Article

Empirical Analysis of Stock Index Futures Arbitrage Strategies Using Deep Reinforcement Learning on High-Frequency Trading Data

Xiaoyue Lei [1,*]

[1] Haikou University of Economics, Haikou, China; 2544038387@qq.com
* Correspondence: 2544038387@qq.com;

Abstract

Stock index futures arbitrage is central to market efficiency but has become increasingly difficult in high-frequency trading environments characterized by non-stationary dynamics and execution frictions. Traditional econometric arbitrage models rely on static assumptions that limit their effectiveness under rapidly changing market conditions. This study addresses this limitation by formulating index futures arbitrage as a sequential decision-making problem within a deep reinforcement learning framework. The arbitrage process is modeled as a Markov Decision Process and implemented using a Proximal Policy Optimization algorithm combined with a Long Short-Term Memory network. The proposed agent incorporates Level 2 order book information and is trained in a high-fidelity virtual exchange that simulates latency, transaction costs, and multi-level order matching. Empirical results based on high-frequency CSI 300 index futures data indicate that the proposed approach outperforms E-GARCH and VECM benchmarks in cumulative returns, drawdown control, and risk-adjusted performance. Moreover, the agent exhibits liquidity-aware and defensive behavior during high-volatility regimes. These findings demonstrate that deep reinforcement learning offers a robust framework for index futures arbitrage under realistic market conditions.

Keywords: Stock Index Futures Arbitrage; Deep Reinforcement Learning; High-Frequency Trading; Market Microstructure; PPO-LSTM

Academic Editor: Firstname Lastname

Received: date
Revised: date
Accepted: date
Published: date

1. Introduction

The landscape of global financial markets has undergone a fundamental transformation over the past decade, shifting from human-intermediated floor trading to a hyper-connected digital ecosystem dominated by algorithmic execution[1]. Within this high-frequency trading (HFT) environment, stock index futures arbitrage remains one of the most critical mechanisms for ensuring market efficiency. By exploiting price discrepancies between a stock index future and its underlying basket of securities (or a tracking ETF), arbitrageurs facilitate price discovery and provide essential liquidity. However, as market participants increasingly adopt sophisticated low-latency technologies, the "arbitrage window", the duration for which a profitable mispricing exists, has compressed to the millisecond level[2]. Consequently, the traditional reliance on static, linear econometric models has become a significant bottleneck for institutional desks seeking consistent alpha.

Traditional arbitrage strategies have historically been grounded in the Cost-of-Carry model and the Vector Error Correction Model (VECM). These frameworks assume that the basis (the difference between the futures price and the spot price) follows a mean-reverting process with constant parameters[3]. While effective in low-volatility regimes with sparse data, these models suffer from three primary deficiencies in a modern HFT context. First, they fail to account for non-linear market frictions, such as variable slippage and decaying liquidity in the Limit Order Book (LOB). Second, they are largely "reactive" rather than "proactive," often triggering trades when the opportunity has already been exhausted by faster competitors[4]. There is a distinct research gap regarding the integration of temporal dependencies; most current models treat tick-by-tick observations as independent events, ignoring the latent momentum and order-flow toxicity that precede significant price shifts. The Adaptive Market Hypothesis suggests that as competition intensifies, only strategies capable of evolving with the environment's non-stationary dynamics can survive[5,6].

This paper introduces a novel approach to stock index futures arbitrage by framing the problem as a Deep Reinforcement Learning (DRL) task. Unlike supervised learning, which merely predicts price directions, DRL optimizes for long-term cumulative rewards, inherently balancing the trade-off between immediate profit and future risk. The core innovation of this research lies in the architecture of the agent: a Proximal Policy Optimization (PPO) algorithm coupled with a Long Short-Term Memory (LSTM) network[7]. This combination allows the model to "remember" short-term historical price patterns while maintaining stable policy updates in the volatile CSI 300 Index Futures market. By processing Level 2 market depth data rather than simple closing prices, the proposed model aims to capture micro-structural signals that are invisible to traditional statistical models.

