

---

## **Intelligent Cloud-Based Deep Reinforcement Learning Architectures for Dynamic Portfolio Risk Prediction and Adaptive Asset Allocation**

**Henry P. Lockwood**

University of Helsinki, Finland

### **ABSTRACT**

The accelerating complexity of global financial markets, characterized by high-frequency trading, heterogeneous investor behavior, geopolitical shocks, and increasingly interconnected asset classes, has rendered traditional portfolio optimization and risk management paradigms insufficient for real-time decision making. Classical approaches rooted in static optimization and equilibrium-based assumptions, while foundational, fail to account for the nonstationary, nonlinear, and adversarial nature of modern financial environments. In response to these challenges, deep reinforcement learning has emerged as a powerful paradigm capable of learning adaptive decision policies directly from sequential market interactions, enabling dynamic portfolio rebalancing and risk-sensitive asset allocation under uncertainty. At the same time, the migration of financial analytics into cloud-native infrastructures has enabled scalable data ingestion, distributed learning, and near-real-time deployment of intelligent trading systems, thereby transforming the operational context in which algorithmic portfolio management occurs.

This study develops a comprehensive theoretical and methodological framework for intelligent cloud-based deep reinforcement learning systems dedicated to dynamic portfolio risk prediction and adaptive portfolio control. Drawing on advances in recurrent and actor-critic reinforcement learning, stochastic policy optimization, hyperparameter tuning, and multimodal data fusion, the paper situates recent developments within a coherent architectural perspective that links financial theory, machine learning, and cloud computing. Central to this discussion is the integration of intelligent cloud frameworks that allow reinforcement learning agents to continuously ingest market data, retrain risk models, and deploy updated policies in a distributed and resilient manner, as exemplified by recent research on cloud-native deep reinforcement learning for portfolio risk prediction (Mirza et al., 2025).

Through an extensive synthesis of the literature on reinforcement learning-based portfolio optimization, the study examines how risk can be modeled not merely as a static constraint but as an evolving state variable learned by an agent interacting with the market environment. The paper further explores how deep neural architectures, including recurrent networks and stochastic policy models, enable agents to capture long-range temporal dependencies, tail-risk dynamics, and regime shifts that are invisible to conventional variance-based models. The cloud dimension is analyzed not simply as a computational convenience but as a structural enabler of continuous learning, model governance, and large-scale deployment across heterogeneous asset universes.

Methodologically, the article develops a text-based but detailed design of a cloud-integrated reinforcement learning pipeline for portfolio risk prediction, incorporating environment modeling, reward shaping, off-policy learning, and automated hyperparameter optimization. The results are interpreted in relation to the broader literature, highlighting how cloud-enabled deep reinforcement

**learning architectures can achieve superior responsiveness to market volatility, improved drawdown control, and enhanced adaptability to structural breaks when compared with both classical optimization and non-cloud-based learning systems.**

**The discussion critically evaluates the epistemological and practical implications of delegating financial risk management to autonomous learning agents, addressing issues of interpretability, stability, regulatory oversight, and ethical responsibility. By positioning intelligent cloud frameworks as the next evolutionary step in financial decision systems, the article argues that deep reinforcement learning-driven risk prediction is not merely a technological innovation but a paradigmatic shift in how portfolio theory itself is operationalized in the digital age.**

## **KEYWORDS**

**Deep reinforcement learning, cloud computing, portfolio risk prediction, adaptive asset allocation, financial machine learning, dynamic portfolio optimization**

## **INTRODUCTION**

The intellectual history of portfolio theory has been shaped by a persistent tension between mathematical elegance and market realism. Early models, most notably those derived from the mean-variance framework, conceptualized portfolio selection as a static optimization problem in which rational investors allocate capital to maximize expected return for a given level of variance-based risk, an approach that remains deeply influential in both academic finance and industry practice (Cornuejols and Tutuncu, 2006). Yet the real-world behavior of financial markets has repeatedly demonstrated that risk is not a stable or exogenous quantity but an emergent property of evolving interactions among market participants, institutional constraints, technological innovations, and macroeconomic forces. This discrepancy has motivated decades of research aimed at developing adaptive, data-driven, and temporally aware portfolio management systems, culminating in the contemporary surge of interest in deep reinforcement learning for financial decision making (Harris, 2024).

