

Article

Dynamic Portfolio Construction in High-Frequency Markets Using Microstructure-Aware Deep Temporal Models

Xiaoyi Meng ^{1,*}

¹ University of California, Irvine, 92618, USA

* Correspondence: Xiaoyi Meng, University of California, Irvine, 92618, USA

Abstract: High-frequency financial markets exhibit rapidly changing liquidity and strong nonstationarity, which reduce the effectiveness of portfolio methods based on static return and covariance estimates. This study presents a microstructure-aware framework for dynamic portfolio allocation that maps high-frequency market states directly to portfolio weights. The model uses order-flow intensity, depth imbalance, and volume asymmetry as key inputs and combines recurrent and convolutional temporal components to capture both persistent dependencies and short-lived microstructure shocks. Portfolio weights are determined under a joint return–risk objective. The framework is evaluated on ten years of high-frequency U.S. equity data. Results show higher risk-adjusted performance than mean–variance and risk-parity benchmarks, with a 14%–22% improvement in Sharpe ratio and more stable drawdown behavior. These findings suggest that integrating market microstructure information into temporal allocation models can improve exposure management in intraday portfolio optimization.

Keywords: high-frequency trading; portfolio construction; market microstructure; deep temporal models; order flow; risk-adjusted performance

1. Introduction

Portfolio construction determines how information extracted from financial markets is transformed into trading positions and, ultimately, into realized returns. In most portfolio optimization frameworks, this transformation relies on estimates of expected returns and risk that are assumed to be sufficiently stable over the rebalancing horizon. Such assumptions are reasonable when prices are sampled at coarse intervals and market conditions evolve slowly. At very short horizons, however, they become increasingly fragile. In high-frequency trading environments, returns are shaped by transient liquidity conditions, discrete order submissions, and pronounced intraday patterns. Empirical studies show that dependence in returns and volatility at sub-minute horizons decays rapidly, often within a limited number of observations. As a result, return and covariance estimates based on recent data are highly sensitive to noise, and portfolio weights derived from these estimates can vary substantially over time [1,2]. This instability poses a fundamental challenge for applying conventional portfolio optimization methods in high-frequency settings.

Evidence from recent empirical work suggests that short-horizon price movements are closely linked to market microstructure rather than price dynamics alone. Information contained in the limit order book, such as order-flow imbalance, queue imbalance, and depth asymmetry, has been shown to explain part of short-term price variation beyond what is captured by past returns [3,4]. Importantly, the contribution of these signals is not

Published: 04 March 2026



Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

constant. Their predictive strength varies with volatility, trading activity, and market depth, indicating that microstructure effects are strongly state dependent [5]. These findings imply that portfolio strategies operating at high frequency must adapt to changing market conditions rather than rely on fixed relationships. To capture such time-varying behavior, deep learning models have been increasingly adopted in high-frequency financial analysis. Recurrent neural networks, particularly long short-term memory (LSTM) models, are commonly used to represent temporal dependence and intraday structure, and often outperform linear benchmarks in short-horizon forecasting tasks [6]. Convolution-based temporal models have also attracted attention, as they are effective at extracting local patterns and abrupt changes from high-frequency sequences. Temporal Convolutional Networks (TCNs), which use causal and dilated convolutions, have been reported to provide more stable forecasts than standard recurrent architectures in several multiscale settings [7,8]. Attention mechanisms are frequently combined with these models to control the influence of noisy or redundant features when the input space is high dimensional [9]. Despite these modeling advances, most existing studies focus on predictive accuracy rather than portfolio performance. Forecast quality is typically evaluated using statistical criteria, while the link between predictions and portfolio allocation is treated in a simplified manner. In practice, improvements in prediction accuracy do not automatically translate into superior investment outcomes. Transaction costs, portfolio turnover, and drawdown risk can substantially offset predictive gains. Empirical comparisons show that prediction-based allocation rules may fail to outperform classical benchmarks, such as minimum-variance or risk parity portfolios, in out-of-sample tests, especially over longer horizons [10,11]. Several recent studies attempt to address this gap by incorporating portfolio objectives directly into the learning process. End-to-end frameworks embed risk measures, such as variance or conditional value at risk, into model training and report improvements in risk-adjusted returns relative to traditional optimization approaches [12,13]. Reinforcement learning methods further formulate portfolio rebalancing as a sequential decision problem and have shown promising results in simulated environments and low-frequency applications [14,15]. However, their effectiveness in high-frequency equity markets remains limited. These methods typically require large training samples, are sensitive to reward specification, and face difficulties in separating predictive signals from execution-related noise, which reduces their reliability in long-horizon empirical evaluations.

