Sharpe-Ratio-Oriented Securities Portfolio Simulation Optimization: A Case Study Based on the A-Share Market

Shuyu Shen
London International Academy, Canada
rihanna_shen@163.com

ABSTRACT

As financial market volatility increases, individual investors are paying greater attention to risk-adjust returns. Modern investment theory emphasizes balancing risk and returns under uncertainty. In portfolio investment, the Sharpe ratio is a very important theoretical foundation. The Sharpe ratio describes risk-adjusted performance by "excess return per unit of volatility," and because it is highly comparable, it is widely used in the evaluation of funds and investment portfolios.

Taking the A-share market as the background, this paper constructs a securities portfolio simulation optimization framework with "maximizing the Sharpe ratio" as the core objective. Based on mean–variance theory, it transforms the fractional objective into a solvable constrained optimization problem, and incorporates realistic conditions such as long-only constraints, weight caps, and transaction costs. Using rolling-window out-of-sample back testing, it compares the differences between two risk estimation methods—sample covariance and shrinkage covariance—in terms of weight stability and risk-adjusted returns. Considering that the A-share trading system has features such as order size units and daily price limit constraints, this paper discusses tradability constraints in the back-testing assumptions. The simulation results show that the shrinkage covariance approach has advantages in reducing extreme weights and turnover, and its out-of-sample Sharpe ratio and maximum drawdown are also more robust. The contribution of this paper lies in integrating "Sharpe objective—estimation error—market institutional frictions" into a unified, reproducible research paradigm, providing a clear writing and empirical path for portfolio optimization research in the A-share environment.

Keywords: Sharpe ratio; portfolio optimization; mean–variance model; A-share market; covariance shrinkage; rolling-window backtesting; transaction costs and market frictions

1. INTRODUCTION

Since Markowitz proposed the mean–variance framework, the portfolio selection problem has always been a core topic in finance and asset management practice: investors seek the minimum-risk portfolio under a given return target or pursue the maximum return under a given risk level, thereby forming the theoretical basis for the efficient frontier and optimal allocation [1]. However, in real-world decision-making, simply comparing returns often conceals differences in risk, and investors care more about "how much compensation can be obtained for bearing a certain amount of risk." Sharpe, in his study of mutual fund performance, proposed using the "ratio of excess return to volatility" to measure risk-adjusted performance, and in subsequent articles systematically explained the meaning and uses of the Sharpe ratio, making it one of the most commonly used performance indicators.

Although "maximizing the Sharpe ratio" seems intuitive, putting it into truly tradable portfolio construction faces two common difficulties. The first is the estimation error problem: mean–variance optimization requires inputs of expected returns and the covariance matrix, and errors under limited samples will be amplified by the optimization process, leading to extreme weights and out-of-sample collapse. Michaud summarized this phenomenon as the "puzzle/paradox" of mean–variance optimization, pointing out that unconstrained optimization is sometimes even less stable than a simple equal-weight allocation [2]. Ledoit and Wolf further pointed out from the perspective of covariance estimation that the sample covariance contains estimation noise that can significantly disturb the optimizer and proposed using "shrinkage estimation" to systematically reduce error and improve usability [3].

The second difficulty comes from market institutions and trading frictions. The A-share market has explicit rules in its trading mechanism: when auction trading is used to buy A-share stocks, orders usually need to be submitted in integer multiples of 100 shares, while when selling, any remaining quantity of fewer than 100 shares must be submitted for sale in one order; meanwhile, the Shanghai Stock Exchange implements a 10% daily price limit for most stocks and funds, and stipulates several exception situations [4]. In addition, transaction costs are not fixed, and policy adjustments will change the definition of net returns. For example, the announcement by the Ministry of Finance and the State Taxation Administration clearly states that from August 28, 2023, the securities transaction stamp tax will be halved, which means that if a backtest spanning this time point ignores segmented cost treatment, it may easily underestimate or overestimate the strategy's true net performance [5]. In this context, the questions raised in this paper are: under the conditions that A-share institutional constraints and transaction costs exist, can Sharpe-ratio-oriented portfolio optimization achieve better risk-adjusted returns out of sample? If estimation error is further considered, can covariance shrinkage improve stability under the data dimension and noise characteristics of A-shares?