Reinforcement learning, as formalized in the foundational work on dynamic programming, provides a conceptual and mathematical apparatus for modeling sequential decision problems in uncertain environments, where an agent learns optimal actions through trial and error guided by feedback from the environment (Bellman, 1957). In finance, this paradigm maps naturally onto the portfolio management process, in which a trading agent repeatedly observes market states, executes trades, and receives rewards in the form of profits, losses, or risk-adjusted performance metrics. Unlike traditional optimization methods that assume known distributions or stationary relationships, reinforcement learning allows the agent to infer latent market dynamics and adapt its strategy as new information arrives, a property that has proven especially valuable in nonstationary and high-dimensional financial settings (Pigorsch and Schafer, 2021).

The advent of deep learning further transformed this landscape by enabling reinforcement learning agents to approximate complex value functions and policies using multi-layer neural networks capable of extracting hierarchical representations from raw market data (Goodfellow et al., 2016). Deep reinforcement learning thus overcame the curse of dimensionality that had long constrained earlier tabular or linear methods, making it feasible to operate in environments characterized by hundreds of assets, multiple time scales, and noisy, nonlinear signals. Recent studies have demonstrated that deep reinforcement learning agents can outperform

---

---

classical benchmarks in portfolio optimization tasks, particularly when augmented with recurrent architectures that capture temporal dependencies and regime shifts (Almahdi and Yang, 2017; Hieu, 2020).

However, as deep reinforcement learning models grow in complexity and data requirements, their effective deployment increasingly depends on cloud-based infrastructures that provide scalable computation, distributed storage, and continuous integration of new data streams. In this context, intelligent cloud frameworks have emerged as a crucial enabling layer for operationalizing deep reinforcement learning in financial markets, allowing models to be trained, validated, and updated in near real time across geographically distributed trading environments (Mirza et al., 2025). Such frameworks are not merely technical conveniences but fundamental components of a new socio-technical system in which portfolio risk prediction and asset allocation are embedded within a constantly evolving digital ecosystem.

The concept of portfolio risk itself undergoes a profound transformation in this setting. Traditional risk measures, such as variance, beta, or even more sophisticated metrics like value at risk, are typically computed retrospectively and assume a fixed probabilistic structure that often fails during periods of market stress. Deep reinforcement learning, by contrast, treats risk as an implicit function of state, action, and time, learned through interaction with the environment rather than imposed a priori (Aboussalah and Lee, 2020). This shift enables the construction of agents that not only react to realized volatility but anticipate future risk by learning patterns associated with drawdowns, liquidity shocks, and cross-asset contagion. The integration of such agents into cloud-based systems further allows these learned risk representations to be continuously updated as new data arrives, enhancing both responsiveness and robustness (Mirza et al., 2025).

Despite these advances, the literature remains fragmented across disciplines, with finance scholars, machine learning researchers, and cloud architects often addressing related problems in isolation. Studies on reinforcement learning for portfolio optimization frequently focus on algorithmic performance without fully considering the infrastructural and governance implications of deploying such models at scale (Espiga-Fernandez et al., 2024). Conversely, research on cloud computing in finance tends to emphasize data management and compliance rather than the epistemological consequences of embedding adaptive learning agents into financial decision processes (Nawathe et al., 2024). This fragmentation has created a literature gap in which the architectural integration of deep reinforcement learning, portfolio risk prediction, and intelligent cloud systems remains under-theorized, despite its growing importance for both academia and industry.

The present article addresses this gap by developing a unified analytical framework that situates deep reinforcement learning-based portfolio risk prediction within the broader context of cloud-enabled financial infrastructures. By synthesizing insights from dynamic programming, neural network theory, and financial economics, the study aims to elucidate how intelligent cloud frameworks can support not only computational efficiency but also conceptual advances in how risk and return are modeled, learned, and acted upon in real time. The framework is grounded in recent empirical and theoretical contributions, including the demonstration that cloud-based deep reinforcement learning systems can dynamically adjust portfolio risk predictions in response to market volatility and structural changes (Mirza et al., 2025).

A central argument advanced here is that the migration of portfolio optimization into intelligent cloud environments marks a paradigmatic shift comparable to the original introduction of mean-variance analysis or the later adoption of algorithmic trading. In this new paradigm, portfolio management is no longer a static problem solved periodically by human analysts but a continuous learning process executed by autonomous

---

agents embedded in a distributed digital infrastructure. This transformation raises profound questions about transparency, accountability, and the nature of financial expertise, which must be addressed alongside technical considerations of model accuracy and computational performance (Ndikum and Ndikum, 2024).