This study focuses on the problem of dynamic portfolio construction in high-frequency markets. We propose a deep learning framework that integrates market microstructure information into portfolio decision-making. The model combines LSTM components to capture longer-term temporal dependencies with Temporal Convolutional Network modules to extract localized patterns and abrupt order-flow changes. An attention mechanism allows the model to adjust the contribution of different signals as market conditions evolve. Portfolio weights are generated through a joint return–risk objective, aligning model training with economically relevant performance measures rather than purely statistical criteria. The proposed framework is evaluated on a decade-long U.S. high-frequency dataset, with performance assessed in terms of Sharpe ratio, maximum drawdown, and return stability, and compared against classical mean-variance and risk-based allocation benchmarks. The results provide evidence on how microstructure-aware temporal modeling can improve portfolio performance in high-frequency trading environments.

2. Materials and Methods

2.1. Sample and Study Scope

The empirical analysis uses high-frequency equity data from the U.S. stock market over a ten-year period from January 2014 to December 2023. The sample includes 50 large-cap stocks selected from the S&P 500 index to ensure continuous trading activity and sufficient market depth. Transaction records and full limit order book snapshots are collected at one-second intervals during regular trading hours (09:30–16:00 EST). After

standard data cleaning procedures, the final dataset contains approximately 1.2×10^9 trade observations and 3.6×10^9 order book records. The selected assets cover multiple industry sectors and display different volatility levels, allowing the analysis of portfolio behavior under diverse microstructure conditions.

2.2. Experimental Design and Benchmark Construction

The study follows a controlled portfolio comparison framework. The deep learning-based allocation model is treated as the experimental strategy, while mean-variance optimization and risk parity serve as benchmark methods. All strategies rely on the same input data, rebalancing schedule, and transaction cost assumptions. Portfolio weights are updated at fixed intraday intervals using rolling estimation windows. The benchmark strategies are implemented using conventional covariance estimation and volatility scaling techniques commonly reported in empirical finance studies. This design ensures that performance differences are attributable to modeling structure rather than data access or trading rules.

2.3. Measurement Procedures and Quality Control

Market microstructure indicators are computed directly from raw trade and order book data. Order-flow intensity is defined as the signed difference between buyer-initiated and seller-initiated volumes. Market depth measures are calculated using aggregated bid and ask quantities across multiple price levels. To reduce the influence of microstructure noise, all input variables are standardized using rolling window normalization for each asset. Data quality control includes the removal of abnormal trading days, correction of timestamp inconsistencies, and exclusion of intervals affected by trading halts. The proportion of missing observations remains below 0.5% and does not materially affect the empirical results.

2.4. Data Processing and Model Formulation

Input features are organized into overlapping time windows to capture both short-term market movements and longer temporal dependencies. Let $\mathbf{x}_t \in \mathbb{R}^d$ denote the feature vector at time t . Portfolio returns are computed as

$$R_{t+1} = \mathbf{w}_t^\top \mathbf{r}_{t+1}$$

where \mathbf{w}_t represents the portfolio weight vector and \mathbf{r}_{t+1} denotes the asset return vector. Model estimation is guided by a joint return-risk objective function, defined as

$$L = -E[R_{t+1}] + \lambda \text{Var}(R_{t+1})$$

where λ controls the relative importance of return and risk. Model parameters are estimated using gradient-based optimization. Hyperparameters are fixed across assets to limit overfitting.

