International Journal of Scientific Research in Engineering and Management (IJSREM)
Volume: 07 Issue: 10 | October - 2023 SJIF Rating: 8.176 ISSN: 2582-3930

Enhancing Portfolio Optimisation with Deep Reinforcement Learning: A Comparative Analysis

J Akshaya

Computer Science Engineering with Big Data Visualisation, SRM Institute of Science and Technology, Vadapalani

Abstract:

The aim of this paper is to understand and analyse how Portfolio Optimisation is enhanced while utilising Deep Reinforcement Learning (DRL) versus while using traditional optimisation methods. One of the objectives is also to focus on how DRL algorithms can navigate through the limitations of the traditional portfolio optimisation approach. The comparing of the two models will help gain a better understanding of the performance of the two models. This would also lead us to better insights on the potential advantages and limitations of DRL while it returns risk-adjusted returns for the portfolio.

1. Introduction:

Portfolio optimisation is a crucial milestone in making investment decisions and risk management. It is a perpetual challenge to achieve balance between maximising returns and minimising risks in a portfolio. In volatile and complex environments such as the financial market, optimisation methods such as Mean-Variance Optimisation and traditional methods often struggle to provide robust solutions. It is within this dynamic context that we turn to the potential of cutting-edge advancements such as Deep Reinforcement Learning (DRL). This research is to explore and understand the adaptability and capacity of DRL, to comprehend intricate patterns that would revolutionise portfolio optimisation.

The aim is to conduct a comparative analysis, where the traditional methods are weighed against DRL, to analyse the limitations and advantages. One of the focuses of this study would be to unveil the strengths of DRL in addressing complexities of nonlinear relationships and volatile market conditions whilst providing a deeper understanding of its application in achieving better results for risk-adjusted returns for the portfolio investments.

2. Literature Review:

2.1 Traditional Portfolio optimisation methods:

Traditional theory is based off of attaining maximum returns while minimising risks. Mean-Variance Optimisation(MVO) is a foundational method in traditional portfolio theory, brought about by an economist Harry Markowitz. MVO is based on simple yet powerful idea that investors should target to construct portfolios that maximise anticipated returns while simultaneously minimising portfolio volatility. There are a few core principles:

1. Expected Returns: MVO begins by estimating the expected returns of an individual assets within the space. These returns are typically calculated from historical data or expert projections.
2. Risk Assessment: The method then quantifies the risk associated with each of the asset by calculating variances and covariances. It measures how the returns of one asset moves in relation to the returns of another.
3. Efficient Frontier: MVO employs these expected returns and risk measures to create an efficient frontier. The efficient frontier represents a se too portfolios that offer the maximum expected return for a given level of volatility. The frontier provides a range of options, each with a unique risk-return profile.

MVO provides a valuable framework for constructing diversifies portfolios, and it remains a central concept in traditional portfolio theory as well. However it is also essential to beware of its limitations. One of the primary challenges associated with traditional portfolio methods is their sensitivity to input parameters. Small changes to the input data, such as expected returns and covariances, can lead to significantly different portfolio allocations, making them susceptible to estimation errors. Moreover MVO depends on the assumption that assets and covariances follow Gaussian distribution. This assumption falls short of taking into consideration the complex and nonlinear nature of the financial market, where sudden shifts and non-Gaussian behaviour are common.

2.2 Introduction of Deep Reinforcement Learning:

As a response to the limitations of the traditional portfolio optimisation methods, DRL is a suitable alternative in the subfield of artificial intelligence that leverages neural networks and reinforcement learning techniques. Neural networks acknowledged for their capacity to capture intricate patterns and feature representations. These multi-layered networks allow DRL algorithms to model complex relationships within the data. Reinforcement learning on the other hand, is an acclaimed framework for sequential decision-making. Agents in a reinforcement learning system learn through interactions with an environment, receiving rewards and penalties for their actions.

In this particular scenario, an RL agent can select to hold, buy, or sell a share, and to guarantee that the RL model is working optimally, it is assessed using market benchmark standards.

3. Deep Reinforcement Learning for Portfolio Optimisation:

DRL holds a significance for enhancing portfolio optimisation by introducing a dynamic, data-driven approach to decision making. To explain the adaption of DRL to portfolio optimisation, we can use the NIFTY50 index and its corresponding companies as a reference, where we can visualise a framework where the portfolio allocation become an action taken by the DRL agent.

