Finance Research Letters 70 (2024) 106266

Contents lists available at ScienceDirect

Finance Research Letters

journal homepage: www.elsevier.com/locate/frl

Portfolio optimization by enhanced LinUCB

He Ni [a,b], Qin Zhang [a,b,*], Xingjian Guo [b], Sultan Sikandar Mirza [c]

[a] *Tailong Finance School, Zhejiang Gongshang University, Hangzhou, Zhejiang, China*
[b] *School of Finance, Zhejiang Gongshang University, Hangzhou, Zhejiang, China*
[c] *School of Accounting, Finance and Economics, University of Waikato, Hamilton, New Zealand*

ARTICLE INFO

Keywords:
Portfolio management
Contextual LinUCB
Bandit learning

ABSTRACT

We study portfolio optimization through reinforcement learning, employing the contextual LinUCB approach to build an investment portfolio with changing market conditions. Contextual information enhances the LinUCB technique recommendation process, maximizing the efficacy of the model and the cumulative return. The recommendation criteria are to achieve the maximum level of confidence regarding the uncertainty surrounding the upward trend of asset returns within the pool of assets. Employing stocks from the Chinese A-share market, Hong Kong market, and US market as empirical tests, our findings demonstrate that the proposed model consistently outperforms the benchmark indices in our backtesting process.

1. Introduction

Modern portfolio theory suggests that investors should invest in a well-diversified portfolio using the mean–variance optimization framework (Markowits, 1952). Investors are, however, heterogeneous in their risk preferences and often make sub-optimal investment decisions (Bucciol and Miniaci, 2015). For instance, conservative investors rely on observed data to avoid suboptimal decisions but impede the exploration of potentially superior decisions, known as the *exploitation* approach. In contrast, aggressive investors prefer an *exploratory approach*, recognizing that existing data may be insufficient to determine the most advantageous decision, leading to suboptimal decisions. The risk preference of an investor may vary in different scenarios, even when considering the same individual. Therefore, it is unlikely that a universally applicable strategy can be accepted by all investors. Instead, the optimal approach must recognize the irrationality exhibited by human investors, taking into account the evolving preferences of investors and addressing both *exploitation* and *exploratory* challenges (Gomes et al., 2021).

This study aims to contribute to the field of cross-disciplinary finance and artificial intelligence by investigating the application of reinforcement learning (RL) techniques to recommend assets for constructing profitable investment portfolios. The objective is to propose a LinUCB strategy for portfolio optimization that leverages an advanced reinforcement learning framework. A key challenge addressed is the development of a robust preference function capable of quantifying and incorporating uncertainty regarding the future performance of assets. Through an analysis of the alignment between algorithmic logic and portfolio selection processes, we aim to contribute to the ongoing discourse surrounding the integration of machine learning techniques in finance. The prior literature suggests that two major benefits of machine learning for portfolio management are model flexibility and weight selection. The reinformance learning further pushes the frontier of portfolio management by allowing the model automatically adjusts based on pre-defined objective functions, i.e. minimizing risk or maximizing reward. A comprehensive literature review is provided in Appendix A.

* Correspondence to: No.18, Xuezheng Str., Zhejiang Gongshang University, Xiasha University Town, Hangzhou, 310018, China.
E-mail address: zhangqinwenya@mail.zjgsu.edu.cn (Q. Zhang).

https://doi.org/10.1016/j.frl.2024.106266
Received 1 December 2023; Received in revised form 1 August 2024; Accepted 10 October 2024

Available online 16 October 2024

To address the challenge of balancing exploration and exploitation that lead to suboptimal investment decision, we consider using one of the multi-arm bandit algorithms, namely the contextual Linear Upper Confidence Bound (LinUCB) algorithm of Li et al. (2010). The algorithm uses contextual information such as historical portfolio performance, investor risk preferences, and other relevant information to improve decision-making. The outcome of the decision depends on the specific circumstances. Following Li et al. (2010), we specify that the outcome of each arm can be represented linearly based on contextual information. In the empirical studies with stocks from the Chinese A-share market, Hong Kong market, and US market, we select stocks that exhibit optimal historical performance, taking into consideration investors' risk preferences. This involves considering stocks that have demonstrated consistent profitability and stability over time (exploitation) and exploring stocks with substantial growth potential that are likely to yield favorable performance in the future (exploration). To solve the exploration–exploitation dilemma, we estimate the expected utility by the upper confidence bound at a given degree.[1] We show that the performance of the recommended portfolio encompasses the benchmark market indices or peer group averages and minimum regret.

