

Interpretable Intraday Signals for Real-Time Open-to-Close Directional Prediction: Evidence from Chinese Equity Indices

Yingxuan Huang

School of Arts and Sciences, Rutgers University-New Brunswick, New Brunswick, NJ, 08901, USA
huangyingxuan000@gmail.com

ABSTRACT

This study tests whether a practitioner-style intraday heuristic can be formalized into reproducible, statistically testable signals for real-time stock-index directional prediction. At 13:30 each trading day, the primary task predicts whether the final close will exceed the daily open. Although the final close remains unknown at the prediction time, part of the open-to-close path has already been observed, so the task differs from forecasting only the remaining-session direction. The heuristic links two 30-minute windows, Zone 2 (10:01–10:30) and Zone 5 (13:01–13:30), to the final open-to-close direction of Chinese equity indices and treats a Zone 5 breakout of the morning range as a confirmation signal. We encode this judgment as three discrete variables, ZONE2_SIGNAL, ZONE5_SIGNAL, and BREAK_SIGNAL, and estimate an interpretable logistic regression using only information available by 13:30. The model is trained on 2015–2024 observations for the Shanghai Stock Exchange (SSE) Composite Index and the China Securities Index (CSI) 300 Index, evaluated on their 2025 observations, and applied without re-estimation to the Shenzhen Stock Exchange (SZSE) Component Index. The formalized framework improves 2025 holdout accuracy from 65.43% for the deterministic rule to 69.55%, with an area under the receiver operating characteristic (ROC) curve (AUC) of 0.768. The same-cutoff 13:30 price-position benchmark remains stronger, with 83.54% accuracy and an AUC of 0.893. As a supplementary remaining-session boundary test, the same specification is re-estimated to predict whether the final close exceeds the 13:30 close. Its 2025 holdout accuracy is 54.12%, with an AUC of 0.481, and it does not outperform the majority-class benchmark. The evidence therefore supports a transparent signal-decomposition framework for real-time open-to-close directional prediction, but not reliable post-13:30 forecasting, trading profitability, causal mechanisms, or market-wide generality.

Keywords: Intraday Pattern Recognition, Breakout Signal, Real-Time Directional Prediction, Logistic Regression, Out-Of-Sample Validation

1. INTRODUCTION

Short-horizon market direction remains one of financial economics’ most persistent and contested prediction problems. Weak-form market efficiency implies that historical prices should contain little systematic information about future returns [1]. By contrast, behavioral finance, technical analysis, and market-microstructure research suggest that investor behavior, information arrival, trading frictions, and institutional features may create temporary patterns that remain statistically detectable even when they are difficult to exploit economically [2–11]. This tension is especially relevant intraday, where information flow, liquidity, and trading behavior may vary across the trading session.

Most empirical studies address this issue either by developing predictive models or by testing established technical-analysis rules [4–6,12,13]. Less attention has been given to practitioner-style intraday judgments that guide discretionary decisions but are rarely expressed as testable scientific objects. Such rules often remain anecdotal: definitions vary across individuals, implementation is difficult to replicate, and predictive value is seldom evaluated under controlled statistical conditions. As a result, even potentially useful trading intuition contributes little to cumulative knowledge unless it can be audited, replicated, and compared across samples. Existing studies show that technical rules and chart patterns can be tested statistically, but less is known about how a practitioner-style intraday heuristic can be decomposed into auditable component signals while preserving its original decision structure.

This paper addresses that methodological gap. Rather than asking whether a particular trading rule can be turned into a profitable strategy, the study asks whether an informal intraday judgment can be translated into a measurable statistical representation. Specifically, it examines whether the Zone 2–Zone 5 heuristic, encoded as fixed signals using information available by 13:30, contains directional information about the final open-to-close outcome. The primary task is therefore

a real-time directional prediction at 13:30 each trading day, formulated as a binary classification of whether the still-unknown final close will be above or below the daily open. A supplementary analysis replaces the outcome with the direction from 13:30 to the close, providing a remaining-session boundary test without changing the signal definitions or information cutoff.