1. Environment Setup:

The foremost step is to define the environment in which the DRL agent would be operating in. In this context, the environment is going to represent the financial market comprising of the NIFTY50 index and its constituent companies. The environment includes historical price data for these assets, relevant risk factors, and other parameters.

2. State representation:

The state in the DRL framework represents the current market conditions, which include the historical performances of the NIFTY50 index, and the individual companies. The state should encompass key financial indicators, such as historical returns data, volatility, correlations, and possible external factors such as economic indicators and news sentiment.

3. Action Space:

The action space defines the choices available to the DRL agent. In the context of portfolio optimisation, the action would correspond to portfolio allocations. The agent would decide how to distribute capital among the NIFTY50 index and its constituents. These allocations can be binary (invest or not) or continuous (proportion of capital allocation to each asset).

4. Reward Mechanism:

The reward mechanism is crucial in RL, as it guides the agent's learning process. In portfolio optimisation, the reward function evaluates the performance of the portfolio allocation chosen by the agent. The reward function can be a function of risk-adjusted returns, where a higher return with lower risk is the result. For example, Sharpe ratio can be used as the reward metric, with the agent learning to maximise it.

5. Q-Network architecture:

The Q-Network, a fundamental building brick of DRL, is responsible for estimating the expected cumulative rewards for each possible action in a given state. Here, Q-Network should be capable of handling the complexity of allocating capital across multiple assets.

6. Exploration vs. Exploitation:

Balancing exploration (trying different portfolio allocations) and exploitation (selecting the best-known allocation) is vital. The agent must explore various allocation strategies to discover optimal solutions while exploiting its existing knowledge. Techniques like epsilon-greedy exploration can be employed.

7. Training and Learning:

The DRL agent learns by interacting with the environment over multiple episodes. During training, it explores various portfolio allocations, receives rewards, and updates its Q-network to improve decision-making. The agent's objective is to learn an optimal policy for portfolio allocation that maximises the expected cumulative rewards over time.

8. Evaluation and Testing:

After training, the DRL agent's performance is evaluated on out-of-sample data to assess its ability to generalise to unseen market conditions. The agent's portfolio optimisation decisions are compared with traditional approaches to quantify its performance.

By adapting DRL to portfolio optimisation in the context of the NIFTY50 index and its constituent companies, the DRL agent becomes a dynamic, data-driven portfolio manager. This approach offers the potential to outperform traditional methods by capitalising on the adaptability and predictive power of deep reinforcement learning.

4. Data and Methodology:

In the pursuit of enhancing portfolio optimisation through DRL, meticulous attention to data and methodology is pivotal. The success of the comparative analysis hinges on the judicious selection of datasets and the deployment of robust DRL algorithms designed to adapt and excel in the volatile financial market.

4.1 Data Selection:

Historical Asset Price Data:

Historical asset price data holds a position of paramount importance. The NIFTY50 index, comprises of fifty of India's leading companies, this serves as a primary reference universe. The historical daily closing prices of these constituent companies have been meticulously curated, spanning a substantial period to encapsulate diverse market conditions. This comprehensive dataset enables us to construct an intricate understanding of the performance of each asset and the broader index.

Risk Factors:

In search of comprehensive portfolio optimisation, we must extend our dataset to encompass a plethora of risk factors. These risk factors encompasses market indicators such as interest rates, inflation, and economic growth, amongst other. By integrating these risk factors into our analysis, we recognise the systemic influences that are underlying to market movements, contributing to a more robust portfolio optimisation model.

Additional Data Sources:

In tandem with price and risk factor data, our analysis benefits from additional sources. These sources extend to encompass sentiment data, harnessing the collective sentiment expressed in financial news and social media. Furthermore, economic indicators play a critical role, serving as leading indicators of market trends. The symbiosis of these datasets enriches our analysis with a holistic perspective, fostering a more encompassing understanding of the financial landscape.

4.2. DRL Algorithm Selection:

The choice of DRL algorithms, foundational to our portfolio optimisation paradigm, reflects the crux of our methodology. Recognising the intricate nature of financial markets, we have selected algorithms that manifest adaptability, agility, and predictive prowess.

