Unknown" for categorical fields or zero for numeric fields). In this case, the commodity-related fields were replaced with "Unknown" to ensure no null values remained in the categorical columns, while numerical values were replaced with zero or treated as outliers for further analysis.

- **Currency and Formatting Adjustments**: In the cleaning process, commas in numerical columns (such as "Opening Price" and "Closing Price") were removed, and the columns were converted into float type to standardize the data.

- **Generalized Cleaning**: Certain economic columns like GDP, Exchange Rate, and Inflation Rate were cleaned to remove any non-numeric characters (such as commas) and converted into numerical formats.

3.1.2. Handling Outliers

Outliers in the dataset can distort statistical analyses and model training. Therefore, identifying and handling outliers is crucial for ensuring the quality and accuracy of the analysis.

Procedure:

- While not explicitly demonstrated in the code, outlier detection typically involves analyzing the distribution of numerical features. Common methods include visualizing data distributions (histograms, box plots) and using statistical methods like Z-scores or IQR-based filtering to identify extreme values.

If outliers are detected, they can be treated by either capping, transforming, or removing them, depending on the context and the impact they have on model performance.

3.1.3. Feature Engineering

Objective: Feature engineering helps in extracting useful information from raw data to improve model performance by creating new features or transforming existing ones.

Procedure:

- **Rate of Return**: A new feature, Return, was engineered by calculating the percentage change in the closing price between consecutive days. This helps in understanding the daily price movements and is essential for financial models.
- **Moving Average**: The 7-day moving average of the Closing_Price was added to capture short-term trends and smooth out volatility, providing insights into market direction.
- **Volatility**: The rolling standard deviation (Volatility) of the Closing_Price was calculated over a 7-day window to measure market volatility. This is important for modeling risk or uncertainty in financial markets.

3.1.4. Numerical Feature Scaling

Scaling numerical features ensures that variables with different magnitudes or units do not disproportionately influence the model's performance.

Procedure:

- The StandardScaler from sklearn was applied to scale the numerical columns such as Opening_price, Closing_Price, RSI, and economic indicators. This process standardizes the values, converting them to a distribution with a mean of 0 and a standard deviation of 1. This is important for algorithms sensitive to the scale of data, such as linear regression or neural networks.

3.1.5. Categorical Feature Encoding

Encoding categorical features allows the machine learning model to handle non-numeric data effectively.

Procedure:

- **One-Hot Encoding**: The OneHotEncoder was applied to categorical features like Commodity, Commodity_Type, and Origin. This method creates binary columns for each category in the original categorical columns, which is important for algorithms like decision trees, logistic regression, and neural networks that cannot process categorical data directly.

3.1.6. Final Dataset Preparation

The final dataset preparation consolidates all the preprocessing steps into a form suitable for model training.

Procedure:

- After applying the cleaning, feature engineering, scaling, and encoding processes, the dataset was transformed using a pipeline. The ColumnTransformer allowed categorical and numerical data to be processed separately but integrated into a unified dataset.
- The processed dataset was then stored in a DataFrame, which includes both the transformed numerical values and the one-hot encoded categorical features.

3.1.7. Train-Test Split

Splitting the data into training and testing sets ensures that the model is evaluated on unseen data, providing an unbiased estimate of its performance.

Procedure:

- The dataset was split based on a date cutoff (2021-07-21), ensuring that all data before this date was used for training and data after this date was used for testing. This is a typical approach in time-series forecasting tasks, where future data cannot be used to predict past outcomes.

4. Results and Discussion

4.1. Presentation of Research Findings

This study presents compelling evidence of the effectiveness of DRL algorithms in optimizing trading strategies for the ECX. By simulating market conditions, the performance of DQN, DDQN, and A2C was evaluated using metrics such as cumulative returns, Sharpe Ratios, Sortino Ratios, and Maximum Drawdown. Among the models, DDQN demonstrated the highest adaptability to dynamic market conditions, achieving more stable and profitable trading outcomes. While DQN struggled with consistency, A2C provided moderate improvements but lacked the scalability and robustness required for volatile environments.

4.2. Research Results and Analysis

DQN Networks:

DQN showed inconsistent performance, with substantial variability between training and testing phases. While the Sharpe and Sortino Ratios improved during testing, the model struggled with high Maximum Drawdown values, indicating poor risk management. The expected returns improved but remained volatile, rendering DQN unsuitable for robust trading scenarios.