2.5. Performance Evaluation and Statistical Analysis

Portfolio performance is evaluated using annualized Sharpe ratio, maximum drawdown, and return volatility. Statistical inference is conducted using block bootstrap methods to account for serial dependence in high-frequency returns. For each strategy, 1,000 bootstrap replications are generated to construct 95% confidence intervals. Portfolio turnover and transaction costs are also reported to assess practical feasibility. Performance measures are aggregated across assets and time periods to reduce the influence of isolated market events.

3. Results and Discussion

3.1. Risk-Adjusted Performance over the Full Out-of-Sample Horizon

Over the ten-year U.S. high-frequency backtest, the proposed dynamic allocation framework shows a persistent advantage over mean-variance and risk-parity baselines when evaluated on risk-adjusted metrics and downside behavior. The Sharpe-ratio lift (14%–22%) is accompanied by a more stable return path, which indicates that the gain is not driven by a small number of extreme episodes but by repeated improvements in day-

to-day risk efficiency. This pattern is in line with recent portfolio-RL evidence showing that strategy quality depends not only on reward design but also on how the policy forms a market “state” that remains informative under distribution shifts and changing correlations [16]. The separation between adaptive and static allocations is typically largest during liquidity-stressed intervals, when covariance estimates become unstable and microstructure conditions dominate short-horizon dynamics, as illustrated in Figure 1.

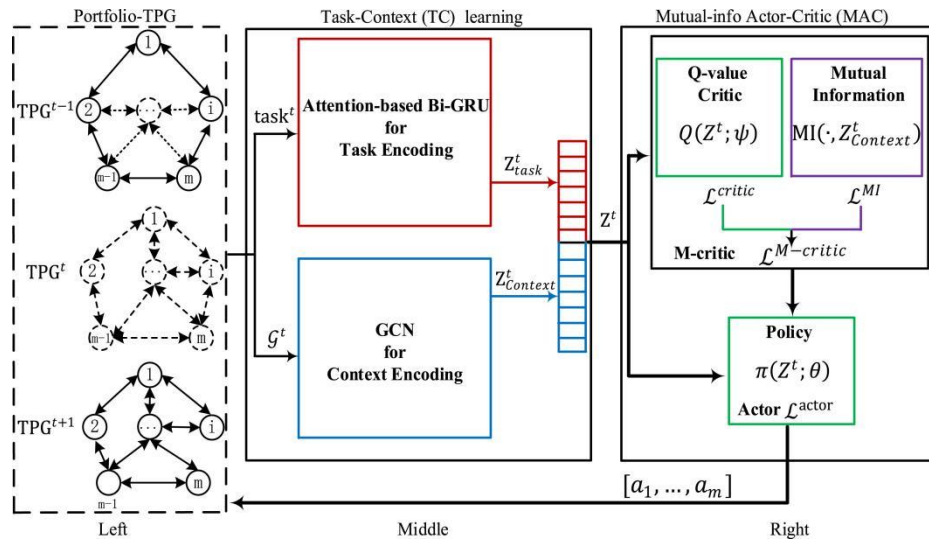


Figure 1. Policy and representation-learning pipeline for portfolio management.

3.2. Economic Meaning of Microstructure Features and Signal Stability

The results indicate that microstructure conditioning improves portfolio behavior in ways that are economically interpretable. When order-flow intensity, depth imbalance, and volume asymmetry are included, the learned weights react less to transient price noise and more to shifts in liquidity supply–demand, which improves net performance after costs by lowering “unpaid” turnover. This is consistent with the broader limit-order-book (LOB) literature, which reports that LOB-based representations carry incremental information beyond trade-time bars and that model performance can deteriorate when feature representations do not generalize across regimes [17]. In portfolio terms, the most meaningful change is not higher raw volatility-taking, but better timing of exposure reductions during adverse order-flow conditions, which directly connects microstructure signals to drawdown control rather than only to point forecasts.

3.3. Architecture-Level Effects under High-Frequency Noise

Comparisons across architectural variants suggest that the benefit does not come from a single temporal block but from complementary modeling of short bursts and slower dependence. TCN-style components capture short-lived microstructure shocks and local trend fragments, while recurrent components better reflect persistence in order-flow regimes and intraday structure; attention improves the weight-update quality by concentrating capacity on the most informative slices, which reduces sensitivity to noisy quote updates, as illustrated in Figure 2.