The contribution is therefore narrower than proposing a new trading strategy, but more methodological and diagnostic: it evaluates whether practitioner reasoning can be expressed as fixed signals, tested under a real-time information cutoff, and interpreted through a simple model. This framing positions the analysis as an interpretable framework for studying intraday directional states, rather than as a direct trading prescription. Against this background, the study differs from conventional indicator-based or high-dimensional forecasting research in what it evaluates: a practitioner-style intraday judgment translated into predefined, auditable, and statistically testable variables. Logistic regression is used as a measurement instrument because its coefficients, odds ratios, and predicted probabilities remain directly interpretable in relation to the underlying signals ^[14].

The heuristic examined here links two specific 30-minute windows in the Chinese equity-index trading session to the final open-to-close direction of the trading day. Zone 2 (10:01–10:30) occurs after the initial burst of opening activity, while Zone 5 (13:01–13:30) follows the midday trading break. The rule also treats a Zone 5 break above the morning high or below the morning low as a confirmation signal. These judgments are encoded as three discrete variables: ZONE2_SIGNAL, ZONE5_SIGNAL, and BREAK_SIGNAL. The decision point is fixed at 13:30, so each predictor uses only information available by that time.

The empirical design deliberately prioritizes interpretability and disciplined validation. The objective is not to search for the strongest possible intraday predictor or to replace the rule with an opaque forecasting mechanism. Instead, the model estimates whether the three components contain measurable directional information and how their relative contributions differ. The design uses chronological holdout testing, cross-index transfer validation, rolling evaluation, sensitivity checks, and a same-cutoff 13:30 price-position benchmark to assess the stability and boundary of the signal evidence.

This study makes three contributions. First, methodologically, it converts a Zone 2–Zone 5 practitioner rule into a fixed signal system with an explicit 13:30 information cutoff. Second, empirically, it measures the directional information in the rule's components using logistic coefficients, odds ratios, probability ranking, and pattern diagnostics. Third, from a validation perspective, it evaluates the open-to-close analysis through chronological holdout testing, cross-index transfer validation, rolling out-of-sample checks, sensitivity analysis, and a same-cutoff benchmark, while using the post-13:30 target as a supplementary boundary test. The principal contribution lies in this process of signal construction and validation rather than in the classifier itself. This distinction matters because many forms of practitioner knowledge in financial markets are neither fully specified forecasting models nor directly testable variables. The paper therefore treats qualitative intraday judgment as an empirical object that can be reproduced, evaluated, and challenged.

The focus on fixed intraday windows has a market-microstructure motivation because information flow, trade informativeness, and trading intensity can vary across the session ^[7–11]. The 10:01–10:30 window follows the opening adjustment period, while the 13:01–13:30 window captures the first half-hour after the midday break. If post-lunch order revisions or renewed information processing matter, a Zone 5 breakout may act as a confirmation signal rather than as a generic momentum input. The specific windows are not treated as theory-derived intervals; they are practitioner-inspired intervals to be tested empirically.

2. RELATED WORK AND HYPOTHESES

The literature offers both support for testing technical rules and strong reasons for caution. Brock et al. ^[4] and Lo et al. ^[6] support the idea that price patterns can be made testable, but Sullivan et al. ^[5] show why data snooping can inflate apparent rule performance. Work on Chinese indices reports informative technical rules yet also treats profitability and stability as sample dependent ^[12]. More directly relevant here, intraday-pattern evidence suggests that predictability may depend on market structure, trading-session design, and sample period ^[9–12]. We therefore test a narrow fixed-window signal and use the rule as a classification signal rather than as a ready-made trading strategy or a broad claim about universal technical-analysis profitability. Directional classification fits the primary target because the dependent variable is the sign of the open-to-close move, not the index level ^[15]. The remaining-session target is reserved for the supplementary forward test. Logistic regression is used because it provides probabilistic classification while allowing each engineered signal to be interpreted through its coefficient and odds ratio ^[14].

Two implications from this literature constrain the paper’s claims. First, evidence on chart rules varies across markets, sample periods, and trading-cost assumptions ^[13]. Second, volume, order flow, and liquidity can alter the meaning of a price movement ^[16–18]. For that reason, the empirical claim here is limited to directional classification. It is not a claim about net trading returns or economic profitability.

Prior studies have made established chart patterns and technical trading rules empirically testable ^[4–6,12,13]. Less attention, however, has been given to the prior methodological step of translating a specific practitioner-defined intraday judgment into a reproducible construct with a fixed decision time. This study addresses that narrower gap by specifying the Zone 2–Zone 5 heuristic explicitly, separating its component signals, and evaluating the resulting representation through chronological holdout, cross-index transfer, rolling out-of-sample, and prominence-sensitivity checks. Its contribution is therefore primarily methodological: it provides an auditable framework for evaluating informal market judgment while preserving a strict real-time information boundary. Prior literature therefore provides three foundations for this study: the testability of price-based patterns, the role of intraday market structure, and the need for interpretable validation when converting qualitative judgment into quantitative variables.