The remainder of the article develops this argument through a detailed exploration of the methodological, empirical, and theoretical dimensions of cloud-based deep reinforcement learning for portfolio risk prediction. By grounding each stage of the analysis in the existing literature while extending it through integrative reasoning, the study seeks to contribute not only to the technical understanding of these systems but also to the broader discourse on the future of financial decision making in an era of intelligent machines and pervasive cloud computing (Sun et al., 2024).

## **METHODOLOGY**

The methodological foundation of intelligent cloud-based deep reinforcement learning for portfolio risk prediction rests on a multilayered architecture that integrates financial modeling, machine learning, and distributed computing into a coherent operational system. At its core, this architecture conceptualizes portfolio management as a sequential decision process in which an agent observes a multidimensional market state, selects an allocation action, and receives a reward that reflects both realized returns and evolving measures of risk, an approach deeply rooted in the theory of dynamic programming and optimal control (Bellman, 1957). However, unlike classical formulations that rely on explicit state transition probabilities and closed-form solutions, the deep reinforcement learning paradigm approximates the optimal policy through iterative interaction with the environment, using neural networks to model complex and nonlinear relationships among assets, market signals, and risk dynamics (Goodfellow et al., 2016).

The first methodological layer concerns the formalization of the financial environment as a reinforcement learning problem. In this context, the environment is defined by a state space that encapsulates current and historical market information, including asset prices, returns, volatility indicators, macroeconomic variables, and potentially alternative data sources such as sentiment or news, as advocated in multimodal portfolio optimization research (Nawathe et al., 2024). Actions correspond to portfolio allocation vectors that specify the proportion of capital invested in each asset, possibly including leverage or margin positions when supported by the trading framework (Gu et al., 2023). The reward function is designed to capture the investor's objectives, typically combining measures of profitability with penalties for risk exposure, drawdowns, or transaction costs, thereby embedding financial preferences directly into the learning process (Almahdi and Yang, 2017).

A distinctive feature of contemporary approaches is the dynamic modeling of risk within the reward structure, rather than treating it as a static constraint. By penalizing outcomes associated with large losses, high volatility, or extreme drawdowns, the reinforcement learning agent learns to internalize risk as part of its value function, adjusting its behavior in anticipation of adverse market conditions (Aboussalah and Lee, 2020). This perspective aligns with recent work demonstrating that deep reinforcement learning agents can effectively learn risk-sensitive policies that outperform traditional mean-variance strategies, particularly in turbulent markets (Acero et al., 2024).

The second methodological layer involves the selection of deep reinforcement learning algorithms suitable for continuous, high-dimensional portfolio control. Actor-critic methods, and in particular stochastic policy gradient approaches, have gained prominence due to their ability to handle continuous action spaces and incorporate

exploration through entropy maximization (Haarnoja et al., 2018). These methods maintain separate neural networks for the policy, which maps states to actions, and the critic, which estimates the expected return of state-action pairs, enabling stable and efficient learning in complex environments (Fujimoto et al., 2018). Recurrent neural networks are often integrated into these architectures to capture temporal dependencies and latent market regimes, reflecting the empirical reality that financial time series exhibit long memory and structural breaks (Aboussalah and Lee, 2020).

The integration of these learning algorithms into an intelligent cloud framework introduces a third methodological layer concerned with data ingestion, model training, and deployment. Cloud infrastructures provide elastic computational resources that allow reinforcement learning agents to be trained on vast historical datasets and continuously updated as new market data streams in, a capability that is essential for maintaining model relevance in fast-changing markets (Mirza et al., 2025). In such frameworks, data pipelines collect and preprocess real-time price feeds, macroeconomic indicators, and alternative data, feeding them into distributed training clusters where neural networks are updated using off-policy learning and experience replay. The trained models are then deployed as microservices that generate portfolio allocation recommendations or execute trades directly, closing the loop between learning and action (Sun et al., 2024).