		$M(T)$					
<i>side</i> →		<i>ask</i>			<i>bid</i>		
<i>volume</i> →		V_1	...	V_L	V_1	...	V_L
<i>price</i> →		P_1	...	P_L	P_1	...	P_L
<i>time</i>	$t - h$	$v_1(t - h)$...	$v_L(t - h)$	$v_1(t - h)$...	$v_L(t - h)$
		$p_1(t - h)$...	$p_L(t - h)$	$p_1(t - h)$...	$p_L(t - h)$
	$t - h + 1$	$v_1(t - h + 1)$...	$v_L(t - h + 1)$	$v_1(t - h + 1)$...	$v_L(t - h + 1)$
		$p_1(t - h + 1)$...	$p_L(t - h + 1)$	$p_1(t - h + 1)$...	$p_L(t - h + 1)$
	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
		\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
	t	$v_1(t)$...	$v_L(t)$	$v_1(t)$...	$v_L(t)$
		$p_1(t)$...	$p_L(t)$	$p_1(t)$...	$p_L(t)$

Figure 2. Market observation construction from LOB sequences.

3.4. Comparison with Prior Work and Practical Constraints

Relative to recent studies that emphasize prediction accuracy, these results support an end-to-end view in which representation learning is evaluated by downstream allocation quality under realistic frictions. Prior work on deep RL for portfolio management stresses that adding global context representations can improve policy quality and transfer, but also notes that performance can be fragile when the state definition is weak or when market conditions shift [18]. In a high-frequency context, the same concern becomes sharper: capacity limits, latency, and execution costs can compress the apparent edge, so the practical value of microstructure-driven allocation depends on controlling turnover and aligning the rebalancing cadence with signal half-life. The most defensible takeaway is therefore conditional: deep temporal structures add value when they improve state stability and reduce unnecessary trading, but the strategy should be validated under stricter execution models, cross-venue tests, and regime-specific stress scenarios before being treated as deployable at scale.

4. Conclusion

This study investigated dynamic portfolio construction in high-frequency trading settings using a market microstructure-aware learning framework that links observed market states directly to portfolio weights. By incorporating order-flow intensity, depth imbalance, and volume asymmetry, together with complementary temporal modeling components, the proposed approach achieved consistently higher risk-adjusted performance than mean-variance and risk-parity benchmarks over ten years of U.S. equity data. The results show that the improvement is mainly driven by smoother exposure adjustment under changing liquidity conditions, rather than by increased risk taking. This finding highlights the importance of stable state representation when return distributions are time-varying and short-horizon signals decay rapidly. From a methodological perspective, the study contributes evidence that combining microstructure information with structured temporal modeling can improve portfolio decisions in nonstationary environments. From an application perspective, the framework provides a practical reference for intraday portfolio strategies that seek to balance responsiveness and turnover. Nevertheless, the analysis remains subject to several limitations, including dependence on market-specific microstructure characteristics, simplified execution assumptions, and potential capacity effects under large trade sizes. Further work should therefore examine cross-market robustness, incorporate explicit execution and impact modeling, and assess performance under extreme stress conditions to better evaluate the general applicability of microstructure-driven portfolio allocation methods.