The hypotheses are kept close to the measured variables. H1 tests whether Zone 2 direction is associated with the daily sign. H2 tests whether Zone 5 direction adds later-session information. H3 asks whether a breakout from the morning range strengthens the directional signal. H4 compares logistic weighting with the deterministic rule. H5 evaluates the result under the 2025 holdout, rolling out-of-sample tests, the SZSE cross-index sample, and alternative prominence settings. All five hypotheses concern the open-to-close analysis. The post-13:30 analysis is reported separately as a supplementary boundary test rather than as an additional hypothesis. These tests organize the empirical design; they do not amount to a general test of market efficiency.

3. DATA AND SIGNAL ENGINEERING

3.1 Data and validation design

One-minute open, high, low, close, volume, and turnover records are obtained from the MyQuant platform for the SSE Composite, CSI 300, and SZSE Component indices from 2015–2025. Each index-day contains 240 records, and both focal windows contain 30 records. The primary indices provide 1,283, 520minute observations. No required-field missing values or duplicate timestamps are observed.

Dates and minute labels are standardized before grouping. A trading day is retained only when its timestamps are unique, all required fields are present, the full 240-minute session is available, and the four focal boundaries (10:01, 10:30, 13:01, and 13:30) exist. This filter prevents an incomplete window from being treated as a valid pattern. Table 1 summarizes the resulting training, chronological holdout, and cross-index transfer partitions.

Table 1. Sample partition.

Partition	Indices	Period	N
Training	SSE; CSI 300	2015–2024	4,861*
Chronological holdout	SSE; CSI 300	2025	486
Cross-index	SZSE	2015–2025	2,674

*One exact open-close flat day has no binary target and is excluded, leaving 4,861 estimation observations. For the supplementary post-13:30 target, three training observations are flat between 13:30 and the close, leaving 4,859 training observations; no such flat observation occurs in the 2025 holdout. One flat SZSE day is excluded from the forward SZSE sample, leaving 2,673 observations. The SZSE sample is never used for fitting, threshold selection, or prominence selection.

3.2 Variables and timing control

Zone 2 is 10:01–10:30 and Zone 5 is 13:01–13:30. The information cutoff and decision time are fixed at 13:30. ZONE2_SIGNAL and ZONE5_SIGNAL equal 1, -1, or 0 for up, down, or sideways patterns. BREAK_SIGNAL equals 1 for a Zone 5 high strictly above the 09:31–11:30 morning high, -1 for a Zone 5 low strictly below the morning low, and 0 otherwise. A diagnostic count reports nine index-days satisfying both breakout conditions; under the prespecified new-high priority rule, these observations are assigned BREAK_SIGNAL = 1. The primary outcome, TARGET, equals 1 when the final close is above the daily open and 0 when it is below. It is the primary target for the real-time directional prediction

made at 13:30 each trading day. The supplementary outcome, FORWARD_TARGET, equals 1 when the final close is above the 13:30 close and 0 when it is below; exact 13:30-to-close flat observations are excluded. It isolates the direction of the remaining session after 13:30. The final close is used only to construct these outcomes and never enters a predictor.

To construct same-cutoff benchmarks, we define two price-position variables. MIDDAY_POSITION is a discrete directional indicator equal to 1, -1, or 0 according to whether the 13:30 close is above, below, or equal to the daily open; it is used to calculate the open-to-close benchmark's directional accuracy. OPEN_TO_1330_RETURN is the continuous return from the daily open to the 13:30 close and is used as its ranking score for AUC. Both variables use only information available by 13:30 and provide conservative same-time benchmarks for the primary analysis. In the supplementary forward test, their sign and continuous value are also used as a mechanically defined pre-13:30 momentum-continuation benchmark.