Deep Q-Networks (DQN):

DQN, an embodiment of value-based reinforcement learning, features prominently in our methodological arsenal. DQN excels in its ability to navigate complex action spaces, making it eminently suited for the portfolio optimisation task. With deep neural networks at its core, DQN facilitates the mapping of states to action values, encapsulating the expected cumulative rewards of each potential portfolio allocation. By leveraging DQN, we empower our agent to make data-driven, context-aware portfolio allocation decisions.

Policy Gradient Methods:

In parallel with DQN, policy gradient methods augment our methodological suite. Policy gradients operate under the premise of directly optimising policy functions, making them adaptable and resilient in evolving market conditions. This class of algorithms, including Proximal Policy Optimisation (PPO) and Trust Region Policy Optimisation (TRPO), offers a complementary perspective on portfolio optimisation. Through these methods, the agent learns not only the value of actions but also the ideal policies to pursue under varying market states.

In unison, DQN and policy gradient methods underpin our portfolio optimisation framework, enabling the agent to learn, adapt, and excel in the dynamic and often unpredictable world of financial markets.

This meticulous selection of data sources and DRL algorithms sets the stage for our comparative analysis, designed to shed light on the advantages and potential limitations of DRL in contrast to traditional portfolio optimisation techniques. By integrating diverse datasets and employing advanced DRL methodologies, we aim to unveil the transformative potential of this cutting-edge approach in the context of asset management and investment strategies.

5. Experimental Design:

The essence of our pursuit to enhance portfolio optimisation through the lens of Deep Reinforcement Learning (DRL) resides in the precision of experimental design. The orchestrated confluence of training and testing periods, well-defined portfolio constraints, and a judicious selection of evaluation metrics serves as the bedrock of our inquiry.

5.1. Training and Testing Periods:

Training Period:

The delineation of our experimental framework commences with a distinct demarcation between training and testing periods. The training phase, often referred to as the "learning period," encompasses a substantial historical span. During this phase, the DRL agent engages with the dynamic financial market, exploring diverse portfolio allocation strategies, and learning from the feedback obtained. The objective is to equip the agent with the ability to discern optimal allocation decisions while adapting to evolving market conditions.

Testing Period:

Upon the culmination of the training period, the testing phase ensues. The testing period represents a segment of historical data previously unseen by the DRL agent, thereby simulating real-world, out-of-sample market conditions. During this phase, the agent's learned policy is rigorously evaluated, thereby ascertaining its aptitude to generalise and perform effectively in dynamic financial landscapes.

5.2. Portfolio Constraints:

Risk Tolerance:

Portfolio optimisation invariably encompasses the incorporation of risk constraints. These constraints encapsulate an investor's risk tolerance, demarcating the maximum acceptable level of risk for a portfolio. In our experiments, we introduce varying risk tolerance levels, thereby reflecting the diversity of investor profiles. The agent is thus tasked with optimising portfolios that adhere to these predefined risk constraints.

Position Limits:

To ensure pragmatic portfolio allocations, position limits are imposed to restrict the concentration of capital in a single asset or a group of assets. These position limits are inherently conducive to the principles of diversification and risk management, and they underscore the practicality of our portfolio optimisation strategy.

5.3. Evaluation Metrics:

Sharpe Ratio:

The Sharpe ratio is a quintessential metric, denoting the risk-adjusted return of a portfolio. It encapsulates the excess return earned relative to the risk taken and serves as a pivotal measure of portfolio performance. Our analysis entails a meticulous scrutiny of the Sharpe ratio, facilitating a comparative assessment of portfolio performance vis-à-vis traditional methodologies.

Maximum Drawdown:

Maximum drawdown is a pivotal risk metric that delves into the extent of loss experienced by a portfolio from its historical peak to its trough. This measure encapsulates the potential downside risk inherent to a portfolio. In our analysis, we scrutinise the maximum drawdown, enabling us to discern the resilience of portfolios under adverse market conditions.

Performance Benchmarking:

In addition to intrinsic metrics, we establish a performance benchmark by juxtaposing the DRL-optimized portfolios with portfolios optimised through traditional methods. This benchmarking facilitates a comprehensive comparative analysis, providing insights into the relative performance of DRL in the realm of portfolio optimisation.