We contribute to the literature as follows. Methodologically, we propose a portfolio optimization method based on reinforcement learning with LinUCB, which extends the modern portfolio theory. The proposed framework benefits from using contextual information such as historical price and investors' individual characterizes to address exploitation and exploration, leading to improved portfolio efficiency. Empirically, we contribute to the growing literature on the applications of reinforcement learning in portfolio optimization. We test the effectiveness of this method using stocks listed on multiple markets. Our research enhances our understanding of how reinforcement learning techniques can benefit ordinary investors who can achieve above-average returns by using the LinUCB algorithm without investing in funds or seeking professional financial advice. Furthermore, our approach can easily adapt to the bond and mutual fund markets.

2. Research design

2.1. Problem setting

Given a pool of assets $a_k, k \in 1, 2, \dots, N$ with a known price history and potential investors with reasonable but varying risk preferences, we aim to propose a pre-trained strategy utilizing a contextual bandit learning algorithm. Given the variability in the dynamics of asset price movements, it is reasonable to assume that speculative investors might consider altering their positions at regular intervals of T days.[2] The approach has the capability to provide recommendations for a series of assets once at a time T according to investors' risk appetite. In the proposed bandit learning algorithm, we hereby define T days as a unit length of a learning loop during which the parameters used for the next recommendation will be updated, whilst the position of assets will not be changed.

We consider generating a numerical preference for each asset a, which we denote $H_t(a) \in \mathbb{R}$. The larger the preference, the more likely the asset will be taken/invested. Similarly, the investment probability on the asset a can be determined according to a soft-max distribution (i.e. Gibbs or Boltzmann distribution):

$$\mathrm{P}_t(a) \doteq \frac{e^{H_t(a)}}{\sum_{k=1}^{N} e^{H_t(k)}} \tag{1}$$

The formulation of a learning algorithm for soft-max action preferences can be reached through applying the concept of stochastic gradient ascent. Following the selection of the asset a and its subsequent return R_t, the preferences of the assets are subject to an updating process known as Stochastic Gradient Ascent.

$$\begin{aligned} H_{t+1}(a) &= H_t(a) + \beta \frac{\partial \mathbb{E}\left[R_t\right]}{\partial H_t(a)} \\ &= H_t(a) + \beta \left(R_t - \bar{R}_t\right)(1 - P_t(a)) \end{aligned} \tag{2}$$

where $\beta > 0$ is a pre-determined step size and $\bar{R}_t \in \mathbb{R}$ is the (accumulated) average of the return up to but not including t. If the current return R_t is higher than the average $\bar{R}_t$, then the probability of taking the asset a in the future increases, otherwise the probability decreases.

Given that each asset, a_k, has a significantly longer historical record compared to the time interval T in this study, it is imperative to consider the contextual information derived from its prior performance while estimating the initial preference $H_\tau(a_k)$[3] up to time τ.

$$H_\tau(a_k) \doteq \underbrace{Q_\tau(a_k)}_{\text{Expected return}} + \underbrace{\gamma \left(U(w_X, X) + w_c \sqrt{\frac{\ln \tau}{N_\tau(a_k)}} \right)}_{\text{Uncertainty associate with future return}} \tag{3}$$

[1] The LinUCB approach relies on the capacity to calculate upper confidence bound estimates for the anticipated returns based on the relevant information.

[2] It refers to a time interval that an investor would adjust according to his/her risk exposure. For simplicity, we treat it as an independent input like the risk preference.

[3] The time stamp, denoted as τ, is regularly updated following each time interval. Hence, in the context of a stock, τ represents the cumulative number of days between an arbitrary base (starting) date and the commencement of a bandit learning process. It denotes the historical data accessible to offer contextual knowledge.

where $Q_\tau(a_k)$ is a weighted average of the historical returns of the asset a_k, used as a representation of greedy actions or exploitation components within the bandit algorithm. It serves as an initial reference point for the preference $H_\tau(a_k)$. The other component of $H_\tau(a_k)$ pertains to the uncertainty surrounding[4] the precision of the return estimation. This refers to the exploration or non-greedy actions that an investor may consider in order to attain optimal outcomes.

The upper confidence bound (UCB) action selection approach is based on the notion that the inclusion of the exploration term, shown in Eq. (3), serves as an indicator of the level of uncertainty or variance in estimating future performance of the asset a_k. The parameter γ serves as an indicator of the investor's risk preference and measures the trade-off between exploitation and exploration. X represents the contextual information, while the variable w_X denotes the corresponding weights, which will be elaborated on in the next section. w_c is a constant known for the Chernoff-Hoeffding bound. It can help derive exploration strategies that ensure the chosen actions converge to the optimal ones while minimizing regret. $N_\tau(a_k)$ represents the number of days that have taken between the date of the initial public offering (IPO) of the asset a_k and the time τ.

