Proceedings of ICFTBA 2026 Symposium: Strategic Management & Business Analytics: Data-Driven Performance Models
DOI: 10.54254/2754-1169/2026.LD35881

Do Technical Indicators Generate Tradable Value? Evidence from Market States, Price-Path Images, and Transaction Costs

Zifan Wang

FinTech, School of Finance, Central University of Finance and Economics, Beijing, China
hplwzf@126.com

Abstract. This paper examines whether technical information can generate economically meaningful stock portfolio returns under transaction cost constraints. Using daily data for CSI 300 constituent stocks from April 2, 2019, to December 30, 2025, this paper compare basic price-volume variables, traditional technical indicators, market-state variables, StockJump and VPIN microstructure variables, OHLC image CNN, and Raw OHLCV CNN models. Instead of relying mainly on classification metrics such as AUC and Accuracy, this paper evaluate models through out-of-sample quintile portfolios, focusing on long-short returns, Sharpe ratios, turnover, and transaction-cost-adjusted performance. The results show that traditional technical indicators, market-state variables, and microstructure variables provide limited incremental value beyond basic price-volume information. The 1-day prediction task exhibits some cross-sectional sorting ability, but its returns are largely eroded by high turnover after 10 bps transaction costs. In contrast, the 5-day OHLC image CNN retains a 9.34% long-short net annualized return and a 0.761 net Sharpe after costs, outperforming structured feature models and Raw OHLCV CNN. These findings suggest that technical information should be evaluated by portfolio-level economic value, and that image-based price-path representation better captures nonlinear information over longer horizons.

Keywords: Technical indicators, Market states, OHLC images, Convolutional neural networks, Transaction costs

1. Introduction

Technical indicators have long been widely used in financial markets for stock screening, market timing, and quantitative trading. Recent studies show that technical indicators may still contain predictive information: Shi et al. [1] re-examine their ability to forecast stock returns, Mostafavi et al. [2] compare indicators across trend, momentum, volatility, and volume dimensions, and Wang et al. [3] integrate technical indicators with chip factors and news. However, for quantitative investment, the key question is not only whether technical indicators improve prediction accuracy, but whether they can generate stable portfolio returns after transaction costs.

DOI: 10.54254/2754-1169/2026.LD35881

Existing research leaves three gaps. First, many studies rely mainly on Accuracy, AUC, or other classification metrics, while less attention is paid to cross-sectional sorting ability and transaction-cost-adjusted portfolio value. Neuhierl et al. [4] argue that probability forecasts should be evaluated within asset-pricing and portfolio-sorting frameworks. Second, traditional technical indicators are constructed from price and volume information and may lose nonlinear information embedded in price paths. Third, return predictability may vary across market environments. Cakici et al. [5] and Horvath et al. [6] show that market conditions and regime structures matter for predictive relationships.

This paper builds a unified framework to compare different layers of technical information. Starting from basic price-volume variables, this paper sequentially add traditional technical indicators, market-state variables, and microstructure variables such as Jump and VPIN to form Set1 to Set4. This paper then transform the past 20 trading days of OHLC price paths into image representations and use a CNN to learn nonlinear price-path patterns. A Raw OHLCV CNN is also included as a control model to distinguish the effect of deep learning architecture from the effect of image-based price representation.

The empirical results show that structured technical information has some cross-sectional sorting ability under zero transaction costs, but its incremental improvement over basic price-volume variables is limited. In contrast, the OHLC image CNN delivers stronger portfolio performance, while Raw OHLCV CNN performs worse, suggesting that the gain comes mainly from image-based price-path representation rather than CNN architecture alone. After 10 bps transaction costs, 1-day prediction profits are largely eroded by high turnover, whereas the 5-day OHLC image CNN maintains positive long-short net returns and Sharpe ratios.

This paper contributes in three ways. First, it compares basic price-volume information, traditional technical indicators, market states, microstructure variables, and image-based price-path representation in one framework. Second, it shifts the evaluation focus from prediction accuracy to portfolio economic value under transaction costs. Third, by comparing OHLC image CNN with Raw OHLCV CNN, it shows that price-path images provide a useful nonlinear representation of technical information for stock return prediction.

2. Literature review

2.1. Stock return prediction and probability forecasting

Return prediction has been the primary problem in both asset pricing and quantitative investing, whose basic idea is to obtain effective signals from high-dimensional data that are useful for explaining or predicting returns' cross-sectional variation. According to Kelly et al. [7], in the return prediction setting, nonlinearities and interrelationships in the data may contain information which cannot be captured by linear models, hence theoretically motivating the usage of complex models. On the other hand, Cakici et al. [5] and Hanauer et al. [8] investigate machine learning applications in cross-sectional return prediction in international markets and emerging markets, respectively, illustrating that return predictability is highly dependent on market settings and sample sensitiveness. From the methodology viewpoint, Neuhierl et al. [4] employ probability forecasting viewpoint, stressing the fact that predictions made from models could be employed directly into portfolio sorting and asset pricing setting, instead of only being the target of classification accuracy assessment. Chen et al. [9] further show that the model design and out-of-sample performance are key elements for machine learning based return predictions, instead of solely focusing on prediction accuracy in fixed test sets. Inspired by these studies, this paper uses the model-output upward

DOI: 10.54254/2754-1169/2026.LD35881

probability as the basis for cross-sectional sorting, evaluates the predictive ability of different models by constructing Long-Short portfolios, and further examines the economic value under transaction cost constraints, rather than using classification metrics such as Accuracy and AUC as the sole criteria for model quality.