Automated hyperparameter optimization constitutes a critical methodological component of this pipeline, as deep reinforcement learning models are notoriously sensitive to choices of learning rates, network architectures, and exploration parameters. Modern frameworks leverage Bayesian optimization and other adaptive search techniques to tune these hyperparameters efficiently, thereby improving model performance and stability without excessive manual intervention (Akiba et al., 2019). When embedded in a cloud environment, such optimization processes can be parallelized across multiple compute nodes, accelerating convergence and enabling rapid experimentation with alternative model configurations (Ndikum and Ndikum, 2024).

Despite these advantages, the methodology is not without limitations. The reliance on large volumes of historical data raises concerns about overfitting and the potential for learned policies to exploit spurious correlations that do not persist out of sample, a problem that has been widely discussed in the reinforcement learning finance literature (Harris, 2024). Moreover, the opacity of deep neural networks complicates the interpretability of portfolio decisions, posing challenges for risk managers and regulators who require transparent justifications for trading behavior (Espiga-Fernandez et al., 2024). Cloud-based deployment further introduces issues of data security, latency, and system resilience, which must be addressed through robust engineering and governance frameworks (Mirza et al., 2025).

Nevertheless, by combining advanced reinforcement learning algorithms with intelligent cloud infrastructures, the proposed methodology offers a powerful and flexible approach to dynamic portfolio risk prediction and adaptive asset allocation. It represents a synthesis of theoretical insights from finance and machine learning with practical considerations of large-scale deployment, providing a foundation for the empirical and interpretive analyses that follow (Sun et al., 2024).

## RESULTS

The interpretive results emerging from the application of intelligent cloud-based deep reinforcement learning to portfolio risk prediction reveal a consistent pattern of enhanced adaptability, risk sensitivity, and operational

scalability when compared with both classical optimization techniques and standalone machine learning models. Across the literature, studies report that reinforcement learning agents embedded in cloud infrastructures demonstrate a superior capacity to track evolving market conditions, adjust exposure in response to volatility, and manage drawdown risk, outcomes that align with the theoretical expectations of dynamic programming and stochastic control (Mirza et al., 2025).

One of the most salient findings concerns the ability of deep reinforcement learning agents to internalize complex risk dynamics through their reward structures and state representations. Rather than relying on exogenous risk metrics, these agents learn to associate certain market configurations with future adverse outcomes, enabling preemptive adjustments in portfolio allocation that reduce exposure to impending turbulence (Aboussalah and Lee, 2020). Empirical analyses reported in the literature suggest that such agents exhibit lower maximum drawdowns and more stable return profiles than mean-variance or Black Litterman-based strategies, particularly during periods of heightened volatility (Sun et al., 2024).

The cloud dimension further amplifies these effects by allowing continuous retraining and real-time deployment of updated policies. In intelligent cloud frameworks, reinforcement learning models are not frozen after an initial training phase but are perpetually refined as new data arrives, leading to a form of online learning that is better suited to nonstationary financial environments (Mirza et al., 2025). This capability has been associated with improved responsiveness to regime shifts, such as sudden changes in interest rates or geopolitical events, which often render static models obsolete (Ndikum and Ndikum, 2024).

Another important result pertains to the scalability of these systems. Cloud-based architectures enable the management of high-dimensional portfolios involving hundreds of assets, a task that would be computationally prohibitive for traditional optimization methods or on-premises learning systems (Pigorsch and Schafer, 2021). Distributed training and inference allow deep reinforcement learning agents to evaluate vast numbers of potential allocation strategies in parallel, effectively exploring a larger portion of the action space and discovering more nuanced trade-offs between risk and return (Gu et al., 2023).

From a methodological standpoint, the incorporation of automated hyperparameter optimization within the cloud pipeline has been shown to enhance both performance and stability. By systematically searching the hyperparameter space, these systems reduce the likelihood of suboptimal configurations that could lead to unstable learning or excessive risk-taking, a finding that underscores the importance of integrating tools such as adaptive optimization frameworks into financial reinforcement learning workflows (Akiba et al., 2019).

Despite these positive outcomes, the results also highlight persistent challenges. Deep reinforcement learning agents can exhibit periods of erratic behavior, particularly during early training or when confronted with unprecedented market conditions, raising concerns about their reliability in live trading environments (Harris, 2024). Cloud-based deployment mitigates some of these risks by enabling rapid rollback and model updates, but it does not eliminate the fundamental uncertainty inherent in learning-based systems (Mirza et al., 2025).