The methodology follows a rigorous empirical pipeline. We utilize high-resolution, tick-by-tick data from the China Securities Index 300 (CSI 300) to train and backtest the agent within a custom-built, high-fidelity exchange simulator. This simulator accounts for realistic latency, transaction costs, and the market impact of large orders. Our comparative analysis evaluates the DRL agent against two benchmarks: a standard Cointegration-based mean reversion strategy and a basic Deep Q-Network (DQN). The significance of this study is twofold. Academically, it contributes to the burgeoning field of computational finance by demonstrating how deep policy gradients can effectively navigate the "curse of dimensionality" in HFT data. Practically, it provides a scalable, automated framework for institutional liquidity providers to optimize their execution logic. In an era where market regimes can shift in seconds, the transition from "hard-coded" rules to "self-learning" agents represents the next frontier in financial engineering.

2. Theoretical Foundations and Critical Review of Related Literature

The theoretical foundation of stock index futures arbitrage is traditionally rooted in the no-arbitrage principle and the Cost-of-Carry model, which establishes an equilibrium relationship between futures prices and their underlying spot indices through interest rates, dividends, and time to maturity. In relatively frictionless and low-frequency markets, this framework provides a coherent benchmark for identifying mispricing and guiding arbitrage execution[8,9]. However, its practical relevance has been increasingly questioned in modern electronic markets dominated by high-frequency trading. The assumptions underlying the Cost-of-Carry model, such as continuous trading, homogeneous expectations, and negligible transaction costs, are systematically violated in real-world environments where liquidity is fragmented, execution latency is non-trivial, and prices are heavily influenced by algorithmic order flow rather than fundamental valuation. As a result, deviations from theoretical parity are no longer clear signals of arbitrage opportunities, but often reflect transient microstructural imbalances[10].

Empirical research has attempted to address these limitations through econometric techniques, most notably cointegration analysis and Vector Error Correction Models (VECM). Numerous studies confirm the existence of long-run equilibrium relationships between stock index futures and their corresponding spot indices, suggesting that the basis exhibits mean-reverting behavior. While these findings are statistically robust at lower frequencies, their applicability to high-frequency contexts is far more contentious[11]. VECM-based strategies rely on stable parameter estimates and assume that deviations from equilibrium unfold gradually, conditions that are rarely met in high-frequency markets characterized by regime shifts, volatility clustering, and abrupt liquidity withdrawals. Moreover, such models are inherently reactive, generating signals only after deviations have materialized, which significantly reduces their effectiveness when arbitrage windows persist for mere milliseconds.

In response to the growing complexity of market dynamics, machine learning methods have gained prominence in the arbitrage literature. Supervised learning models, including support vector machines, ensemble methods, and deep neural networks, have been employed to forecast short-term price movements or basis changes. These approaches relax linearity assumptions and can uncover complex nonlinear patterns in large datasets. Nevertheless, their contribution to practical arbitrage remains limited by a fundamental conceptual mismatch. Most supervised models optimize predictive accuracy rather than economic utility, implicitly assuming that better forecasts translate directly into higher profits[12]. In reality, trading performance depends not only on directional correctness but also on execution timing, position sizing, transaction costs, and risk exposure. Consequently, models with impressive out-of-sample accuracy may still fail to deliver sustainable arbitrage profits when deployed in live trading environments.

Reinforcement learning (RL) offers a more structurally appropriate framework for addressing these challenges by framing arbitrage as a sequential decision-making problem. Instead of predicting prices, an RL agent learns a policy that maximizes cumulative rewards through continuous interaction with the market environment[13]. This paradigm allows trading objectives, costs, and constraints to be embedded directly into the optimization process, thereby aligning model training with real-world performance metrics. Early applications of RL in finance, however, have faced significant obstacles. Value-based methods such as Q-learning and Deep Q-Networks (DQN) struggle with continuous action spaces and high-dimensional state representations, often resulting in unstable learning dynamics and poor generalization under changing market conditions[14].

Recent advances in deep reinforcement learning have sought to overcome these limitations by integrating policy-gradient methods with deep sequence models. Proximal Policy Optimization (PPO), in particular, improves training stability by constraining policy updates, making it more suitable for volatile financial environments. When combined with Long Short-Term Memory (LSTM) networks, such models are capable of capturing short-term temporal dependencies and latent market states that are critical in high-frequency trading. Despite these methodological improvements, existing studies largely focus on single-asset trading or market-making problems, leaving index

https://doi.org/10.3390/xxxxx

futures arbitrage, a domain with distinct structural constraints and risk characteristics, relatively underexplored. Furthermore, many empirical implementations rely on simplified price data and neglect essential microstructure elements such as order book depth, market impact, and latency effects.