2.2. Technical indicators and technical analysis

Technical indicators are common in stock selection, market timing, and quantitative investment. Research has been done to show that there is still information left in technical indicators: Shi et al. [1] explore the predictability of return by technical indicators; Mostafavi et al. [2] classify the technical indicators into trend, momentum, volatility, and volume indicators; Wang et al. [3] consider technical indicators along with chip factors and news. Albahli et al. [10] make use of the DenseNet model to integrate technical indicators for stock prediction and prove that technical indicators can be mined for nonlinear information by deep learning. But for quantitative investment, what needs to be considered is whether technical indicators can earn a profit through stable return on the portfolio. This paper therefore evaluates technical information through portfolio performance rather than classification accuracy alone.

2.3. Application of machine learning and deep learning in financial prediction

Machine learning and deep learning methods are increasingly used in financial prediction because they can model high-dimensional nonlinear relationships and complex feature interactions. Saberironaghi et al. [11] and Giantsidi et al. [12] give an overview of their applications in the prediction of stock prices and stress the importance of considering the properties of the used datasets and evaluation procedure for models' performance. According to Wang et al. [13], it is not complexity that plays a key role, but rather consistency with the task. Therefore, this paper uses CNN not for model complexity itself, but to learn local morphological patterns from OHLC price paths. Raw OHLCV CNN is included as a control to separate the effect of image representation from CNN architecture.

2.4. Market states and state dependence

An increasing number of studies suggest that stock markets do not always operate in a uniformly stable environment, and that return formation mechanisms, risk levels, and investor behavior may change across different market states. Kijkarncharoensin et al. [14] and Gupta et al. [15] identify market states using Hidden Markov Models and related frameworks, while Horvath et al. [6] use Wasserstein-distance-based clustering to show that market structures are non-stationary. These studies motivate the inclusion of market-state variables in return prediction. In this paper, however, market states are used as one additional feature group in Set3 rather than as a full regime-switching analysis.

2.5. Positioning of this paper relative to existing literature

In summary, existing literature has extensively discussed the return prediction problem from multiple perspectives including stock return prediction, technical indicators, machine learning, market states, and model interpretability, but there remains room for further improvement. First, traditional technical indicators, market state variables, microstructure variables, and price path image representations have rarely been systematically compared within a unified framework, and

DOI: 10.54254/2754-1169/2026.LD35881

the marginal incremental value of different levels of technical information still lacks consistent conclusions. Second, the currently existing literature tends to utilize prediction accuracy or costless predictiveness as the dominant criteria of performance evaluation, paying less attention to transaction costs, portfolio turnover rate, and the economic values of the Long-Short portfolio constructed using technical information, which implies a slight gap between the statistics-based predictiveness and implementability. Third, the question regarding the stability of technical information in predicting stock returns at different prediction horizons and under different market states needs to be addressed. With respect to the above research landscape, we build a unified comparison framework including Set1 to Set4, the OHLC Image CNN and the Raw OHLCV CNN, starting from the simplest price-volume information up to price path image information, and make analysis systematically based on three aspects: predictive power, economic value of the portfolio, and cross-period stability, aiming to explore the process of generating investment gains using technical information for stock return prediction..

3. Methodology

3.1. Research questions and empirical design

This paper examines the 1-day-ahead and 5-day-ahead directional movement of the stocks in CSI 300 from April 2, 2019 to December 30, 2025. Specifically, the main hypothesis of the study is whether traditional technical indicators, market-state variables, and nonlinear higher-order information of micro-structure have predictive ability that is incremental to the basic price-volume information, and whether the information could be utilized in generating economically significant investment portfolios. Thus, this study focuses on the out-of-sample portfolio performance with transaction cost consideration, not only on the statistical accuracy.

According to the different information dimensions and information construction logic, this paper categorizes the technical information into four progressive levels, including basic price-volume information, traditional technical indicator information, market state macro information, and nonlinear higher-order technical information with three components, which are StockJump, VPIN, and price path based on images. This multi-level information system provides a comprehensive representation of micro-trading information and market operational information of each stock.