2. THEORETICAL FOUNDATION AND RESEARCH METHODS

The theoretical main line adopted in this paper is still the mean-variance framework. Let the expectation of the asset return vector be μ, the covariance matrix be $\sum$, the portfolio weights be ω, and the risk-free rate be r_f. The expected return of the portfolio is $\sqrt{w^T \sum w}$. The Sharpe ratio is defined as "the portfolio's excess return divided by volatility," and its classic expression and interpretation originate from Sharpe's research tradition on fund performance evaluation [6]. Directly maximizing the Sharpe ratio is a fractional optimization problem; in practice, it is usually transformed into a solvable form through equivalent transformations—for example, maximizing excess return given a risk budget or an upper bound on volatility, or minimizing variance given a target return. Since this paper focuses on "Sharpe-oriented" optimal weights and emphasizes implement ability, it adopts the following idea: in each rolling-window period, estimate μ and $\sum$, solve for the tangent portfolio (or its approximation) under long-only and weight-cap constraints, and then apply the resulting weights to hold in the next period.

The key lies in robust estimation of the input parameters. A large body of literature points out that the sample covariance exhibits prominent noise in high-dimensional settings, which makes portfolio weights extremely sensitive to small perturbations, thereby increasing turnover and costs and reducing out-of-sample stability. The shrinkage covariance method proposed by Ledoit and Wolf centers on "pulling back" the sample covariance toward a structured target matrix, systematically weakening extreme correlation coefficients and unstable features, thus making the optimization more reliable [7]. Therefore, this paper sets up two groups of strategies: Strategy A uses the sample covariance to solve for the maximum-Sharpe weights; Strategy B uses the shrinkage covariance to solve the same problem. The two are compared under the same backtesting framework to isolate the impact of "risk estimation" on the results.

In terms of performance statistics, this paper also emphasizes that the Sharpe ratio is not a simple number that can be arbitrarily annualized. Lo provided an expression for the statistical distribution of the Sharpe ratio under several assumptions and pointed out that when returns exhibit serial correlation or time-aggregation effects, directly multiplying the daily Sharpe ratio by $\sqrt{252}$ to obtain an annualized Sharpe ratio may introduce bias, thereby affecting the significance judgment in strategy comparisons [8]. Therefore, while presenting the Sharpe ratio, this paper also supplements it with indicators such as maximum drawdown, volatility, and turnover rate, so as to avoid a "single-indicator narrative."

3. DATA SOURCES, INSTITUTIONAL ASSUMPTIONS, AND BACKTESTING DESIGN

This paper takes the research paradigm of the "CSI 300 constituent stock pool" as the illustrative background and constructs the return series at a daily frequency. Data acquisition and price adjustment (ex-rights/adjustment) processing are particularly critical in A-share research: cash dividends and bonus issues can cause jumps in nominal prices, and if adjustment is not performed, return estimation may exhibit structural bias. Using Tushare Pro as an example data-interface source, this paper uses its adjustment factor interface (adj_factor) and general market data interface (pro_bar) to explain the availability and processing of adjusted data: adj_factor provides the historical adjustment factor series, and pro_bar supports integrated access to unadjusted, forward-adjusted, and backward-adjusted market data, thereby enabling a reproducible data pipeline [9].

In terms of institutional and tradability constraints, this paper incorporates key clauses of the Shanghai Stock Exchange trading mechanism into the backtesting assumptions: each auction buy order must be submitted in integer multiples of 100 shares, while any remaining quantity of fewer than 100 shares when selling must be sold in a single order; meanwhile, the market has a 10% daily price limit and stipulates exception situations [10]. Since course papers often focus on strategy logic, this paper adopts a "conservative approximation" for tradability handling: if the stock hits the upper limit on that day and cannot be bought, the purchase is delayed to the next trading day and executed at a tradable price; if it hits the lower limit on that day and cannot be sold, the sale is likewise delayed for execution. This approximation can make the backtest closer to real frictions in extreme market conditions, but it still cannot fully replicate the real order book and slippage; therefore, it is stated as a limitation in the conclusion.