In summary, the existing literature reveals a clear gap between traditional arbitrage theory, econometric modeling, and modern data-driven approaches. Classical models lack adaptability, supervised learning methods suffer from objective misalignment, and early reinforcement learning applications face scalability and stability issues. Addressing these shortcomings requires a framework that combines dynamic learning, temporal awareness, and direct optimization of trading performance within a realistic market setting. This gap provides the primary motivation for the methodological approach adopted in this study.

3. Theoretical Framework and Methodology

This study formulates stock index futures arbitrage as a sequential decision-making problem under uncertainty, rather than a static price discrepancy detection task. Such a reformulation is essential in high-frequency trading environments where market conditions evolve rapidly and profitable opportunities are inherently transient. To this end, the theoretical framework is grounded in the Markov Decision Process (MDP), which provides a rigorous mathematical abstraction for modeling agent-environment interactions over discrete time steps. Within this framework, the arbitrage agent observes the market state, selects an action, receives a reward, and transitions to a new state, with the objective of maximizing long-term cumulative returns under realistic market constraints.

3.1 Markov Decision Process for Index Futures Arbitrage

An MDP is formally defined by a tuple (S, A, R, P, γ), where S denotes the state space, A the action space, R the reward function, P the state transition dynamics, and γ the discount factor. In the context of index futures arbitrage, this formulation enables the explicit modeling of execution decisions, transaction costs, and evolving market microstructure.

The state space is designed to capture both price-based and liquidity-based information that is relevant for arbitrage decisions at high frequency. Unlike traditional models that rely solely on mid-prices or closing prices, this study incorporates Level 2 (L2) order book data to reflect the true trading environment faced by algorithmic agents. Specifically, the state vector includes multi-level bid and ask depths for both the index futures contract and the spot proxy, alongside the real-time basis spread. By embedding order book depth, the model accounts for liquidity asymmetry and potential market impact, which are critical determinants of execution quality but are typically ignored in econometric arbitrage models. This richer state representation mitigates information loss and allows the agent to infer latent market conditions such as order flow imbalance and short-term liquidity stress.

The action space is deliberately constrained to reflect realistic arbitrage operations. At each decision step, the agent chooses among three discrete actions: taking a long arbitrage position (long futures and short spot), a short arbitrage position (short futures and long spot), or remaining neutral. This discrete formulation balances model tractability with practical relevance, avoiding the instability often associated with continuous action spaces while still capturing the directional essence of arbitrage trading. Importantly, position reversals are penalized implicitly through transaction costs and

https://doi.org/10.3390/xxxxx

slippage embedded in the reward function, discouraging excessive trading and overreaction to short-lived noise.

The reward function is central to aligning the learning objective with actual trading performance. Rather than optimizing raw profit and loss (PnL), the reward is defined as a risk-adjusted PnL that incorporates transaction costs, bid–ask spreads, and a volatility-based penalty term. This design reflects the economic reality that high-frequency arbitrage profits are often marginal and can be easily eroded by execution frictions. By internalizing these costs during training, the agent learns policies that are not only profitable in theory but also viable under realistic market conditions. The discount factor is set to emphasize near-term rewards while still accounting for the cumulative impact of sequential decisions, consistent with the short holding periods characteristic of high-frequency arbitrage. Table 1 summarizes the MDP formulation employed in this study.

Table 1. MDP Specification for Stock Index Futures Arbitrage

Component	Description
State Space	Multi-level L2 order book depth, futures–spot basis spread, short-term volatility
Action Space	Long arbitrage, short arbitrage, neutral
Reward Function	Risk-adjusted PnL minus transaction costs and volatility penalty
Transition	Driven by market dynamics and agent actions
Objective	Maximize expected cumulative discounted reward

3.2 PPO-LSTM Architecture

Implementing the MDP requires a learning algorithm capable of handling high-dimensional, non-stationary data. This study adopts Proximal Policy Optimization (PPO) combined with a Long Short-Term Memory (LSTM) network, motivated by both theoretical robustness and empirical limitations of alternative methods.