DDQN Networks:

DDQN emerged as the best-performing model, significantly surpassing DQN and A2C across all performance metrics. It achieved remarkable improvements in cumulative returns (18,219.26) and Sharpe Ratio (11.64), while maintaining excellent risk management as evidenced by a minimal Maximum Drawdown of -66.37. Its stability score indicated consistent performance, making it the most reliable model for ECX trading operations.

Actor-Critic Networks (A2C):

The Actor-Critic algorithm exhibited stable yet modest improvements in testing, with minimal increases in profitability. While it demonstrated adaptability to market dynamics, the model's suboptimal performance in Sharpe and Sortino Ratios and poor drawdown values limited its utility in high-risk environments.

Weighted Average Approach:

As a baseline model, the Weighted Average approach provided balanced but less profitable results. It showed moderate stability and risk management, making it suitable for low-risk scenarios but inadequate for volatile trading environments like the ECX.

4.3. Key Performance Metrics

- **Sharpe Ratio:** DDQN recorded the highest risk-adjusted returns (11.64), demonstrating its ability to maximize profits while minimizing risks.
- **Maximum Drawdown:** DDQN minimized capital losses significantly (-66.37), outperforming all other models.
- **Expected Returns:** DDQN achieved exceptional profitability (18,219.26), while DQN and A2C struggled with consistency.
- **Sortino Ratio:** DDQN showed strong downside risk management (63.00), far exceeding the performance of other models.
- **Stability:** DDQN demonstrated consistent performance, with a stability score of 0.0074, indicating reliable operation under dynamic conditions.

4.4. Discussion and Interpretation of Results

The findings confirm that DDQN is the most effective model for the ECX, excelling in revenue generation, cost efficiency, operational reliability, and risk management. DQN's limitations in handling volatile markets underscore its unsuitability for complex trading scenarios, while A2C displayed moderate improvements but lacked profitability. The Weighted Average model provided a balanced alternative but lagged in adaptability and profitability compared to DDQN.

4.5. Implications of the Findings

The research highlights the transformative potential of DRL in modernizing commodity trading systems. DDQN's exceptional performance suggests its capability to automate trading processes, enhance market transparency, and improve decision-making for ECX stakeholders. Key implications include:

- **Revenue Generation:** DDQN maximizes profitability while mitigating risks, making it ideal for the ECX's trading environment.
- **Cost Efficiency:** The model effectively reduces losses, minimizing operational costs.
- **Market Efficiency:** By automating decision-making, the system fosters responsive and transparent market operations.
- **Economic Benefits:** Lower transaction costs and improved price discovery benefit smallholder farmers and traders, promoting equitable participation.

5. Conclusion

This research demonstrates the potential of DRL-based intelligent trading systems in addressing systemic inefficiencies in the ECX. Among the tested models, DDQN emerged as the most robust, excelling in profitability, risk management, and scalability. The study establishes a framework for integrating AI into Ethiopia's commodity trading system, paving the way for modernized, data-driven market operations aligned with global standards.

5.1. Recommendations

1. **Adoption of DRL Systems:** The ECX should implement DDQN-based trading models to improve decision-making and risk management.
2. **Infrastructure Development:** Investments in real-time data collection and processing systems are critical for scalability.
3. **Policy Support:** Policymakers should incentivize AI-driven innovation in commodity trading.
4. **Capacity Building:** Training programs should equip stakeholders with the knowledge to use AI-driven systems effectively.

5.2. Future Research Directions

Future studies should focus on integrating multi-agent reinforcement learning (MARL) to simulate complex market interactions, real-time implementations of DRL models, and the incorporation of global economic trends to enhance robustness. Additional research into hybrid models combining A2C and DDQN could provide more balanced performance in profitability and risk management.

- Application of DRL in Futures Trading

Hedging Strategies: Using DRL to optimize hedging mechanisms for mitigating risks in volatile markets.

Margin Management: Developing models for dynamic margin management to ensure compliance and optimize returns.

- Multi-Agent Reinforcement Learning (MARL)

Exploration of MARL frameworks where multiple agents interact in simulated futures markets could enhance strategy development by incorporating competition and collaboration dynamics.

- Incorporating Macro-Economic Indicators

Future research should integrate macroeconomic variables, such as interest rates and geopolitical data, to enhance DRL models' predictive accuracy. This could involve incorporating time-series analysis for robust market forecasting.

- Risk Assessment and Regulation

Risk Management Models: Creating systems to monitor exposure in highly leveraged markets.

Algorithmic Auditing Tools: Developing frameworks for ensuring transparency and fairness in DRL-based trading systems.