To precisely isolate the marginal incremental contributions of various types of technical information, this paper uniformly adopts an expanding-window rolling out-of-sample prediction framework for empirical research. This paper abandons the single evaluation paradigm of traditional studies that uses fitting metrics such as AUC and statistical accuracy to determine model performance, and no longer relies on minor fluctuations in statistical metrics to define model superiority. Instead, based on the out-of-sample upward probabilities output by each model, stocks within the sample pool are cross-sectionally sorted by predicted upward probability, and standardized long-only and long-short hedging portfolios are constructed. Combined with risk-return performance under zero-cost and transaction-cost scenarios, the true incremental value and economic viability of various types of technical information are comprehensively evaluated.

3.2. Data sources and sample construction

The data used in this study is daily individual stock data of the CSI 300 constituent stocks from 2019 to 2025. Core data fields include individual stock opening price, highest price, lowest price, closing price, trading volume, daily return rate, turnover rate, and other basic price-volume

DOI: 10.54254/2754-1169/2026.LD35881

indicators, ultimately constructing a stock-day unbalanced panel sample that provides standardized data support for subsequent model training and empirical testing.

This paper defines the directional prediction label based on the future interval compound return of individual stocks. Let the daily return of stock i on trading day t be r_{i,t}, and the compound return over the next k trading days is defined as: $itr_{i,t}h$

$$R_{i,t}^{(h)} = \prod_{k=1}^{h}\left(1 + r_{i,t+k}\right) - 1, \quad h \in \left\{1,5\right\} \tag{1}$$

On this basis, this paper constructs a binary classification label for the future directional movement:

$$y_{i,t}^{(h)} = 1\left(R_{i,t}^{(h)} > 0\right), \quad h \in \left\{1,5\right\} \tag{2}$$

where the label equals 1 if the stock achieves a positive compound return over the next k trading days, and 0 if the future interval compound return is non-positive.

Beyond the basic price-volume variables, this paper incorporates two types of higher-order micro-trading variables, StockJump and VPIN, to characterize price jumps, abnormal volatility, order flow imbalance, and latent informed trading risk, and to examine whether they provide marginal predictive information. StockJump is a stock price jump characteristic indicator, and VPIN (Volume-Synchronized Probability of Informed Trading) reflects the short-term capital competition intensity and trading sentiment of individual stocks. Both types of indicators adopt industry-standardized calculation methods as microstructure information supplements beyond traditional price-volume indicators.

This paper employs a stock-day time-series aligned merging approach to accomplish multi-source data integration. Using stock code (Code) and trading day date (Date) as the unique linking keys, the basic daily price-volume market data, derived technical indicator data, market state variable data, and StockJump and VPIN higher-order micro-trading data are precisely matched, with unified time-series and cross-sectional dimensions. Invalid observations with missing fields, abnormal codes, or misaligned dates are removed, ultimately constructing an unbalanced panel dataset with unified dimensions, continuous time series, and complete indicators, ensuring data consistency for subsequent feature construction and model training. This research remove observations with insufficient historical windows, missing features, abnormal trading gaps, or future-return windows crossing validation or test boundaries to avoid look-ahead bias.

These studies motivate the inclusion of market-state information in return prediction. In this paper, market-state variables are introduced as an additional feature group in Set3 to test whether market-level conditions provide incremental predictive information beyond firm-level price-volume and technical variables.

3.3. Feature set system design

To progressively isolate the marginal incremental contributions of different types of technical information and precisely quantify the differential predictive value of basic price-volume indicators, traditional technical indicators, market state variables, and higher-order micro-trading features, this

DOI: 10.54254/2754-1169/2026.LD35881

paper constructs four progressive feature sets. Each set is incrementally layered from basic information to higher-order nonlinear information, enabling hierarchical examination of the predictive ability of different dimensions of technical information. The specific composition, variable attributes, and research functions of each feature set are shown in Table 1.

Table 1. Feature set composition and research functions

Feature Set	Content	Function	Number of Features
Set1	Basic price-volume variables	Control for returns, volatility, volume, turnover, liquidity and other basic information	13
Set2	Set1 + Traditional technical indicators	Test whether traditional technical indicators provide increment beyond basic price-volume information	30
Set3	Set2 + Market state variables	Test whether overall market returns, volatility, and trading states provide additional information	38
Set4	Set3 + StockJump + VPIN	Test whether jumps and order flow imbalance provide additional information	77

It is important to note that Set1_Base in this paper is not untreated raw market data, but rather a baseline control variable group (baseline controls) that has been screened, cleaned, and standardized, systematically encompassing core basic trading dimensions such as individual stock returns, volatility, trading liquidity, and market capitalization. This establishes a comprehensive foundational information benchmark that effectively avoids basic information omission bias, providing a rigorous and uniform comparison standard for subsequent incremental value testing of each feature set.