3.3 Reproducible pattern recognition

The pattern-recognition procedure translates an existing practitioner heuristic into a deterministic, reproducible encoding scheme; it does not search for an optimal technical-pattern set. Local peaks are detected from each 30-minute close sequence with `scipy.signal.find_peaks` in SciPy 1.15.3, and troughs are detected from its negative ^[19]. The baseline prominence is set to 0.20 times the window price range; distance, width, height, and other peak-selection parameters are left unset. Start, extrema, and end points are chronologically merged. Consecutive movements map to monotonic up/down, V, inverted V, N, inverted N, W, and M patterns. Longer sequences are assigned complex up, complex down, or sideways according to the end-to-start change. This mapping uniquely determines each direction signal.

The prominence threshold is recomputed within each window, so it scales with the window's own price range instead of an index-level absolute price increment. Endpoints are retained because an interval can be directional even without an interior extremum. When multiple extrema are present, their temporal order defines the pattern sequence.

4. EMPIRICAL MODELS

The rule benchmark follows the common-direction signal when Zones 2 and 5 agree. Under disagreement, a new high confirms an upward classification for a down/up pair, and a new low confirms a downward classification for an up/down pair; otherwise, the Zone 2 direction is retained. Sideways cases are unclassified. Thus, the breakout condition is used only to resolve directional disagreement between Zones 2 and 5; when both zones point in the same direction, the rule benchmark follows that common direction without an additional breakout requirement.

The logistic model estimates the probability that the final close exceeds the daily open:

$$\ln\left(\frac{P_i}{1-P_i}\right) = \beta_0 + \beta_1 B_i + \beta_2 Z_{2i} + \beta_3 Z_{5i} \quad (1)$$

In Equation (1), P_i denotes the probability that the final close exceeds the daily open for index-day observation i ; B_i , Z_{2i} , and Z_{5i} denote BREAK_SIGNAL, ZONE2_SIGNAL, and ZONE5_SIGNAL, respectively; and β_0 and β_1 – β_3 denote the intercept and corresponding signal coefficients.

The implementation uses scikit-learn 1.6.1 LogisticRegression ^[20]. The model retains the specified package settings: L_2 regularization, $C = 1.0$, `fit_intercept = True`, `class_weight = None`, and the `lbfgs` solver. Features are not standardized, the decision threshold is 0.5, and no test or SZSE observation is used to re-estimate coefficients.

This interpretable specification is chosen deliberately because the research question concerns whether the formalized components contain measurable directional information, rather than whether a more complex algorithm can maximize predictive accuracy. Model simplicity is therefore a feature of the research design, not a claim of performance superiority. More flexible nonlinear models may improve predictive fit, but they would make it more difficult to attribute directional information to specific practitioner-defined components.

Accuracy is calculated at the 0.5 probability threshold. AUC is also reported to assess ranking ability without relying on that single threshold. For the 2025 holdout metrics, Wilson intervals are used for accuracy and bootstrap intervals are used for AUC.

For coefficient inference, the same training matrix is re-estimated with an unregularized statsmodels 0.14.4 Logit model. Standard errors are clustered by trading date across 2,431 date clusters, allowing the two primary-index observations on the same date to share an arbitrary within-date error correlation. The model uses the identical three predictors and target definition, so the coefficient signs, cluster-robust significance tests, and odds ratios provide a statistical interpretation of

the same signal set; they should not be read as identical to the regularized scikit-learn coefficients. Multicollinearity is checked using variance inflation factors (VIF).

A one-variable baseline logit containing only BREAK_SIGNAL is also estimated. Comparing this baseline with the full model allows the incremental contribution of ZONE2_SIGNAL and ZONE5_SIGNAL to be evaluated using McFadden pseudo- R^2 , the Akaike information criterion (AIC), and the Bayesian information criterion (BIC). In addition, a 13:30 price-position benchmark is evaluated. Its directional version predicts an upward open-to-close day when the 13:30 close is above the daily open and a downward day when it is below the daily open. Its continuous version uses OPEN_TO_1330_RETURN as the ranking score for AUC calculation. Although intentionally simple, this benchmark is stringent because it uses information available at the same 13:30 cutoff. It is included as a conservative reference for determining whether the engineered signals add directional information beyond that already summarized by the index's contemporaneous position relative to the daily open.

For the supplementary forward test, the same three-signal specification is re-estimated with FORWARD_TARGET using the 2015–2024 primary-index observations and applied to the 2025 holdout without changing the predictors, information cutoff, regularization settings, or 0.5 threshold. It is compared with the deterministic rule, a training-sample majority-class baseline, and a pre-13:30 momentum-continuation benchmark. The fitted model is also evaluated in expanding-window tests and on the SZSE sample. This analysis is used to define the boundary of the primary evidence rather than to introduce an additional hypothesis.