Overall, the results indicate that intelligent cloud-based deep reinforcement learning represents a significant advance in the domain of portfolio risk prediction, offering a level of adaptivity and scalability that is unattainable with traditional approaches. At the same time, they underscore the need for careful system design, monitoring, and governance to ensure that these powerful tools are deployed in a responsible and resilient manner (Espiga-Fernandez et al., 2024).

---

## DISCUSSION

The implications of integrating deep reinforcement learning with intelligent cloud frameworks for portfolio risk prediction extend far beyond incremental improvements in predictive accuracy or computational efficiency. At a theoretical level, this integration challenges foundational assumptions of financial economics by redefining risk, rationality, and optimality in terms of adaptive learning processes rather than static equilibria (Benhamou et al., 2020). In classical portfolio theory, risk is typically treated as a quantifiable attribute of asset returns, often assumed to be stable or at least slowly varying over time. Deep reinforcement learning, by contrast, conceptualizes risk as an emergent property of the interaction between an agent and a dynamic market environment, learned through experience and encoded in the agent's value function and policy network (Aboussalah and Lee, 2020).

This shift has profound epistemological consequences. If risk is not an externally imposed constraint but a learned representation, then portfolio management becomes an inherently subjective process, shaped by the agent's training data, reward design, and exploration strategy. In an intelligent cloud environment, these factors are further influenced by infrastructural choices, such as data sources, retraining frequency, and deployment pipelines, suggesting that financial outcomes are increasingly mediated by technological architectures (Mirza et al., 2025). This perspective invites a reevaluation of what it means for a portfolio to be optimal, as optimality is no longer defined solely by adherence to a mathematical criterion but by the evolving performance of a learning system embedded in a socio-technical context (Ndikum and Ndikum, 2024).

Scholarly debates in this area often revolve around the trade-off between adaptivity and interpretability. Proponents of deep reinforcement learning argue that the ability to capture nonlinear dependencies, regime shifts, and tail-risk dynamics justifies the use of complex and opaque models, particularly in markets where traditional assumptions are routinely violated (Sun et al., 2024). Critics, however, caution that the black-box nature of deep neural networks undermines transparency and accountability, making it difficult for investors and regulators to understand why a particular portfolio decision was made or to assess the associated risks (Espiga-Fernandez et al., 2024). Intelligent cloud frameworks exacerbate this tension by enabling rapid and automated deployment of updated models, potentially outpacing the ability of human overseers to evaluate their behavior (Mirza et al., 2025).

Another point of contention concerns the stability of learning-based systems in adversarial and reflexive markets. Financial markets are not passive environments but are shaped by the actions of participants, including algorithmic traders whose strategies may interact in unpredictable ways. Reinforcement learning agents trained in historical or simulated environments may behave differently when deployed in live markets, where their own actions can influence prices and liquidity (Harris, 2024). Cloud-based architectures facilitate rapid adaptation, but they also risk amplifying feedback loops if multiple agents respond to the same signals in similar ways, potentially contributing to systemic instability (Pigorsch and Schafer, 2021).

The literature also highlights important ethical and regulatory considerations. As deep reinforcement learning agents assume greater responsibility for portfolio management, questions arise about liability, fiduciary duty, and the distribution of risk between investors and system providers (Ndikum and Ndikum, 2024). Intelligent cloud frameworks complicate these issues by dispersing decision making across distributed networks and automated pipelines, challenging traditional notions of control and accountability (Mirza et al., 2025).

Despite these challenges, the integration of deep reinforcement learning and cloud computing offers unprecedented opportunities for advancing the science and practice of portfolio risk management. By enabling continuous learning, large-scale experimentation, and real-time deployment, intelligent cloud frameworks create a living laboratory in which financial theories can be tested and refined in situ (Sun et al., 2024). Future research can build on this foundation by exploring hybrid models that combine reinforcement learning with interpretable financial structures, as well as governance mechanisms that ensure ethical and stable operation (Benhamou et al., 2020).

## CONCLUSION

The convergence of deep reinforcement learning and intelligent cloud computing represents a transformative development in the field of portfolio risk prediction and asset allocation. By reframing portfolio management as a continuous learning process embedded in a scalable digital infrastructure, this paradigm offers a powerful alternative to static and assumption-laden traditional models. Grounded in the dynamic programming tradition yet enriched by modern neural network architectures, cloud-based reinforcement learning systems demonstrate a remarkable capacity to adapt to market volatility, anticipate risk, and optimize performance in complex financial environments (Mirza et al., 2025).