2.2. Contextual information

We consider both investor and asset features (contexts) that include investors' risk preferences and historical return information such as realized volatility, associate skewness and kurtosis, and trading volume. In the LinUCB algorithm with N arms (N assets), each asset a_k receives a return $R_{a_k}(t)$ at each time step $t = 1, 2, \ldots$. We choose an optimal portfolio at a time, where optimality is defined as the best achievable accumulated returns.[5] The key is to identify the most suitable features and their associated weights to construct contextual information that effectively forecasts the uncertainty of future movements in asset returns in Eq. (3).

Since our aim is to perform short-term predictions of stock returns, we do not consider using macro factors such as the market environment, as suggested by Welch and Goyal (2008) and Christiansen et al. (2012) that the predictive power of macroeconomic fundamentals in forecasting short-term future volatility is limited. Our attention is on internal features such as skewness, kurtosis, trade volume, and institution holding. We utilize the subsequent selection criteria based on internal features.:

- Finance literature suggests that the skewness of risk premium is statistically and economically significant in short-term stock return prediction. Bufalo et al. (2022) contribute to the field of financial modeling by introducing a skew-geometric Brownian motion framework and offer a more accurate tool for predicting returns in portfolios with skewed assets. Indeed, many relevant studies in the literature show evidence that extended mean–variance-skewness models, multi-objective decision approaches, and neural network-based models all contribute to more robust portfolio optimization by accounting for the asymmetry in return distributions (Orlando and Bufalo, 2021; Patton, 2004; Zhang et al., 2021; Wang et al., 2018).
- Higher kurtosis in financial return distributions indicates a higher likelihood of increased volatility and albeit infrequent but severe consequences of significant fluctuations in the market, commonly known as tail events (Consigli, 2002; Van Oordt and Zhou, 2016). Therefore, in measuring the uncertainty of assets that may impact their future performance, kurtosis is needed.
- Increased trading volume typically corresponds to greater liquidity, resulting in more pronounced fluctuations in prices. Indeed, technical analyses view trading volume as a fundamental element in asset pricing (Lo and Wang, 2000; Amihud and Noh, 2021)
- The behavior of institutional investors has the capacity to impact market sentiment, interpreted as a valuable indicator to mitigate the occurrence of significant price volatility and to enhance overall market liquidity (Edelen et al., 2016). We use significant holdings of shares by institutional investors to account for the behavior of institutional investors.

To further quantify the uncertainty of stock performance, we employ the evaluation of both the volatility and risk inherent in the price fluctuations of a certain stock. We consider various prevalent techniques and metrics, including historical volatility, implied volatility, Value at Risk, Average True Range (ATR), and VIX index, among others (Baillie and DeGennaro, 1990; Bekaert and Hoerova, 2014). The uncertainty attributed to contextual information is estimated by a linear regression of the aforementioned features:

$$U(w_X, X) = \alpha(X) \sum_{i=1,2,\ldots} w_{X_i} X_i \tag{4}$$

where $\alpha(X)$ is a pre-trained adaptive parameter to calibrate the weights of contextual information. A threshold ζ is defined according to prior knowledge of the asset, when the volatility of the asset is larger than the threshold value, the $\alpha(X)$ can depress the uncertainty in order to avoid risky actions:

$$\alpha(X) = \begin{cases} \alpha_0, & V_{w_X X} \le \zeta \\ \alpha_0 + 1 - \exp(V_{w_X X} - \zeta), & V_{w_X X} > \zeta \end{cases} \tag{5}$$

[4] The idea of uncertainty includes both the upper and lower bounds of the expected return. We focus exclusively on the upper limit due to the intrinsic aggressiveness of the bandit learning-based method compared to the usual mean–variance investment methodology.

[5] The metric can be replaced by any well-defined utility or measures such as Sharpe ratio.