From the perspective of indicator construction logic and information nesting relationships, the weak predictive incremental effect of Set2 relative to Set1 possesses full theoretical rationality and logical necessity. The various traditional technical indicators incorporated in this paper are essentially secondary processing and derivative calculations of basic price, return, volume, and volatility information, without independent original information sources. Specifically, the RSI indicator is calculated based on the difference between historical gains and losses of individual stocks, the MACD indicator is constructed from the difference between moving averages of stock prices over different periods, Bollinger Band width directly maps the historical volatility level of individual stocks, and the ATR indicator characterizes the intraday price amplitude features of individual stocks. Since the Set1 baseline control group already fully covers all underlying core information including returns, volatility, volume, and intraday trading, the information content of traditional technical indicators is essentially fully explained and encompassed by Set1 variables. As a result, with respect to the full set of baseline constraints Set1 imposes, there is no way for Set2's conventional technical indicators to generate any useful marginal gain in terms of performance.

3.4. Model specification and rolling training mechanism

In this paper, the SGD Logistic Model is used as the unified prediction model where the format of inputs, prediction outputs, and the training process of model are standardized to guarantee the horizontal comparison of empirical results under various feature sets and time ranges. The specific parameters for the core model are set as follows: the input features are set as 20 days' data of an individual stock in its past, while the output predictions are defined as 1-day ahead and 5-day ahead

DOI: 10.54254/2754-1169/2026.LD35881

short-term direction classification. The testing period will cover from the first half year of 2022 to the second half year of 2025 with semi-annual rolling.

To prevent overfitting and data snooping, a dynamic expanding window method is used for dividing datasets in this paper. In each semi-annual testing window, the half-year just before the test window will be set as the validation window for choosing model hyperparameters and optimal model structure; all the previous samples will be regarded as the expanding training set.

In data preprocessing stage, this paper uses intraday cross-sectional quantile ranking to conduct feature standardization, avoiding problems of variable scale difference and extreme values. The scaler used for standardization will be fitted solely using the training set sample, thus preventing any possibility of information leakage of test set. At the same time, the rules for sample selection will be very strict, where in the case that the calculation period of sample return for prediction target horizon overlaps with the validation/test period, the sample will be directly discarded.

3.5. Image-based price path extension model

In order to overcome the information limitations of traditional structured features and dig out the potential of predicting with nonlinear technical information such as stock price path and volatility path, this paper converts OHLC price information of the past 20 trading days into images or tensors and uses CNN to learn price path information, thereby realizing intelligent mining and modeling of price path information. In this approach, we convert the time series trading information of a single stock into standardized images, and automatically discover price path features without defining technical pattern rules in a data-driven way, which can not only discover the structure of price volatility, amplitude feature of intraday price, but also the matching relationship between volume and price, which is not covered by traditional indicators.

The image model uses each stock's OHLC information over the past 20 trading days as its primary input to represent the short-term price path. The Raw OHLCV CNN is included as a control model that directly uses the raw price-volume sequence as input, allowing the analysis to distinguish the contribution of CNN architecture from the contribution of image-based price-path representation. The prediction tasks cover both 1-day-ahead and 5-day-ahead binary directional classification, and all out-of-sample predicted upward probabilities enter the unified portfolio backtesting framework.

The model training procedure is done in accordance with the unified semi-annual rolling expanding window out-of-sample empirical framework that is applied in the whole paper, using time-series partitioning criteria, sample selection, anti-data-leakage restrictions, and base model fully consistent, guaranteeing horizontal comparability of empirical results. Model training procedure depends on the available hardware and tries to use GPU acceleration, otherwise fallback to CPU training is done with 5 training epochs maximum per each rolling period.

3.6. Evaluation metric system

In order to tackle the discrepancy between statistical prediction and economic significance, this paper uses a dual assessment approach where portfolio performance is viewed as the main factor while classification measures are used as secondary indicators. Stocks in the sample data set are ranked on the basis of their out-of-sample predicted upward probability for each model and equally partitioned into five quintiles. The quintile containing the highest probability is named Q5, while the quintile with the lowest probability is defined as Q1.

DOI: 10.54254/2754-1169/2026.LD35881

Based on this sorting, two equal-weight portfolios are constructed. The first is a long-only portfolio holding Q5 stocks, which measures the positive stock-selection ability of each signal. The second is a long-short portfolio that goes long Q5 and short Q1, which is used to extract the excess return associated with model-based cross-sectional ranking while reducing exposure to broad market movements.

The empirical analysis reports AUC, Accuracy, and F1 as auxiliary prediction diagnostics, but these metrics cannot capture transaction costs, turnover, or practical portfolio value. Therefore, the main evaluation metrics are long-only and long-short annualized returns, Sharpe ratios, transaction-cost-adjusted net returns, and portfolio turnover. The long-short portfolio is interpreted as a standard academic portfolio construction rather than a directly implementable A-share trading strategy.