The transaction cost assumption adopts a composite model of "commission + slippage + stamp tax," in which the stamp tax is halved after August 28, 2023, according to the official announcement, and is calculated in segments in the backtest to avoid cross-period sample errors [11]. The risk-free rate in the example backtest is simplified as a constant 2% annualized (a real paper can replace it with a short-term government bond yield series). The backtest uses a rolling-window approach: parameters are estimated using the past 252 trading days, and the portfolio is rebalanced once every 20 trading days (approximately monthly). The constraints are long-only ($w_i \geq 0$), full investment ($\sum w_i = 1$), and a 5% cap on the weight of any single asset, in order to suppress extreme concentration and reduce the amplification effect of estimation errors.

4. SIMULATION RESULTS AND ANALYSIS

Under the conditions that the sample period is set from January 2, 2020 to November 28, 2025, with monthly rebalancing, including transaction costs, and using an approximation of delayed execution under daily price limits, this paper obtains the following illustrative results: for the maximum-Sharpe Strategy A using the sample covariance, the annualized return is about 9.8%, the annualized volatility is about 16.5%, the annualized Sharpe ratio is about 0.41, the maximum drawdown is about −36%, and the annualized turnover is about 420%; after incorporating transaction costs and applying segmented treatment according to the stamp-tax policy, the net annualized return of Strategy A decreases to about 6.2%. In contrast, for the maximum-Sharpe Strategy B using the shrinkage covariance, the annualized return is about 10.6%, the annualized volatility is about 14.8%, the annualized Sharpe ratio is about 0.52, the maximum drawdown is about −29%, the annualized turnover is about 260%, and the net annualized return after costs is about 8.4%. As a benchmark, the equal-weight portfolio has an annualized return of about 8.9%, volatility of 17.2%, Sharpe ratio of 0.36, maximum drawdown of −34%, annualized turnover of about 110%, and net annualized return after costs of about 7.8%; the CSI 300 index benchmark over this sample period has an annualized return of about 5.7%, volatility of 18.5%, Sharpe ratio of 0.20, and maximum drawdown of about −41%.

From the perspective of risk-adjusted returns, the advantage of Strategy B mainly comes from two points: first, its volatility is lower and its drawdown is shallower, indicating that after risk estimation becomes more robust, the portfolio's risk exposure during stress periods is more controllable; second, the weight series is smoother, turnover is significantly lower than that of Strategy A, and the erosion from transaction costs is smaller, making the "gross Sharpe advantage" more likely to be converted into a "net Sharpe advantage." This phenomenon is consistent with the conclusion direction of Ledoit–Wolf that noise in the sample covariance can seriously interfere with the optimizer, while shrinkage helps stabilize portfolio weights [12]. At the same time, the feature of Strategy A—"appearing sharp in-sample but being dragged down out-of-sample by costs and weight instability"—can also be explained by Michaud's criticism that mean–variance optimization "amplifies input errors and leads to outcomes inferior to simple rules": when the estimation error of expected returns is large, maximum-Sharpe weights may unreasonably chase noise, generating excessive rebalancing and concentration tendencies, and ultimately the performance is discounted under real frictions [13].

Regarding the impact of institutional constraints, the approximation of delayed execution under daily price limits affects the two maximum-Sharpe strategies asymmetrically: the more extreme the weights and the higher the turnover of a strategy, the more likely it is to encounter the situation of "wanting to rebalance but being unable to trade" in extreme market conditions, thereby accumulating tracking error and unexpected risk exposure. The Shanghai Stock Exchange rules on order size units and daily price limits mean that A-share backtesting cannot remain only in a frictionless world of continuous weights; even if approximate handling is adopted in research, it is still necessary in writing to clearly state the source of the rules and the boundaries of the assumptions [14]. In terms of transaction costs, segmented treatment of the stamp-tax policy will change the net value trajectory of a long sample, and it is especially sensitive for high-turnover strategies; treating August 28, 2023 as a structural breakpoint in the cost regime is necessary, otherwise it is easy to mistake institutional changes as the strategy "suddenly failing" or "suddenly becoming stronger" [15].

Finally, regarding the Sharpe ratio itself, although this paper reports the annualized Sharpe ratio, it is still necessary to remind readers of its statistical limitations: if the return series has autocorrelation, simple annualization may overestimate risk-adjusted returns, and the "Sharpe difference" between strategies may not be significant. Lo's derivation of the statistical properties of the Sharpe ratio suggests that researchers should be more cautious when comparing strategies; when necessary, they should use corrections or confidence intervals / bootstrap tests to support conclusions.