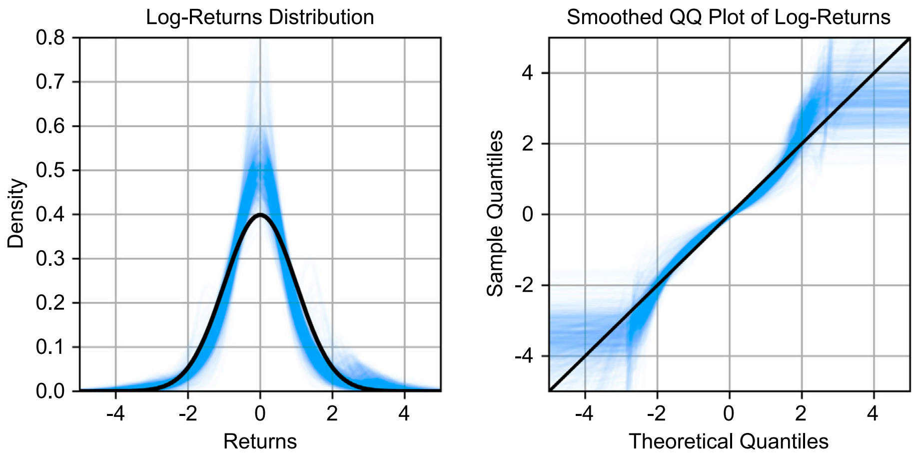

Fig. 1. The log-return distribution (left) and the QQ plot of log-return (right) are depicted. Out of a total of 4788 stocks, a random sample of 500 stocks was examined, revealing that the prospective stocks do not exhibit a normal distribution of log returns.

2.3. Roulette wheel selection

The criteria for selecting a portfolio is to prioritize assets with a better possibility of success, defined as those that have demonstrated strong historical performance, hold favorable contextual information, and are more likely to outperform other assets. In our study, we implemented a fundamental concept of natural selection known as roulette wheel selection, sometimes referred as fitness proportionate selection. The algorithm mimics the concept of a roulette wheel, where an individual's likelihood of being chosen is decided by their fitness rating or investment probability in this research. This extensively utilized process in genetic algorithms enables the identification of individuals that will be selected for reproduction.

Roulette wheel selection is executed by employing a metaphorical representation of a spinning wheel. The wheel is partitioned into pieces that are proportional to the probability of investing. The higher the investment probability of an asset, the larger its corresponding portion on the wheel. The selection of the asset for reproduction is determined by a random spin of the wheel.

Although the approach has components that involve randomness, its primary objective is to direct the selection procedure according to the investment probability. The selection process is iterated M times to create a portfolio consisting of M distinct assets.

3. Empirical study

This study utilizes 4788 stocks listed on the Chinese A-share market[6] as the selected pool of assets due to the notable liquidity of the Chinese stock market and the significant activity observed in initial public offerings (IPOs) in recent years. The daily closing (adjusted) prices of the candidate stocks were gathered from the beginning of 2019 to June 2023, including both the bear and bull markets.

We plot daily log returns in Fig. 1. It reveals the non-normal distribution, indicating the variance is not an appropriate metric for quantifying uncertainty.

To account for the varying volatility of the return distribution, we employ Conditional Value at Risk (CVaR)[7] rather than VaR to proxy the uncertainty associated with contextual information. The regressor comprises the skewness and kurtosis of historical returns, the trading volume, and the percentage of institutional investment.

$$W_c^* = \arg\min_{W_c} \| \underbrace{\frac{1}{1-\beta}\int_{\beta}^{1} \mathrm{VaR}_u(f)df}_{\text{Conditional VaR}} - [X_{ske}, X_{kur}, X_{vol}, X_{int}] \cdot W_c \|_2 \tag{6}$$

$\mathrm{VaR}_u(f)$ denotes the Value at Risk at level $u = 0.7$ of the loss distribution f. $\beta = 0.95$ is a frequently employed selection for a 95% confidence level. The Eq. (6) is utilized to estimate the ideal vector of weights, W_c, of contextual information employing an

[6] The Chinese A-share market is composed of a total of 4788 individual stocks. The Shanghai Stock Exchange currently comprises a total of 2080 stocks, with 410 of them listed on the Science and Technology Innovation Board. The Shenzhen Stock Exchange accommodates 2618 stocks, out of which 1132 are listed on the ChiNext Board. The newly established Beijing Stock Exchange currently houses 90 stocks.

[7] The higher the level of instability in an investment, the more likely it is that VaR may not provide a comprehensive assessment of the associated risks, as it is unconcerned with factors beyond its predetermined threshold. The CVaR model is designed to mitigate the limitations of the VaR concept. The utilization of CVaR in contrast to VaR typically results in a more cautious strategy with regards to risk exposure.