References

1. P. N. Kolm, J. Turiel, and N. Westray, "Deep order flow imbalance: Extracting alpha at multiple horizons from the limit order book," *Mathematical Finance*, vol. 33, no. 4, pp. 1044-1081, 2023. doi: 10.1111/mafi.12413
2. L. Zhang, and L. Hua, "Major issues in high-frequency financial data analysis: A survey of solutions," *Mathematics*, vol. 13, no. 3, p. 347, 2025. doi: 10.3390/math13030347
3. A. Briola, S. Bartolucci, and T. Aste, "Deep limit order book forecasting: a microstructural guide," *Quantitative Finance*, vol. 25, no. 7, pp. 1101-1131, 2025. doi: 10.1080/14697688.2025.2522911
4. Z. Eisler, J. P. Bouchaud, and J. Kockelkoren, "The price impact of order book events: market orders, limit orders and cancellations," *Quantitative Finance*, vol. 12, no. 9, pp. 1395-1419, 2012.
5. N. Hautsch, L. M. Kyj, and P. Malec, "The merit of high-frequency data in portfolio allocation (No. 2011-059)," *SFB 649 Discussion Paper*, 2011.
6. S. Giantsidi, and T. Claudia, "Deep Learning for Financial Forecasting: A Review of Recent Advancements," Available at SSRN 5263710. doi: 10.2139/ssrn.5263710
7. P. Lara-Benítez, M. Carranza-García, J. M. Luna-Romera, and J. C. Riquelme, "Temporal convolutional networks applied to energy-related time series forecasting," *applied sciences*, vol. 10, no. 7, p. 2322, 2020. doi: 10.3390/app10072322
8. Q. Zhao, H. Li, Q. Zhang, and Y. Wang, "A Study on Metal Futures Price Prediction Based on Piecewise Cubic Bézier Filtering for TCN," *Applied Sciences*, vol. 15, no. 17, p. 9792, 2025. doi: 10.3390/app15179792
9. Y. Guo, and X. Chen, "Forecasting the mid-price movements with high-frequency lob: A dual-stage temporal attention-based deep learning architecture," *Arabian Journal for Science and Engineering*, vol. 48, no. 8, pp. 9597-9618, 2023. doi: 10.1007/s13369-022-07197-3
10. K. Kubo, and K. Nakagawa, "Portfolio optimization using deep learning with risk aversion utility function," *Finance Research Letters*, vol. 74, p. 106761, 2025. doi: 10.1016/j.frl.2025.106761
11. J. Zou, F. Wang, and Y. Wu, "Large portfolio allocation using high-frequency financial data," 2018. doi: 10.4310/sii.2018.v11.n1.a12
12. Z. Zhang, S. Zohren, and S. Roberts, "Deep learning for portfolio optimization," *arXiv preprint arXiv:2005.13665*, 2020. doi: 10.3905/jfds.2020.1.042
13. O. Ali, "Deep Learning for Portfolio Optimization: An AI-Driven Approach to Risk-Adjusted Returns," Available at SSRN 5510579.
14. X. Y. Liu, Z. Xia, J. Rui, J. Gao, H. Yang, M. Zhu, and J. Guo, "Finrl-meta: Market environments and benchmarks for data-driven financial reinforcement learning," *Advances in Neural Information Processing Systems*, vol. 35, pp. 1835-1849, 2022.
15. N. Passalis, A. Tefas, J. Kannianen, M. Gabbouj, and A. Iosifidis, "Temporal logistic neural bag-of-features for financial time series forecasting leveraging limit order book data," *Pattern Recognition Letters*, vol. 136, pp. 183-189, 2020. doi: 10.1016/j.patrec.2020.06.006
16. S. Kim, H. Kwon, J. Yoon, and H. Y. Kim, "Dynamic factor-informed reinforcement learning for enhancing portfolio optimization," *Financial Innovation*, vol. 12, no. 1, p. 3, 2026. doi: 10.1186/s40854-025-00803-x
17. R. Paki, and H. Abbasimehr, "PB-TABL: Task Incremental Learning Strategy via Applying Piggyback Architecture on Temporal Attention-Augmented Bilinear Networks for Financial Time-Series Classification," *IEEE Transactions on Big Data*, 2025. doi: 10.1109/tbdata.2025.3598730
18. Z. Wang, B. Huang, S. Tu, K. Zhang, and L. Xu, "Deeptrader: a deep reinforcement learning approach for risk-return balanced portfolio management with market conditions embedding," In *Proceedings of the AAAI conference on artificial intelligence*, May, 2021, pp. 643-650. doi: 10.1609/aaai.v35i1.16144

Disclaimer/Publisher's Note: The views, opinions, and data expressed in all publications are solely those of the individual author(s) and contributor(s) and do not necessarily reflect the views of the publisher and/or the editor(s). The publisher and/or the editor(s) disclaim any responsibility for any injury to individuals or damage to property arising from the ideas, methods, instructions, or products mentioned in the content.