In sum, our experimental design meticulously amalgamates historical data, distinct training and testing phases, portfolio constraints, and a comprehensive selection of evaluation metrics. Through this methodical orchestration, we aspire to unravel the transformative potential of DRL in portfolio optimisation and delineate its relative advantages and potential limitations when juxtaposed against conventional methodologies.

6. Comparative Analysis:

The crux of our research endeavour materialises in the comparative analysis between Deep Reinforcement Learning (DRL)-based portfolio optimisation and conventional, traditional methods. This segment serves as the conduit for discerning the distinct performance attributes of DRL in contrast to time-tested methodologies. In presenting these results, we embark on a rigorous assessment of performance metrics, unravelling key insights into the outperformance and potential limitations of DRL in the realm of portfolio optimisation.

6.1. Performance Metrics:

Our comparative analysis hinges upon a comprehensive spectrum of performance metrics, which serve as the arbiter of portfolio optimisation effectiveness. These metrics encompass, but are not limited to, the following:

Sharpe Ratio:

The Sharpe ratio stands as the epitome of risk-adjusted performance assessment. It encapsulates the excess return garnered in relation to the risk undertaken, serving as an anchor point for our analysis.

Portfolio Volatility:

The inherent risk in portfolio management materialises in the form of volatility. Portfolio volatility is a quintessential measure, illuminating the degree of price fluctuation experienced by the portfolio constituents. This metric offers insight into the stability and risk exposure of the optimised portfolios.

Maximum Drawdown:

Maximum drawdown, a pivotal risk metric, delineates the magnitude of loss experienced by portfolios from their historical zenith to nadir. It provides critical insights into the resilience of portfolio strategies under adverse market conditions.

6.2. Findings and Insights:

In the pursuit of comparative analysis, our research has unearthed a panorama of findings that underscore the transformative potential of DRL in portfolio optimisation.

DRL Outperformance:

In discerning the instances where DRL outperforms traditional methods, we note its ability to navigate nonlinearity and adapt to dynamic market conditions. DRL-optimized portfolios, under varying risk constraints, exhibit superior risk-adjusted returns as evidenced by higher Sharpe ratios. The adaptability of DRL agents in navigating complex action spaces emerges as a pivotal attribute, enabling them to exploit nuanced market trends. DRL excels particularly in volatile market environments, where traditional methodologies often falter.

DRL Limitations:

While DRL showcases remarkable potential, it is not devoid of limitations. Instances of underperformance are discerned under certain market conditions characterised by stable, linear patterns. In such circumstances, the complexities harnessed by DRL may result in portfolios that are relatively less competitive in terms of risk-adjusted returns. Furthermore, DRL agents require substantial training data, rendering them susceptible to underperformance during market conditions that deviate significantly from historical data.

In sum, our comparative analysis alludes to the transformative potential of DRL in portfolio optimisation. The adaptability, resilience, and adaptiveness of DRL-optimized portfolios are reflected in superior risk-adjusted returns, particularly in non-linear and dynamic market scenarios. However, the nuanced interplay of DRL performance and traditional methods underscores the importance of a balanced, pragmatic approach. The integration of DRL into the portfolio optimisation toolkit stands as a testament to innovation in the realm of asset management, with a profound potential to enhance investment strategies.

7. Robustness and Sensitivity Analysis:

The viability of any portfolio optimisation framework hinges on its capacity to withstand the diverse vicissitudes of financial markets. In the light of our exploration into Deep Reinforcement Learning (DRL)-based portfolio optimisation, this segment embarks on a journey to unearth the robustness of DRL under diverse market conditions. Additionally, we delve into a sensitivity analysis, meticulously scrutinising the influence of hyper-parameters on the performance of our DRL models.

7.1. Robustness Analysis:

Market Conditions and DRL Performance:

Our first endeavour in robustness analysis is to assess the adaptability of DRL-based portfolios under varying market conditions. These market conditions span the spectrum from stable, linear patterns to highly volatile, non-linear trends. In volatile markets, DRL demonstrates a remarkable propensity to harness non-linearity and adapt portfolio allocations to maximise risk-adjusted returns. This adaptability stands as a testament to the resilience of DRL in dynamic financial landscapes.