4. Empirical results and analysis

4.1. Testing the predictive ability and cross-sectional sorting ability of technical information

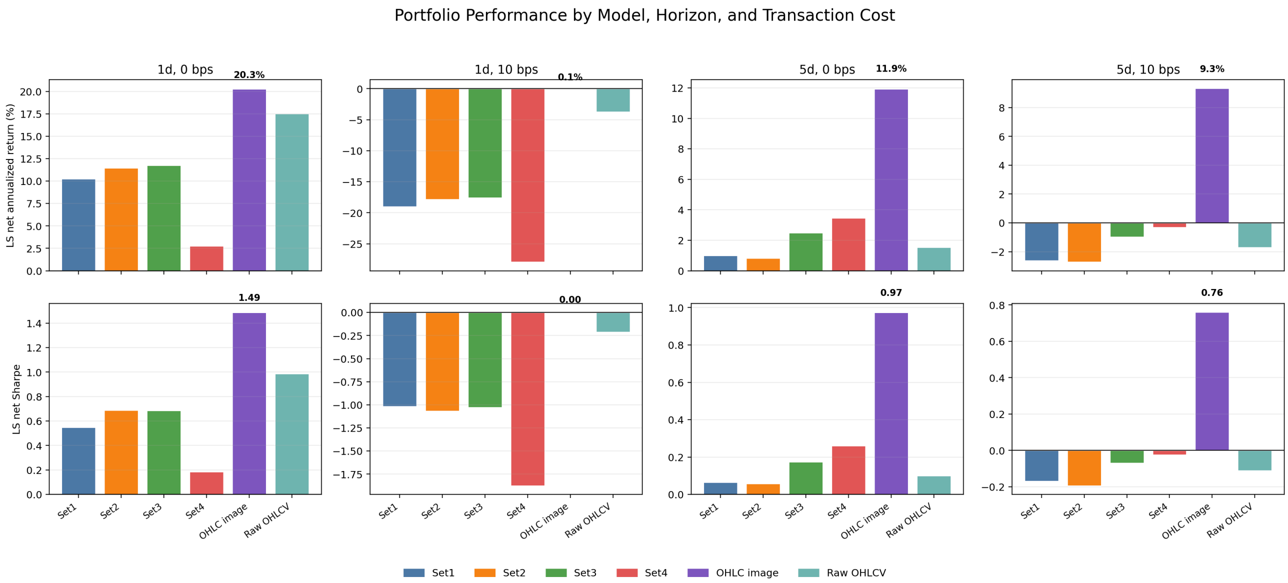

Figure 1. Portfolio performance by model, horizon, and transaction cost

This paper first examines the predictive ability and cross-sectional sorting ability of different levels of technical information under conditions without transaction costs. Since no trading frictions are introduced at this stage, portfolio returns can be regarded as a direct reflection of the model's ability to identify cross-sectional differences in future returns, and thus can reflect the original predictive value embedded in different information sets.All models construct quintile portfolios based on out-of-sample predicted probabilities, and Long-Short portfolios are formed by going long on the highest probability group (Q5) and shorting the lowest probability group (Q1). The main results are visualized in Figure 2.

From the 1-day-ahead prediction results, the four structured feature models Set1 through Set4 all exhibit a certain degree of cross-sectional sorting ability, but returns do not show a stable increasing trend as the information level expands. As the baseline model, Set1 achieves a 10.26% Long-Short annualized return and a 0.550 Sharpe ratio using only basic price-volume information, indicating that basic trading information itself already contains some short-term predictable components. After adding traditional technical indicators, Set2's portfolio return rises to 11.48% with a Sharpe of

DOI: 10.54254/2754-1169/2026.LD35881

0.689; further adding market state variables to form Set3, the Long-Short annualized return is 11.75% with a Sharpe of 0.687, still a limited improvement. Notably, when StockJump and VPIN higher-order micro-trading variables are further added to form Set4, the portfolio return actually declines to 2.79% with a Sharpe of only 0.187.These results indicate that traditional technical indicators, market state variables, and higher-order micro-trading features, while expanding the feature dimension, do not significantly enhance the model's ability to identify cross-sectional differences in future returns. In particular, the improvement of Set2 and Set3 over Set1 is very limited, suggesting that a large number of traditional technical indicators are essentially functional transformations of basic trading information such as prices, returns, volumes, and volatility, and their incremental information content is relatively limited; although market state variables introduce market-level environmental information, their contribution to the relative return ranking of individual stocks is similarly weak. As for the micro-trading structural information characterized by StockJump and VPIN, it may contain stronger noise components, and under the current prediction framework, it has not been effectively transformed into a stable cross-sectional sorting advantage.

However, when the research moves on to the extension from structured feature representation to data-driven feature learning, the findings begin to vary significantly. The OHLC Image CNN model yields 20.29% Long-Short annualized return and a Sharpe ratio of 1.486 for the 1-day-ahead task, outperforming all structured feature models in terms of profitability. Compared to Set1 to Set4, the image model is different in that it does not depend on a pre-built technical indicator system to learn feature representations but directly extracts the pattern features from the price path. This finding shows that there is non-linear information contained in the price trajectory which traditional structured variables have trouble capturing, and that the OHLC image CNN model is capable of doing so efficiently.