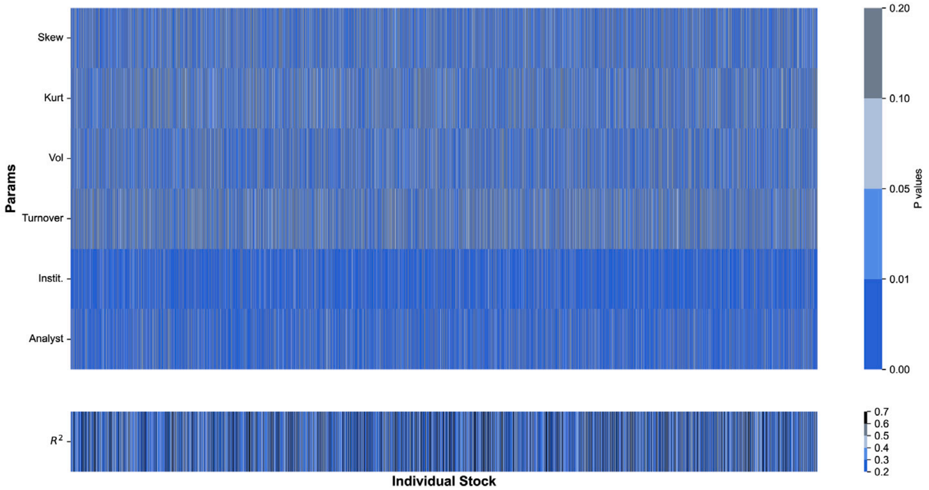

Fig. 2. The color bars represent the significance of the parameters estimated by Eq. (4). For most stocks (i.e., each column represents an individual stock), the parameters are significant($p < 0.05$ or in blue), and the quality of regressions are acceptable (i.e., in terms of high R^2 values shows in the lower part of the figure).

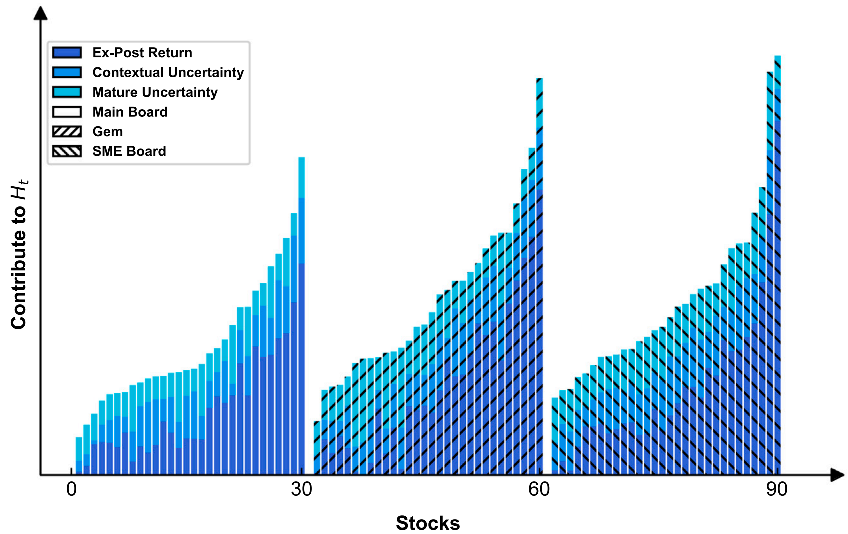

Fig. 3. The distribution of *preference* that consists of ex-post return (dark blue) and uncertainty (deep & light blue).

ordinary least squares regression on CVaR. Fig. 2 shows that the weights in the regressors, W_c, are significant in most stocks. This finding suggests that we may rely on this collection of regressors to effectively assess the uncertainty associated with contextual information, aligning with the literature discussed in Section 2.2.

The initial valuation of the preference H for each stock a_k is determined based on Eq. (1). Based on a randomly selected sample of 30 stocks from three categories, Fig. 3 illustrates the visual representation of variable H and its contributing factors the ex-post averaged return and the associated uncertainty due to contextual and mature information.

Evaluating the performance of the recommended portfolio encompasses benchmark comparison, which involves comparing it against market indices or peer group averages and minimum regret.[8] Fig. 4 shows that the performance of the proposed model can consistently outperform CSI300.[9]

[8] The concept of minimum regret is intricately connected to the idea of downside risk. Regret can be conceptualized as the quantifiable disparity between the results and the ideal results an investor could have achieved with perfect foresight.

[9] The CSI 300 Index (China Securities Index 300) is a major stock market index that tracks the performance of the top 300 A-share stocks listed on the Shanghai Stock Exchange (SSE) and the Shenzhen Stock Exchange (SZSE). It is widely recognized as a key benchmark for the Chinese A-share market and is used by investors.