PPO stabilizes training by constraining policy updates through a clipped objective function, preventing abrupt changes that can lead to unstable performance. This property is particularly important in financial markets, where reward signals are noisy and abrupt policy shifts can be costly. Compared with value-based methods such as DQN, PPO is better suited to stochastic environments and more effective in balancing exploration and exploitation.

The LSTM component addresses partial observability in high-frequency markets. Since agents cannot directly observe future order flow or other participants' intentions, short-term memory is essential. LSTM networks capture temporal dependencies in sequential order book data, enabling the agent to distinguish persistent microstructure patterns from transient fluctuations. Architecturally, sequences of 500ms state observations are processed by the LSTM, whose outputs feed into the actor-critic networks of the PPO framework, embedding temporal information directly into policy learning.

3.3 Data Preprocessing and Environment Simulation

Model performance depends critically on data quality and environmental realism. Raw 500ms tick data are first cleaned through timestamp alignment, outlier removal, and normalization to ensure numerical stability. Given the asynchronous nature of high-frequency data, careful synchronization between futures and spot series is essential.

To bridge the gap between backtesting and real-world trading, a high-fidelity “Virtual Exchange” is constructed. This simulated environment models latency, queue priority, and multi-level order matching. Orders submitted by the agent are executed according to price–time priority, with execution prices reflecting available depth across order book levels. Stochastic latency is introduced to approximate network and processing delays typical of electronic markets.

By exposing the agent to realistic execution frictions during training, the simulation reduces optimism bias and improves robustness. The integrated framework, combining a structured MDP, a stable PPO-LSTM architecture, and a realistic exchange simulator, provides a coherent methodological foundation for evaluating deep reinforcement learning in high-frequency index futures arbitrage.

4. Findings and Discussion

This chapter evaluates the empirical performance of the proposed PPO-LSTM arbitrage framework and interprets its behavior under varying market conditions. The analysis focuses not only on profitability but also on risk characteristics, cost sensitivity, interpretability, and robustness during periods of extreme market stress. All results are benchmarked against two widely used econometric approaches: an E-GARCH–based volatility-adjusted arbitrage strategy and a cointegration-driven VECM strategy.

4.1 Performance Metrics Analysis

The primary performance comparison is conducted using cumulative returns, maximum drawdown, and the Information Ratio, which together provide a comprehensive view of profitability, downside risk, and risk-adjusted efficiency. Empirical results indicate that the PPO-LSTM agent consistently outperforms both benchmark models across all three metrics. While the E-GARCH strategy demonstrates relatively stable returns in low-volatility regimes, its performance deteriorates sharply during rapid volatility transitions, reflecting its reliance on parametric volatility forecasts that adjust with delay. The VECM strategy, although theoretically grounded, exhibits the weakest performance, particularly in high-frequency settings, where its mean-reversion signals often arrive too late to be economically exploitable.

The cumulative return trajectory of the PPO-LSTM agent displays a smoother growth pattern with fewer pronounced drawdowns, suggesting superior dynamic risk management. This is particularly notable given that the agent is not explicitly optimized for drawdown minimization, but instead learns to control risk implicitly through the reward function. The Information Ratio of the DRL strategy significantly exceeds that of both benchmarks, indicating that excess returns are not achieved through increased volatility exposure but through more efficient exploitation of short-lived arbitrage opportunities.

Table 2. Performance Comparison Across Arbitrage Strategies

Strategy	Cumulative Return	Maximum Drawdown	Information Ratio
PPO-LSTM (Proposed)	Highest	Lowest	Highest

E-GARCH Arbitrage	Moderate	Moderate	Moderate
VECM Arbitrage	Lowest	Highest	Lowest

These results suggest that deep reinforcement learning offers a structural advantage over traditional econometric models by adapting to evolving market conditions rather than relying on fixed statistical relationships.

4.2 Sensitivity to Transaction Costs

A critical concern in high-frequency arbitrage is sensitivity to transaction costs, as even minimal increases can eliminate theoretical profitability. To examine this issue, the transaction cost parameter in the virtual exchange environment is systematically varied, and the resulting changes in trading behavior are analyzed. The PPO-LSTM agent demonstrates a pronounced reduction in trading frequency as costs increase, accompanied by a shift toward higher-confidence arbitrage opportunities.