Market Stability and DRL Underperformance:

Under conditions characterised by market stability and linear trends, the adaptability of DRL-based portfolios may lead to underperformance. These instances underscore the need for prudent model selection, as traditional methodologies may suffice under such conditions. The nuanced interplay of market dynamics and DRL performance requires astute calibration.

7.2. Sensitivity Analysis:

Hyper-Parameter Assessment:

Sensitivity analysis is instrumental in revealing the impact of hyper-parameters on the performance of DRL models. These hyper-parameters span a spectrum, encompassing learning rates, exploration strategies, and the depth of neural network architectures. Our analysis elucidates the response of DRL models to variations in hyper-parameters, unveiling their impact on the agent's capacity to learn and adapt.

Learning Rate Sensitivity:

The learning rate, serving as a cardinal hyper-parameter, is examined for its influence on model convergence and learning. A sensitivity analysis reveals the optimal learning rate range that fosters efficient model training without leading to overfitting or convergence issues.

Exploration vs. Exploitation Strategy:

The calibration of exploration vs. exploitation strategies is a focal point in our sensitivity analysis. We scrutinise the impact of varying strategies, including epsilon-greedy and soft actor-critic methods. The selection of these strategies significantly influences the ability of DRL agents to explore diverse portfolio allocations while exploiting known optimal policies.

Neural Network Depth and Width:

The architecture of the neural networks employed in DRL models is pivotal. Sensitivity analysis evaluates the depth and width of these networks, identifying the optimal configurations that balance model complexity and predictive accuracy. The choice of neural network architecture substantiates the adaptability of DRL models to complex decision spaces.

In conclusion, our robustness and sensitivity analysis serves as a critical testament to the adaptability and resilience of DRL-based portfolio optimisation. DRL exhibits remarkable capacity to adapt to diverse market conditions, maximising risk-adjusted returns in non-linear and volatile scenarios. Furthermore, sensitivity analysis illuminates the impact of hyper-parameters on DRL model performance, facilitating optimal model calibration. This meticulous exploration aligns with our commitment to precision and innovation in the realm of asset management and investment strategies.

8. Discussion and Implications:

The culmination of our research journey ushers in a realm of practical implications that resonate beyond the boundaries of theoretical inquiry. In this section, we delve into the practical ramifications of our research findings, particularly pertaining to the application of Deep Reinforcement Learning (DRL)-based portfolio optimisation in real-world asset management and investment strategies. As we traverse this discourse, we also confront the potential advantages and challenges that punctuate the landscape of DRL implementation in the realm of financial management.

8.1. Practical Implications:

Adaptability and Dynamic Decision-Making:

The transformative potential of DRL in portfolio optimisation emerges as a beacon of adaptability. DRL-based portfolios, underpinned by deep neural networks, exhibit a profound capacity to adapt to dynamic market conditions, leveraging historical data to discern optimal allocation strategies. This adaptability translates to real-world asset management, offering investors and portfolio managers a responsive, data-driven tool to navigate the multifaceted world of finance.

Enhanced Risk-Adjusted Returns:

Our comparative analysis underscores the potential for DRL to enhance risk-adjusted returns. DRL-optimized portfolios, particularly adept in non-linear and volatile market scenarios, hold the promise of superior performance. The practical implication is the potential to harness DRL for improved risk management and wealth accumulation.

Balancing Human Expertise and AI Innovation:

The integration of DRL into asset management and investment strategies instigates a paradigm shift. While AI innovation brings a wealth of data-driven decision-making capabilities, it does not obviate the need for human expertise. Practical implementation requires a harmonious interplay of human intuition and AI insights. Portfolio managers become orchestrators of AI-driven strategies, offering a holistic approach to asset management.

8.2. Advantages and Challenges:

Advantages:

1. Adaptability: DRL offers the capacity to adapt to evolving market conditions, maximising portfolio performance in dynamic environments.
2. Data-Driven Insights: DRL leverages historical data and additional sources to glean predictive insights, thereby enhancing decision-making.
3. Superior Risk-Adjusted Returns: In non-linear and volatile scenarios, DRL-based portfolios exhibit the potential for superior risk-adjusted returns.

4. Holistic Portfolio Management: The integration of DRL augments portfolio management, ushering in an era of data-driven and dynamic strategies.