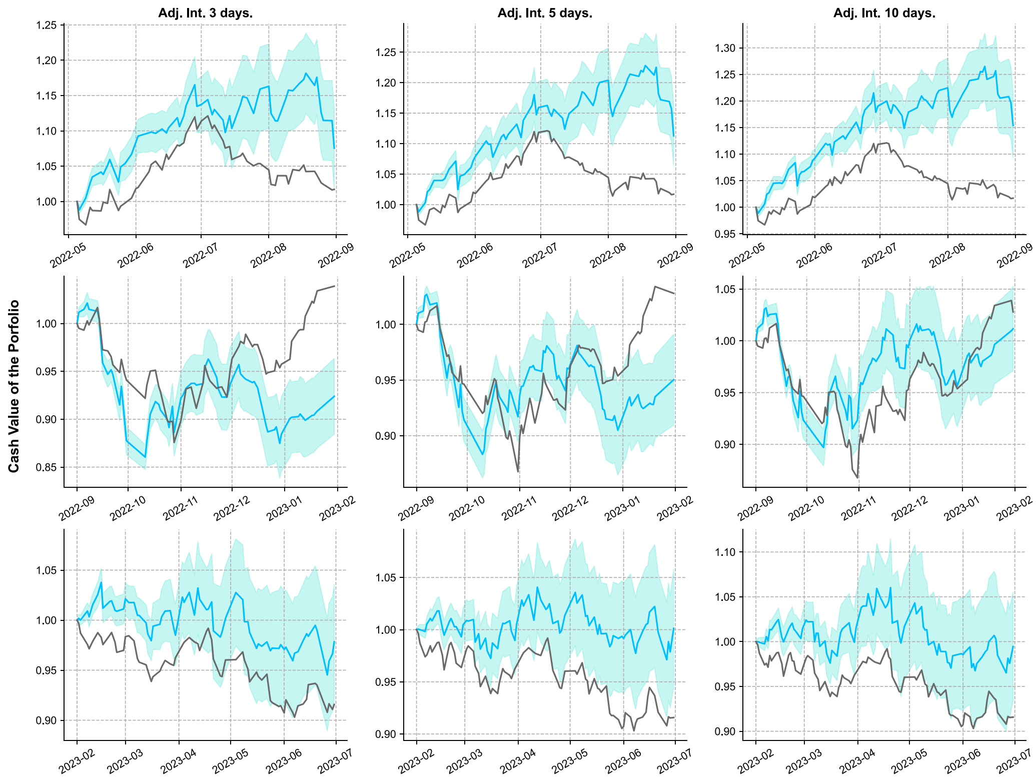

Fig. 4. The performance (cash value) of the *proposed* portfolio (blue) against the CSI300 index (black line). The shadow area depicts around 70% (i.e., mean$\pm 1\sigma$) of possible cash changes. Subfigures depict performance based on various options for the length of the adjustment period (3 days, 5 days, and 10 days) and the testing periods (downward reverse trend, upward reverse trend, and downward trend). The performance of *proposed* portfolios can generally exceed the benchmark index except in the period of sharp reverse (shown in the middle rows), as the blue areas are generally above the black line.

To demonstrate the robustness and consistency of the proposed strategy, we compare the performance of the proposed portfolio optimization algorithms with Hong Kong and US stock markets across different backtesting periods in Table 1. The results indicate that the proposed algorithms outperform the benchmark in terms of average return, Sharpe ratio, maximum drawdown, and information ratio, further validating our main results.

4. Conclusion and future work

In this paper, we propose a reinforcement learning framework for portfolio optimization by employing the contextual LinUCB with changing market conditions. Contextual information such as historical price, volatility and trading volume improve recommendation outcome, leading to maximized cumulative return. We set the recommendation criteria as the maximum level of confidence regarding the uncertainty of asset upward returns within the pool of assets. Our primary contributions are twofold. Methodologically, our work joins the growing literature on AI and finance, extending modern portfolio theory using state-of-art reinforcement learning techniques. Empirically, we demonstrate that contextual LinUCB consistently outperforms the market when applied to stocks listed on the Chinese A-share market, as well as those listed on the Hong Kong and US markets. This provides a challenging yet rewarding environment for testing the effectiveness of portfolio optimization algorithms.

Despite the success attained by the presented method, several areas still need to be studied. The concept of uncertainty can vary, and additional exogenous contextual information might improve the accuracy of estimating uncertainty. This study restricts the scope of uncertainty to encompass only risk, considering that ambiguity may also be persuasive. Similarly, exploitation might be predicated on richer content in addition to the accumulation of historical returns. Current research demonstrates the utility of implementing a prototype reinforcement learning algorithm in the context of portfolio management.

Table 1
Portfolio optimization algorithms with different stock markets.