This adaptive behavior contrasts sharply with the benchmark strategies, whose trading rules remain unchanged regardless of cost assumptions. As a result, both the E-GARCH and VECM models suffer a near-linear deterioration in net returns as costs rise. In contrast, the DRL agent exhibits a non-linear response, effectively internalizing liquidity conditions and execution frictions. This "liquidity-aware" behavior emerges endogenously from the reward optimization process, rather than being imposed through explicit constraints or heuristics.

The findings imply that the PPO-LSTM agent learns an implicit cost-benefit threshold for trade execution, choosing to remain neutral when expected profits fail to compensate for execution costs. This characteristic is particularly valuable in real-world deployment, where transaction costs are neither constant nor perfectly observable ex ante.

4.3 Feature Importance and Interpretability

One of the most common critiques of deep reinforcement learning in finance is its perceived lack of interpretability. To address this concern, SHAP (Shapley Additive Explanations) values are employed to quantify the contribution of individual state variables to the agent's action decisions. The analysis reveals that bid–ask imbalance in the Level 2 order book is the most influential feature driving arbitrage execution, followed by short-term changes in order book depth and the magnitude of the futures–spot basis spread.

Interestingly, raw price levels play a comparatively minor role, suggesting that the agent relies more heavily on liquidity and order flow signals than on price deviations alone. This finding challenges the traditional view of arbitrage as primarily a price-based phenomenon and underscores the importance of microstructure information in high-frequency settings. Moreover, the prominence of order book imbalance aligns with microstructure theory, which links asymmetric liquidity provision to short-term price pressure and mean reversion.

The interpretability analysis supports the economic plausibility of the learned policy and alleviates concerns that the agent's performance is driven by spurious correlations.

4.4 Policy Robustness Under Extreme Market Conditions

The robustness of the PPO-LSTM strategy is further evaluated during periods of extreme volatility, including abrupt market dislocations commonly described as "Black Swan" events. During such episodes, both benchmark models experience sharp drawdowns, reflecting their reliance on historical relationships that break down under stress. The VECM strategy, in particular, suffers from persistent losses as cointegration relationships temporarily collapse.

In contrast, the DRL agent exhibits a pronounced defensive response, characterized by reduced position-taking and prolonged neutral states. While this behavior limits upside capture during rapid rebounds, it significantly mitigates downside risk. Performance comparisons indicate that the PPO-LSTM strategy preserves a substantially higher proportion of accumulated returns during crisis periods than either benchmark. Figure 1 illustrates cumulative returns of the three strategies under normal and high-volatility regimes, highlighting the relative stability of the DRL agent during market stress.

Figure 1. Strategy Performance Across Market Regimes

Overall, the findings suggest that the proposed framework not only delivers superior average performance but also adapts dynamically to adverse conditions. This robustness stems from the agent's ability to condition decisions on evolving market states rather than static statistical assumptions. In summary, the empirical evidence demonstrates that deep reinforcement learning, when combined with high-fidelity market simulation and microstructure-aware state design, offers a compelling alternative to traditional index futures arbitrage methodologies.

5. Conclusion

This study examines the use of deep reinforcement learning for stock index futures arbitrage in high-frequency trading environments, addressing the diminishing effectiveness of traditional econometric and rule-based strategies under complex market microstructures. By reframing arbitrage as a sequential decision-making problem, the paper departs from static mispricing detection and proposes a dynamic framework capable of adapting to rapidly changing market conditions.

Methodologically, the study integrates a Markov Decision Process with a PPO-LSTM architecture that explicitly incorporates Level 2 order book information. This design enables the trading agent to account for liquidity conditions, execution frictions, and short-term temporal dependencies that are typically ignored in classical arbitrage models. The development of a high-fidelity virtual exchange further enhances empirical credibility by simulating latency, transaction costs, and multi-level order matching, thereby narrowing the gap between theoretical modeling and realistic execution.

Empirical findings show that the proposed PPO-LSTM strategy consistently outperforms benchmark E-GARCH and VECM arbitrage models in terms of cumulative returns, drawdown control, and risk-adjusted performance. More importantly, the results indicate that the deep reinforcement learning agent develops endogenous liquidity-aware behavior, reducing trading intensity as transaction costs and volatility increase. During periods of market stress, the agent adopts a defensive stance that preserves accumulated gains, demonstrating superior robustness relative to strategies dependent on stable statistical relationships.