- Real-Time Data and Scalability

Scalability: Enhancing algorithms for large-scale datasets and real-time performance.

Market Adaptability: Creating models capable of adapting to sudden market shifts.

Declarations

Source of Funding

This research did not receive any grants from public, commercial, or not-for-profit sectors.

Competing Interests Statement

The authors declare no competing financial, professional, or personal interests.

Consent for publication

The authors consent to the publication of this research.

Authors' contributions

Both the authors took part in literature review, analysis, and manuscript writing equally.

References

Amha, K., & Gebremeskel, A. (2019). Commodity Trading and Infrastructure Bottlenecks in Ethiopia. African Economic Review, 27(3): 215–228.

Biais, T.M. (2015). Equilibrium High-Frequency Trading. In Econometrica. https://doi.org/10.3982/ecta10857.

Black, T.M. (2018). Risk Management in Commodity Trading. Commodity Markets Review, 22(4): 123–140.

CQG Inc. ([Accessed: Dec. 16, 2024]). CQG Integrated Client: Advanced Charting and Analytics for Traders ([Online]). Retrieved from https://www.cqg.com.

Ding, W.Z. (2018). Deep Reinforcement Learning for Portfolio Management. Journal of Financial Data Science, 3(2): 45–58.

Ethiopian Commodity Exchange (2021). Annual Report ECX. Addis Ababa: ECX.

Ethiopian commodity market Authority (2015). Commodity market manual and stakeholder role. In ECX, Commodity market manual and stakeholder role, Pages 15–56, Addis Ababa: ECX.

FIX Trading Community (2024, Dec [Accessed: Dec. 16, 2024].). Introduction to the FIX Protocol ([Online]). Retrieved from https://www.fixtrading.org.

Gabre-Madhin, E. (2011). The Transformative Role of Electronic Trading in Ethiopia. African Journal of Economic and Management Studies, 2(1): 28–45.

Gabre-Madhin, E.Z. (2012). Ethiopia Commodity Exchange and the Growth of the Agricultural Sector. Journal of Agricultural & Food Industrial Organization, 10(1): 8–27.

Gibbons, G. (2021). Analyzing Commodity Markets with Historical Data: Applications for Developing Economies Systems. In I.T. Social.

Johnson, A. (2020). Advanced Trading Algorithms. IEEE Journal of Trading Systems, 67(4): 45–56.

Konda, V., & Tsitsiklis, J. (2000). Actor-Critic Algorithms. In Advances in Neural Information Processing Systems (NIPS), Pages 1008–1014.

Lee, H. (2021). Enhancing Agricultural Market Efficiencies with DRL. Journal of Commodity Studies, 8(2): 110–125.

Lillicrap, T.H., et al. (2015). Continuous Control with Deep Reinforcement Learning. Proceedings of the International Conference on Learning Representations (ICLR).

MetaQuotes Software Corp. (n.d.). MetaTrader 5: Multimarket Trading Platform ([Online]). Retrieved from https://www.metatrader5.com.

Miller, T., & Black, A. (2018). Risk Management in Commodity Trading. Commodity Markets Review, 22(4): 123–140.

Moody, J., & Saffell, M. (1998). Reinforcement Learning for Trading. Advances in Neural Information Processing Systems, 11(11): 917–923.

Sexton, R. (1988). Price Information and Trade Efficiency in Markets. American Economic Review, 78(5): 1195–1201.

Smith, J. (2018). Introduction to Financial Markets. IEEE Transactions on Finance, 54(3): 123–130.

Sutton, R., & Barto, A. (2018). Reinforcement Learning: An Introduction (2nd Eds.). Cambridge MA, USA: MIT Press.

Trading Technologies International Inc. (n.d.). "X_TRADER: Advanced Trading Tools and Solutions," ([Online]). Retrieved from https://www.tradingtechnologies.com.

Worku, A. (2014). Assessing the Opportunity and Challenges of Ethiopia Commodity Exchange for the Members of Agricultural Product Export. International Journal of Economics & Management Sciences, 3(3): 1–12.

Yang, Z.W. (2020). Real-Time Optimization of Commodity Trading Strategies Using Deep Reinforcement Learning. Journal of Finance and Data Science, 6(2): 124–134.

Zhang, Y.W. (2020). A DDQN-Based Intelligent Decision-Making Framework for Multi-Asset Trading. IEEE International Conference on Autonomic and Secure Computing (DASC), Pages 177–183.