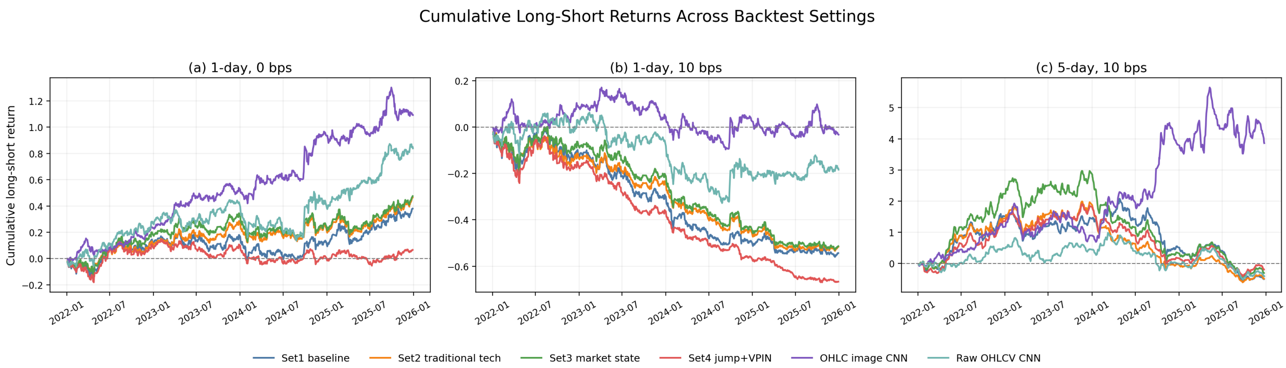

Figure 2. Cumulative long-short returns across backtest settings

4.2. Testing economic value under transaction cost constraints

Forecasting capacity does not imply investment capacity. All forecasts need to go through the actual trading process in order to become profits, and trading cannot avoid the constraints of turnovers and trading costs. After confirming that technical information has some capacity of sorting, this article will apply the trading cost of 10 bps to test their economic feasibility.

The 1-day-ahead task demonstrates the gap between predictive ability and economic value. After adding 10 bps transaction costs, the long-short portfolio returns of all structured models become negative. Specifically, Set1's long-short annualized return drops from 10.26% to -19.01%, Set2 to -17.85%, Set3 to -17.61%, and Set4 to -27.96%. Meanwhile, the long-short turnover rates of all

DOI: 10.54254/2754-1169/2026.LD35881

models remain between 1.16 and 1.22, indicating frequent rebalancing. Transaction costs therefore almost completely offset the gross returns generated before costs.

Even the best-performing OHLC image CNN cannot avoid this problem. Its Long-Short annualized return drops from 20.29% to 0.05%, with a net Sharpe of only 0.004. This result indicates that although the image model can discover stronger short-term prediction signals, the corresponding portfolio still faces extremely high turnover due to the excessively short prediction horizon, and therefore the vast majority of excess returns are ultimately offset by transaction costs.This finding does not mean that the image model is ineffective; rather, it reveals an important fact in financial prediction research: the existence of a signal does not imply its economic value. The 1-day-ahead task results have already demonstrated that price path images can improve cross-sectional sorting ability, but this ability is insufficient to overcome the cost barriers imposed by high-frequency trading.

Based on the above issues, this paper further examines the 5-day-ahead prediction task. A longer prediction horizon implies lower portfolio turnover. The results show that although structured models still maintain some returns under zero-cost conditions, most returns are completely eroded after cost adjustment. The Long-Short annualized returns of Set1 through Set4 under 10 bps conditions are -2.63%, -2.72%, -0.99%, and -0.32%, respectively, none forming stable excess returns.In contrast, the OHLC image CNN exhibits a clear advantage in the 5-day-ahead task. Under zero-cost conditions, its long-short annualized return reaches 11.94% with a Sharpe ratio of 0.974. After adding 10 bps transaction costs, it still retains a positive long-short net annualized return of 9.34% and a net Sharpe of 0.761. Its long-short turnover rate is only 0.517, substantially lower than that of the 1-day-ahead task. This result suggests that the price-path information extracted by the image model remains positive after transaction costs and shows better economic viability over the longer prediction horizon.

Notably, although the Raw OHLCV CNN also uses a CNN architecture, its post-cost long-short return is -1.71%, substantially weaker than that of the OHLC image CNN. This implies that the CNN architecture itself is not the sole source of performance. The more important factor is the image-based representation of price paths, which preserves spatial and morphological information that is difficult to capture from raw numerical sequences.

4.3. Relationship between classification metrics and portfolio returns

To further illustrate the relationship between model predictive ability and investment returns, this paper simultaneously reports traditional classification metrics including AUC, Accuracy, and F1, with results shown in Table 3.