While significant challenges remain in terms of interpretability, stability, and governance, the theoretical and practical advances documented in the literature suggest that intelligent cloud frameworks will play an increasingly central role in shaping the future of financial decision making. As research continues to integrate insights from finance, machine learning, and cloud engineering, the vision of truly adaptive, risk-aware, and resilient portfolio management systems moves ever closer to realization (Sun et al., 2024).

## REFERENCES

1. Hieu, L. T. (2020). Deep Reinforcement Learning for Stock Portfolio Optimization. *International Journal of Intelligent Systems Applications*, 12(4), 35–47.
2. Mirza, M. H., Budaraju, A., Valiveti, S. S. S., Sarma, W., Kaur, H., and Malik, V. (2025). Intelligent Cloud Framework for Dynamic Portfolio Risk Prediction Using Deep Reinforcement Learning. In *Proceedings of the 2025 IEEE International Conference on Computing*, 54–59.
3. Pigorsch, U. and Schafer, S. (2021). High-Dimensional Stock Portfolio Trading with Deep Reinforcement Learning. *Quantitative Finance*, 21(5), 739–753.
4. Akiba, T., Sano, S., Yanase, T., Ohta, T., and Koyama, M. (2019). Optuna: A Next-Generation Hyperparameter Optimization Framework. *ArXiv Preprint*.
5. Ndikum, P. and Ndikum, S. (2024). Advancing Investment Frontiers: Industry-grade Deep Reinforcement Learning for Portfolio Optimization. *International Journal of Intelligent Systems Applications*, 16(3), 89–102.
6. Gu, J., Du, W., Rahman, A. M. M., and Wang, G. (2023). Margin Trader: A Reinforcement Learning Framework for Portfolio Management with Margin and Constraints. *Proceedings of the ACM International Conference on AI in Finance*, 610–618.

7. Espiga-Fernandez, F., Garcia-Sanchez, A., and Ordieres-Mere, J. (2024). A Systematic Approach to Portfolio Optimization: A Comparative Study of Reinforcement Learning Agents, Market Signals, and Investment Horizons. *International Journal of Intelligent Systems Applications*, 16(2), 45–57.
8. Aboussalah, A. M. and Lee, C. G. (2020). Continuous Control with Stacked Deep Dynamic Recurrent Reinforcement Learning for Portfolio Optimization. *Expert Systems with Applications*, 140, 112891.
9. Cornuejols, G. and Tutuncu, R. (2006). *Optimization Methods in Finance*. Cambridge University Press.
10. Sun, R., Stefanidis, A., Jiang, Z., and Su, J. (2024). Combining Transformer-based Deep Reinforcement Learning with Black Litterman Model for Portfolio Optimization. *International Journal of Intelligent Systems Applications*, 16(4), 12–25.
11. Benhamou, E., Saltiel, D., Ungari, S., and Mukhopadhyay, A. (2020). Bridging the Gap between Markowitz Planning and Deep Reinforcement Learning. *ArXiv Preprint*.
12. Haarnoja, T., Zhou, A., Abbeel, P., and Levine, S. (2018). Soft Actor-Critic: Off-Policy Maximum Entropy Deep Reinforcement Learning with a Stochastic Actor. *Proceedings of the International Conference on Machine Learning*.
13. Harris, L. (2024). Deep Reinforcement Learning for Financial Portfolio Optimization. *International Journal of Intelligent Systems Applications*, 16(1), 18–29.
14. Fujimoto, S., van Hoof, H., and Meger, D. (2018). Addressing Function Approximation Error in Actor-Critic Methods. *ArXiv Preprint*.
15. Goodfellow, I., Bengio, J., and Courville, A. (2016). *Deep Learning*. MIT Press.
16. Nawathe, S., Panguluri, R., Zhang, J., and Venkatesh, S. (2024). Multimodal Deep Reinforcement Learning for Portfolio Optimization. *International Journal of Intelligent Systems Applications*, 16(3), 32–48.
17. Acero, F., Zehtabi, P., Marchesotti, N., Cashmore, M., Magazzeni, D., and Veloso, M. (2024). Deep Reinforcement Learning and Mean-Variance Strategies for Responsible Portfolio Optimization. *International Journal of Intelligent Systems Applications*, 16(2), 72–85.