Challenges:

1. Data Requirements: DRL models necessitate substantial historical data for robust training, rendering them less effective in periods of data scarcity.
2. Model Complexity: Complex DRL models can be challenging to interpret and calibrate, necessitating substantial computational resources.
3. Risk of Overfitting: The adaptability of DRL can result in overfitting, especially when applied to historical data that deviates significantly from current market conditions.
4. Implementation Costs: Incorporating DRL into real-world asset management requires investments in computational infrastructure and expertise, presenting cost challenges.

In summary, the practical implications of our research herald a transformative potential for DRL-based portfolio optimisation in the landscape of asset management and investment strategies. The adaptability, data-driven insights, and the promise of superior risk-adjusted returns underscore the allure of this approach. Nevertheless, the careful calibration of model complexities, risk mitigation, and the harmonious integration of AI with human expertise form the crucible for successful implementation. The future, as our research demonstrates, is likely to witness a symbiotic relationship between AI innovation and human ingenuity in the pursuit of optimal financial management.

9. Conclusion:

In conclusion, our comparative analysis between traditional portfolio optimisation methods and Deep Reinforcement Learning (DRL)-based approaches illuminates a transformative potential in the realm of asset management and investment strategies. The practical implications underscore the capacity of DRL to usher in an era of adaptability and data-driven decision-making, yielding the promise of enhanced risk-adjusted returns.

The key findings of our study affirm that DRL-based portfolio optimisation exhibits a remarkable propensity for outperformance, particularly in non-linear and volatile market scenarios. The adaptability of DRL agents, predicated on deep neural networks, infuses dynamism into portfolio management, enhancing the ability to navigate intricate market landscapes. The essence of our research contributes significantly to the field of finance and asset management, offering practitioners and investors a robust and innovative tool to augment their wealth accumulation strategies.

10. Future Directions:

As we stand on the precipice of this transformative juncture, there exist numerous avenues for future research that can advance the frontiers of DRL-based portfolio optimisation:

Refinement of DRL Techniques: Future research may delve into the refinement of DRL algorithms, honing their adaptability, robustness, and interpretability. This quest for optimal DRL models holds the promise of enhancing performance while mitigating challenges.

Expansion to Other Asset Classes: The application of DRL in portfolio optimisation need not be confined to equity investments. The extension to other asset classes such as fixed income, commodities, and alternative investments presents a frontier rife with potential.

Sophistication in DRL Architectures: The exploration of more sophisticated DRL architectures, encompassing deep recurrent networks, attention mechanisms, and unsupervised learning, beckons the field towards uncharted territories. These architectures may uncover hidden patterns and insights in financial data.

Integration of External Data Sources: The incorporation of an even broader spectrum of external data sources, such as unstructured data from news sentiment, social media, and alternative data providers, extends the dimensions of data-driven decision-making.

11. References:

1. Sutton, R. S., & Barto, A. G. (2018). Reinforcement Learning: An Introduction. MIT Press.

2. Fama, E. F., & French, K. R. (1992). The Cross-Section of Expected Stock Returns. Journal of Finance, 47(2), 427-465.

3. Silver, D., Hubert, T., Schrittwieser, J., Antonoglou, I., Lai, M., Guez, A., ... & Hassabis, D. (2017). Mastering Chess and Shogi by Self-Play with a General Reinforcement Learning Algorithm. Nature, 551(7676), 354-359.

4. Merton, R. C. (1971). Optimum Consumption and Portfolio Rules in a Continuous-Time Model. Journal of Economic Theory, 3(4), 373-413.

5. Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., ... & Hassabis, D. (2015). Human-level Control Through Deep Reinforcement Learning. Nature, 518(7540), 529-533.

6. Black, F., & Litterman, R. (1992). Global Portfolio Optimization. Financial Analysts Journal, 48(5), 28-43.

7. Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., ... & Wierstra, D. (2016). Continuous Control with Deep Reinforcement Learning. arXiv preprint arXiv:1509.02971.
8. Sharpe, W. F. (1966). Mutual Fund Performance. Journal of Business, 39(1), 119-138.
9. He, K., Zhang, X., Ren, S., & Sun, J. (2016). Deep Residual Learning for Image Recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (pp. 770-778).
10. Markowitz, H. (1952). Portfolio Selection. Journal of Finance, 7(1), 77-91.