Table 2. Classification prediction metrics

Model	Target	Sample Size	AUC	Accuracy	F1
Set1 Basic Price-Volume	1d	279,413	0.510	0.509	0.481
Set2 Traditional Tech. Indicators	1d	279,413	0.510	0.508	0.481
Set3 Market States	1d	279,413	0.509	0.506	0.480
Set4 Jump+VPIN	1d	279,413	0.510	0.506	0.484
OHLC Image CNN	1d	282,851	0.510	0.508	0.478

DOI: 10.54254/2754-1169/2026.LD35881

Table 2. (continued)

Raw OHLCV CNN	1d	283,145	0.516	0.505	0.527
Set1 Basic Price-Volume	5d	269,889	0.506	0.504	0.493
Set2 Traditional Tech. Indicators	5d	269,889	0.504	0.503	0.493
Set3 Market States	5d	269,889	0.503	0.498	0.507
Set4 Jump+VPIN	5d	269,889	0.504	0.499	0.508
OHLC Image CNN	5d	273,353	0.514	0.503	0.556
Raw OHLCV CNN	5d	273,647	0.511	0.507	0.492

In general, the AUC values of all models are centered around 0.50, with the best AUC value being just 0.516 of the Raw OHLCV CNN in the 1-day-ahead task and 0.514 of the OHLC image CNN in the 5-day-ahead task. In terms of traditional machine learning, such classification performances are quite normal. But on the other hand, the OHLC image CNN is also able to deliver a better portfolio return performance compared with other models. Such a result suggests that in stock cross-sectional prediction tasks, there is no straightforward relationship between classification performance and investment value. The reason is that constructing a portfolio does not depend on the classification result but the cross-sectional ranking of predicted probabilities. If the model can rank the future outperformers and future underperformers correctly, then even though it does not bring about a great improvement in classification accuracy, it can still produce a long-short return.

Based on the analysis from this paper, there is no significant gap between the performances of structured models and image models in terms of AUC, while the gap at the level of portfolio return is huge. This once again shows that the traditional classification performance measures are not sufficient to evaluate the value of the predictive signal for investment purposes. In the field of asset pricing and quantitative investment, the measures that are more important to consider include portfolio return, Sharpe ratio, and net return minus trading cost.

4.4. Semi-annual stability analysis

After confirming that the OHLC image CNN possesses economic value, this paper further examines its stability across different market phases. To avoid the overall sample results being influenced by abnormal performance in specific periods, this paper computes the return performance of each model across different testing intervals using semi-annual rolling windows, and measures model robustness through window-level return distributions.The results suggest that the 5-day-ahead OHLC image CNN is not entirely driven by a single testing window and shows some cross-period robustness. Six out of eight semi-annual windows generate positive returns, with an average net annualized return of 9.43% and an average net Sharpe of 0.777. In contrast, the Raw OHLCV CNN achieves positive returns in only four windows, with an average net annualized return of -1.63%. The structured models also show weaker average post-cost performance.

Table 3. Semi-annual stability statistical results

Model	Target	Valid Windows	Positive-Return Windows	Avg. Net Ann.	Min	Max	Avg. Net Sharpe
OHLC Image CNN	1d	8/8	5/8	-0.15%	-19.95%	16.43%	-0.049

DOI: 10.54254/2754-1169/2026.LD35881

Table 3. (continued)

OHLC Image CNN	5d	8/8	6/8	9.43%	-7.00%	22.39%	0.777
Raw OHLCV CNN	1d	8/8	6/8	-4.09%	-49.94%	15.13%	-0.216
Raw OHLCV CNN	5d	8/8	4/8	-1.63%	-25.78%	17.03%	0.069
Set2 Traditional Tech. Indicators	1d	8/8	1/8	-18.11%	-49.55%	7.08%	-1.099
Set2 Traditional Tech. Indicators	5d	8/8	5/8	-2.66%	-27.30%	30.24%	-0.184
Set3 Market States	1d	8/8	1/8	-17.87%	-50.22%	2.47%	-1.048
Set3 Market States	5d	8/8	5/8	-0.94%	-27.77%	36.42%	-0.067
Set4 Jump+VPIN	1d	8/8	0/8	-28.20%	-55.13%	-11.23%	-2.062
Set4 Jump+VPIN	5d	8/8	5/8	-0.28%	-19.48%	24.21%	-0.007

4.5. Summary

Synthesizing the above analysis, this section forms a chain of evidence. First, under zero transaction costs, traditional technical indicators, market-state variables, and microstructure variables provide limited incremental predictive ability relative to basic price-volume information, while the OHLC image CNN substantially improves cross-sectional sorting ability. Second, after introducing transaction costs, the 1-day-ahead signals are largely eroded by high turnover, showing that predictive ability does not necessarily translate into investment value. Third, when the prediction horizon extends to 5 days, the OHLC image CNN retains positive post-cost performance and shows some cross-period robustness.