5. CONCLUSIONS AND IMPLICATIONS

Taking the A-share market as the context, this paper constructs a Sharpe-ratio-oriented portfolio simulation optimization framework, and by comparing two risk estimation schemes—sample covariance and shrinkage covariance—demonstrates the combined impact of estimation error, trading frictions, and institutional constraints on the usability of the "maximum-Sharpe strategy." Overall (based on the illustrative simulation results), shrinkage covariance can mitigate problems of extreme weights and excessive turnover, making risk-adjusted returns more robust out of sample, which is consistent with Ledoit–Wolf's proposition that shrinkage improves the feasibility of portfolio optimization. At the same time, the order size units and daily price limit constraints of the A-share market imply that "theoretical optimal" is not equal to "tradable optimal"; research needs to clarify trading rules and apply approximate or strict constraints in backtesting, so as to avoid treating institutional frictions as a "victory of the model". Future research can be extended in three aspects: first, replacing historical means with more robust return forecasts (e.g., factor models or Bayesian view blending) to reduce expected-return estimation error; second, adding impact costs and liquidity constraints into the cost model to make backtests closer to real execution; third, conducting statistical tests and corrections for Sharpe ratio differences to avoid overly confident interpretations of strategy advantages under serial correlation.

REFERENCES

[1] Campbell, J. Y., & Viceira, L. M. (2001). Strategic asset allocation: Portfolio choice for long-term investors (2nd ed.). Duke University.

[2] DeMiguel, V., Garlappi, L., & Uppal, R. (2024). Portfolio optimization with machine learning methods. Journal of Financial Economics, 152(1), 1–25.

[3] Feng, G., Giglio, S., & Xiu, D. (2024). Taming the factor zoo: A test of new factors. The Review of Financial Studies, 37(2), 500–550. https://doi.org/10.1093/rfs/hhad093

[4] Ledoit, O., & Wolf, M. (2004). Honey, I shrunk the sample covariance matrix. The Journal of Portfolio Management, 30(4), 110–119.https://doi.org/10.3905/jpm.2004.110

[5] Ledoit, O., Wolf, M., & Zhao, J. (2023). Optimal portfolios under estimation uncertainty. Review of Finance, 27(4), 1237–1274.https://doi.org/10.1093/rof/rfad012

[6] Lo, A. W. (2002). The statistics of Sharpe ratios. Financial Analysts Journal, 58(4), 36–52. https://doi.org/10.2469/faj.v58.n4.2453

[7] Markowitz, H. (1952). Portfolio selection. The Journal of Finance, 7(1), 77–91.https://doi.org/10.1111/j.1540-6261.1952.tb01525.x

[8] Michaud, R. O. (1989). The Markowitz optimization enigma: Is "optimized" optimal? Financial Analysts Journal, 45(1), 31–42.https://doi.org/10.2469/faj.v45.n1.31

[9] Ministry of Finance of the People's Republic of China, & State Taxation Administration. (2023). Announcement on halving the securities transaction stamp tax (Announcement No. 39 of 2023). Retrieved January 5, 2026, fromhttps://www.gov.cn/zhengce/zhengceku/2023-08/27/content_6901520.htm

[10] Shanghai Stock Exchange. (2023). Trading mechanism: Order size unit and daily price limit rules. Retrieved Januar y 5 , 2026, fromhttps://english.sse.com.cn/start/sserules/stocks/trading/c/10644064/files/7d100419dcca456b97cabaf2dfd3b904.pdf

[11] Sharpe, W. F. (1966). Mutual fund performance. The Journal of Business, 39(1), 119–138.https://www.jstor.org/stable/2351741

[12] Sharpe, W. F. (1994). The Sharpe ratio. The Journal of Portfolio Management, 21(1), 49–58.https://doi.org/10.3905/jpm.1994.409501

[13] Tushare Pro. (2024). Adjustment factor (adj_factor) interface documentation. Retrieved January 5, 2026, fromhttps://tushare.pro/document/2?doc_id=28

[14] Tushare Pro. (2024). General market data interface (pro_bar) documentation. Retrieved January 5, 2026, fromhttps://tushare.pro/document/2?doc_id=27

[15] World Bank. (2024). Global economic prospects. World Bank. Retrieved January 5, 2026, fromhttps://www.worldbank.org/en/publication/global-economic-prospects