Market		Avg. Return%		Sharpe R.		Max. Draw.%		Inf. R.
		Port.	Bench.	Port.	Bench.	Port.	Bench.	Port.
CN	2006–2009	52.62	39.60	6.36	4.98	68.88	72.30	3.82
	2010–2015	8.10	0.90	3.26	0.25	47.36	43.50	2.63
	2016–2019	23.38	4.24	4.19	1.25	25.58	32.46	11.99
	2020–2023	2.57	−4.66	0.06	−1.03	32.66	42.94	5.30
HK	2006–2009	29.15	9.99	4.36	3.65	73.39	65.13	5.59
	2010–2015	1.70	0.07	1.02	−0.14	42.13	34.90	0.60
	2016–2019	7.75	7.22	2.01	2.35	32.10	25.42	0.26
	2020–2023	−8.94	−12.09	−1.25	−1.98	36.15	52.19	0.66
US	2006–2009	7.15	−3.18	1.54	0.56	49.92	56.78	5.08
	2010–2015	10.66	10.33	1.02	0.65	25.63	19.39	0.73
	2016–2019	25.96	12.56	2.89	1.98	22.77	19.78	7.85
	2020–2023	23.25	10.00	2.95	2.08	23.17	25.30	8.28

Note: Applying proposed portfolios to the Chinese stock market, Hong Kong stocks, and US stocks across different backtesting periods: 2006 to 2009, 2010 to 2015, 2016 to 2019 and 2020 to 2023. For each of these backtesting periods, conduct 30 backtests to calculate the average returns (%), Sharpe ratio, maximum drawdown rate (%) and information ratio, and compare these metrics with the benchmark indices for each market, which are the CSI 300 Index, Hang Seng Index (HSI), and S&P 500 Index, respectively. In most experiments, it is evident that the proposed portfolio consistently outperforms the benchmark indices, exhibiting higher average returns, greater Sharpe ratio, lower maximum drawdown, and positive & reasonable information ratio.

CRediT authorship contribution statement

He Ni: Writing – original draft, Methodology, Formal analysis, Conceptualization. **Qin Zhang:** Writing – review & editing, Formal analysis. **Xingjian Guo:** Software, Formal analysis, Data curation. **Sultan Sikandar Mirza:** Writing – review & editing, Investigation.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was supported by the characteristic & preponderant discipline of key construction universities in Zhejiang Province (Zhejiang Gongshang University - Collaborative Innovation Center of Statistical Data Engineering Technology & Application), Zhejiang Gongshang University, China “Digital+” Disciplinary Construction Management Project (SZJ2022B001), Zhejiang Gongshang University Tailong Finance School Key Scientific Research Projects (1330XJ3623013) and Zhejiang Gongshang University Young Talent Program (QRK23027). Qin acknowledges financial support from the China Postdoctoral Science Foundation under Grant Number 2024M752860.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.frl.2024.106266.

Data availability

Data will be made available on request.

References

Amihud, Y., Noh, J., 2021. Illiquidity and stock returns II: Cross-section and time-series effects. Rev. Financ. Stud. 34 (4), 2101–2123.
Baillie, R.T., DeGennaro, R.P., 1990. Stock returns and volatility. J. Financ. Quant. Anal. 25 (2), 203–214.
Bekaert, G., Hoerova, M., 2014. The VIX, the variance premium and stock market volatility. J. Econometrics 183 (2), 181–192.
Bucciol, A., Miniaci, R., 2015. Household portfolio risk. Rev. Financ. 19 (2), 739–783.
Bufalo, M., Liseo, B., Orlando, G., 2022. Forecasting portfolio returns with skew-geometric Brownian motions. Appl. Stoch. Models Bus. Ind. 38 (4), 620–650. http://dx.doi.org/10.1002/asmb.2678.
Christiansen, C., Schmeling, M., Schrimpf, A., 2012. A comprehensive look at financial volatility prediction by economic variables. J. Appl. Econometrics 27 (6), 956–977.
Consigli, G., 2002. Tail estimation and mean–VaR portfolio selection in markets subject to financial instability. J. Bank. Financ. 26 (7), 1355–1382.