From the empirical results, the incremental value of traditional technical indicators, market state variables, and higher-order micro-trading variables is overall limited, while the nonlinear feature extraction method based on price path images can more effectively identify cross-sectional differences in future returns and ultimately form excess returns of practical investment significance. The stable performance achieved by the 5-day-ahead OHLC image CNN constitutes the core empirical evidence for the incremental value of technical information.

5. Conclusion

This paper examines the incremental value of different levels of technical information in stock return direction prediction from the perspective of whether statistical predictive ability can be transformed into portfolio returns. The results show that basic price-volume variables already contain substantial short-term predictive information. Adding traditional technical indicators, market-state variables, and microstructure variables such as StockJump and VPIN does not bring consistent portfolio return improvements. This suggests that many structured technical indicators are reprocessed versions of existing price, return, volume, and volatility information, and their marginal value beyond comprehensive baseline controls is limited in the current daily-frequency cross-sectional framework. The main result of this paper is generated by the 5-day ahead OHLC image CNN model. This model keeps a net annualized return of 9.34% with a net Sharpe ratio of 0.761 even with 10 bps transaction cost, and significantly outperforms structured feature models as well as the raw OHLCV CNN. The above result indicates that image-based price path representation is able

DOI: 10.54254/2754-1169/2026.LD35881

to maintain morphological information which is hard for conventional artificial signals to extract, and the longer prediction horizon helps alleviate turnover pressure.

Potential future work might extend the sample size beyond CSI 300 and further investigate the heterogeneity of regimes and the interpretability of the models.

References

[1] Shi, S., Zhang, B., & Peng, X. (2025). Technical indicators and aggregate stock returns: An updated look. *Journal of Multinational Financial Management*, Article 100898. https://doi.org/10.1016/j.mulfin.2025.100898

[2] Mostafavi, S. M., Farooque, M. U., Reza, M. S., et al. (2025). Key technical indicators for stock market prediction. *Machine Learning with Applications*, Article 100631. https://doi.org/10.1016/j.mlwa.2025.100631

[3] Wang, H.-C., Hsiao, W.-C., & Liou, R.-S. (2024). Integrating technical indicators, chip factors and stock news for enhanced stock price predictions: A multi-kernel approach. *Asia Pacific Management Review, 29*(3), 292–305. https://doi.org/10.1016/j.apmrv.2023.10.001

[4] Neuhierl, A., Weber, M., & Zhou, G. (2024). Empirical asset pricing with probability forecasts. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4717935

[5] Cakici, N., Fieberg, C., Metko, D., & Zaremba, A. (2023). Machine learning goes global: Cross-sectional return predictability in international stock markets. *Journal of Economic Dynamics and Control, 155*, Article 104725. https://doi.org/10.1016/j.jedc.2023.104725

[6] Horvath, B., Jacquier, A., & Tankov, P. (2024). Clustering market regimes using the Wasserstein distance. *The Journal of Computational Finance*. https://doi.org/10.21314/jcf.2024.005

[7] Kelly, B., Malamud, S., & Zhou, K. (2024). The virtue of complexity in return prediction. *The Journal of Finance*. https://doi.org/10.1111/jofi.13298

[8] Hanauer, M. X., & Kalsbach, T. (2023). Machine learning and the cross-section of emerging market stock returns. *Emerging Markets Review, 55*, Article 101022. https://doi.org/10.1016/j.ememar.2023.101022

[9] Chen, M., Hanauer, M. X., & Kalsbach, T. (2024). Design choices, machine learning, and the cross-section of stock returns. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.5031755

[10] Albahli, S., Alhassan, F., Khan, M. A., et al. (2023). An improved DenseNet model for prediction of stock market using stock technical indicators. *Expert Systems with Applications, 228*, Article 120903. https://doi.org/10.1016/j.eswa.2023.120903

[11] Saberironaghi, M., Ren, J., & Saberironaghi, A. (2025). Stock market prediction using machine learning and deep learning techniques: A review. *AppliedMath, 5*(3), Article 76. https://doi.org/10.3390/appliedmath5030076

[12] Giantsidi, S., & Tarantola, C. (2025). Deep learning for financial forecasting: A review of recent trends. *International Review of Economics & Finance, 104*, Article 104719. https://doi.org/10.1016/j.iref.2025.104719

[13] Wang, Z. (2025). Machine learning for stock return prediction: Transformers or simple neural networks. *Finance Research Letters, 86*, Article 108783. https://doi.org/10.1016/j.frl.2025.108783

[14] Kijkarncharoensin, A., Kritsadakorn, S., & Chaiyarak, K. (2025). Identifying risk regimes in a sectoral stock index through a multivariate hidden Markov framework. *Risks, 13*(7), Article 135. https://doi.org/10.3390/risks13070135

[15] Gupta, R., Kapoor, S., Gupta, H., & Natesan, S. (2025). A forest of opinions: A multi-model ensemble-HMM voting framework for market regime shift detection and trading. *Data Science in Finance and Economics, 5*(4), 466–501. https://doi.org/10.3934/dsfe.2025019