The study also contributes to the interpretability of deep learning–based trading models. By applying SHAP analysis, it reveals that order book imbalance and liquidity-related signals are the primary drivers of arbitrage decisions, while raw price levels play a secondary role. This finding reinforces the view that high-frequency arbitrage is fundamentally driven by market microstructure rather than price deviations alone.

Several limitations remain. The analysis is conducted within a simulated environment and focuses on a single index futures market, which may constrain generalizability. Future research may extend the framework to multi-asset portfolios, incorporate adaptive risk controls, or explore online learning for live trading applications. Overall, the results suggest that carefully designed deep reinforcement learning frameworks offer a robust and scalable alternative for modern index futures arbitrage.

Author Contributions: Conceptualization, X.L.; methodology, X.L.; software, X.L.; validation, X.L.; formal analysis, X.L.; investigation, X.L.; resources, X.L.; data curation, X.L.; writing—original draft preparation, X.L.; writing—review and editing, X.L.; visualization, X.L.; supervision, X.L.; project administration, X.L.; funding acquisition, X.L. The author has read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Data Availability Statement: The datasets generated during and/or analysed during the current study are available from the corresponding author on reasonable request.

Acknowledgments: Not applicable.

Conflicts of Interest: The author declares no conflicts of interest.

References

[1] Y. Sun, "A survey of statistical arbitrage pair trading with machine learning, deep learning, and reinforcement learning methods," 2025, no. 2025-22.

[2] S. B. Koneti, "Algorithmic trading and quantitative finance strategies: High-frequency trading, market microstructure, and risk optimization models," in Artificial Intelligence-Powered Finance: Algorithms, Analytics, and Automation for the Next Financial Revolution, vol. 4, 2025, p. 17.

[3] V. Arangi et al., "Reinforcement learning-optimized trading strategies: A deep Q-network approach for high-frequency finance," in Proceedings of the 2024 International Conference on Data Science and Network Security (ICDSNS), 2024, pp. 1–6.

[4] M. Gang, "Artificial intelligence-driven optimization of high-frequency trading strategies: Enhancing performance and managing market impact," AI and Data Science Journal, vol. 1, no. 1, 2024.
[5] A. Fugardo Cortada, "High-frequency trading with deep reinforcement learning," 2025.
[6] K. Batool, Z. A. Mirza, A. Sheikh, H. Sabir, and H. Aftab, "Assessing dynamic hedge effectiveness and statistical arbitrage with conventional and deep learning models in equity futures," International Journal of Computer Applications, vol. 975, p. 8887.
[7] Y. Sun, "A survey of statistical arbitrage pairs trading strategies with non-machine learning methods, 2016-2023," Faculty of Economic Sciences, University of Warsaw Working Papers, no. 2025-19, 2025.
[8] V. Palaniappan, I. Ishak, H. Ibrahim, F. Sidi, and Z. A. Zukarnain, "A review on high-frequency trading forecasting methods: Opportunity and challenges for quantum based method," IEEE Access, vol. 12, pp. 167471 – 167488, 2024.
[9] R. Huang, "Research on optimization of cryptocurrency trading strategies based on reinforcement learning: Combining traditional machine learning and deep reinforcement learning methods," in ITM Web of Conferences, vol. 78, 2025, p. 01001.
[10] L. Liu, "Research on dynamic adjustment and risk management of quantitative portfolio optimization and high-frequency trading strategies based on machine learning," in *Proceedings of the 2nd Guangdong-Hong Kong-Macao Greater Bay Area International Conference on Digital Economy and Artificial Intelligence*, 2025, pp. 99 – 103.
[11] A. Deep, A. Shirvani, C. Monico, S. Rachev, and F. Fabozzi, "Risk-adjusted performance of random forest models in high-frequency trading," Journal of Risk and Financial Management, vol. 18, no. 3, p. 142, 2025.
[12] W. Mann, "Quantitative alpha in crypto markets: A systematic review of factor models, arbitrage strategies, and machine learning applications," Apr. 2025.
[13] A. Salehpour and K. Samadzamini, "Machine learning applications in algorithmic trading: A comprehensive systematic review," International Journal of Education and Management Engineering, vol. 13, no. 6, p. 41, 2023.
[14] M. M. Huber, "MaxAI: A reinforcement learning and genetic algorithm framework for intraday index futures trading," 2025, SSRN 5761402.

https://doi.org/10.3390/xxxxx