Edelen, R.M., Ince, O.S., Kadlec, G.B., 2016. Institutional investors and stock return anomalies. J. Financ. Econ. 119 (3), 472–488.
Gomes, F., Haliassos, M., Ramadorai, T., 2021. Household finance. J. Econ. Lit. 59 (3), 919–1000.
Li, L., Chu, W., Langford, J., Schapire, R.E., 2010. A contextual-bandit approach to personalized news article recommendation. In: Proceedings of the 19th International Conference on World Wide Web. pp. 661–670.
Lo, A.W., Wang, J., 2000. Trading volume: definitions, data analysis, and implications of portfolio theory. Rev. Financ. Stud. 13 (2), 257–300.
Markowits, H.M., 1952. Portfolio selection. J. Financ. 7 (1), 71–91.
Orlando, G., Bufalo, M., 2021. Empirical evidences on the interconnectedness between sampling and Asset Returns' Distributions. Risks 9 (5), 88. http://dx.doi.org/10.3390/risks9050088.
Patton, A.J., 2004. On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. J. Financ. Econom. 2 (1), 130–168. http://dx.doi.org/10.1093/jjfinec/nbh006.
Van Oordt, M.R., Zhou, C., 2016. Systematic tail risk. J. Financ. Quant. Anal. 51 (2), 685–705.
Wang, Y., Wei, Y., Wu, C., Yin, L., 2018. Oil and the short-term predictability of stock return volatility. J. Empir. Financ. 47, 90–104.
Welch, I., Goyal, A., 2008. A comprehensive look at the empirical performance of equity premium prediction. Rev. Financ. Stud. 21 (4), 1455–1508.
Zhang, Z., He, M., Zhang, Y., Wang, Y., 2021. Realized skewness and the short-term predictability for aggregate stock market volatility. Econ. Model. 103, 105614.

Further reading

Aboussalah, A.M., Lee, C.-G., 2020. Continuous control with stacked deep dynamic recurrent reinforcement learning for portfolio optimization. Expert Syst. Appl. 140, 112891.
Almahdi, S., Yang, S.Y., 2017. An adaptive portfolio trading system: A risk-return portfolio optimization using recurrent reinforcement learning with expected maximum drawdown. Expert Syst. Appl. 87, 267–279.
Bali, T.G., Cakici, N., 2004. Value at risk and expected stock returns. Financ. Anal. J. 60 (2), 57–73. http://dx.doi.org/10.2469/faj.v60.n2.2610.
Bali, T.G., Cakici, N., Levy, H., 2008. A model-independent measure of aggregate idiosyncratic risk. J. Empir. Financ. 15 (5), 878–896.
Ban, G.-Y., El Karoui, N., Lim, A.E., 2018. Machine learning and portfolio optimization. Manage. Sci 64 (3), 1136–1154.
Brock, W., Lakonishok, J., LeBaron, B., 1992. Simple technical trading rules and the stochastic properties of stock returns. J. Financ. 47 (5), 1731–1764.
Dash, R., Dash, P.K., 2016. A hybrid stock trading framework integrating technical analysis with machine learning techniques. J. Financ. Data Sci. 2 (1), 42–57.
Fama, E.F., 1970. Efficient capital markets. J. Financ. 25 (2), 383–417.
Friesen, G.C., Weller, P.A., Dunham, L.M., 2009. Price trends and patterns in technical analysis: A theoretical and empirical examination. J. Bank. Financ. 33 (6), 1089–1100.
Gu, S., Kelly, B., Xiu, D., 2021. Autoencoder asset pricing models. J. Econometrics 222 (1), 429–450.
Harvey, C.R., Siddique, A., 2000. Conditional skewness in asset pricing tests. J. Finance 55 (3), 1263–1295. http://dx.doi.org/10.1111/0022-1082.00247.
Jang, J., Seong, N., 2023. Deep reinforcement learning for stock portfolio optimization by connecting with modern portfolio theory. Expert Syst. Appl. 218, 119556.
Jiang, J., Kelly, B., Xiu, D., 2023. (Re-) Imag (in) ing Price Trends. J. Finance 78 (6), 3193–3249.
Kaczmarek, T., Będowska-Sójka, B., Grobelny, P., Perez, K., 2022. False safe haven assets: Evidence from the target volatility strategy based on recurrent neural network. Res. Int. Bus. Finance 60, 101610.
Ke, Z.T., Kelly, B.T., Xiu, D., 2019. Predicting Returns with Text Data. Technical Report, National Bureau of Economic Research.
Leigh, W., Purvis, R., Ragusa, J.M., 2002. Forecasting the NYSE composite index with technical analysis, pattern recognizer, neural network, and genetic algorithm: a case study in romantic decision support. Decis. Supp. Syst. 32 (4), 361–377.
Lo, A.W., Mamaysky, H., Wang, J., 2000. Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation. J. Financ. 55 (4), 1705–1765.
Luss, R., d'Aspremont, A., 2015. Predicting abnormal returns from news using text classification. Quant. Finance 15 (6), 999–1012.
Vo, N.N., He, X., Liu, S., Xu, G., 2019. Deep learning for decision making and the optimization of socially responsible investments and portfolio. Decis. Support Syst. 124, 113097.
Wu, M.-C., Lin, S.-Y., Lin, C.-H., 2006. An effective application of decision tree to stock trading. Expert Syst. Appl. 31 (2), 270–274.