5. RESULTS

5.1 Signal evidence and model estimates

As a descriptive full-sample diagnostic, the rule benchmark correctly classifies 3,446 of 5,341 valid observations (64.52%). This diagnostic summarizes the signal structure and is not used to fit the logistic model, choose thresholds, or select the SZSE index. The close is above the open on 87.57% of the 1,400 new-high conditions, while 81.59% of the 1,249 new-low conditions finish below the open. Because nine days meet both conditions, the two counts are not mutually exclusive.

Table 2 reports the unregularized statsmodels Logit estimates with standard errors clustered by trading date. BREAK_SIGNAL has the largest coefficient, 1.7181. ZONE2_SIGNAL and ZONE5_SIGNAL are smaller but still positive at 0.2055 and 0.1548. All three predictors remain statistically significant at the 1% level under the cluster-robust covariance estimator, providing coefficient-level support for H1–H3. Because BREAK_SIGNAL is coded -1, 0, and 1, its odds ratio of 5.574 applies to a one-unit change. A move from -1 to +1 is a two-unit change and corresponds to an odds ratio of about 31.07, holding the other signals constant.

Only five Zone 2 windows and three Zone 5 windows are classified as sideways. The deterministic rule, therefore, produces a direction for nearly all non-flat index days.

These estimates show that BREAK_SIGNAL is the dominant component of the three-signal specification. ZONE2_SIGNAL and ZONE5_SIGNAL have smaller coefficients but remain statistically significant with standard errors clustered by trading date, indicating that both interval-direction variables provide incremental directional information beyond the breakout component.

Table 2. Unregularized logistic coefficients with date-clustered inference and odds ratios.

Variable	Coef.	S.E.	z	p	OR
BREAK_SIGNAL	1.7181	0.0719	23.88	<0.001	5.574
ZONE2_SIGNAL	0.2055	0.0410	5.01	<0.001	1.228
ZONE5_SIGNAL	0.1548	0.0417	3.71	<0.001	1.167
Intercept	0.1797	0.0440	4.09	<0.001	1.197

Note: S.E. denotes standard error clustered by trading date (2,431 date clusters); OR denotes odds ratio. Coefficients are from the unregularized statsmodels Logit inference model.

The baseline comparison shows the incremental contribution of the two interval-direction variables beyond BREAK_SIGNAL. For the BREAK-only model, McFadden pseudo- R^2 is 0.1956, with AIC = 5397.25 and BIC = 5410.23.

With all variables included, the corresponding values improve to pseudo- $R^2 = 0.2044$, AIC = 5341.92, and BIC = 5367.88. VIFs for the three signals are 1.014 (BREAK_SIGNAL), 1.011 (ZONE2_SIGNAL), and 1.003 (ZONE5_SIGNAL), below typical warning thresholds.

5.2 Chronological out-of-sample test

The 2025 chronological holdout provides the paper’s primary out-of-sample test. Table 3 reports the corresponding results. Logistic weighting improves accuracy over the deterministic rule benchmark from 65.43% to 69.55%, an increase of 4.12 percentage points. Across the two primary indices, holdout accuracy is 67.49% for the SSE Composite and 71.60% for the CSI 300.

The same table also reports the conservative 13:30 price-position benchmark. This same-cutoff benchmark achieves 83.54% accuracy and an AUC of 0.893, compared with 69.55% accuracy and an AUC of 0.768 for the three-signal logistic model. The improvement over the deterministic rule indicates that probabilistic weighting provides a more informative representation of the three engineered signals, although the stronger same-cutoff benchmark indicates that the observed price position at 13:30 contains substantial directional information.

This comparison separates two empirical objectives: achieving stronger predictive performance and decomposing the formalized heuristic into interpretable components. The price-position benchmark performs better on the first objective, whereas the three-signal model addresses the second.

Table 3. Main predictive results.

Evaluation	Accuracy	AUC
2025 rule	65.43%	-
2025 13:30 price-position baseline	83.54%	0.893
2025 logistic	69.55%	0.768
SZSE 2015–2025 (mean annual logit performance)	71.51%	0.7939
SZSE 2025 (logit)	74.90%	0.8205

As an uncertainty check, we compute 95% confidence intervals for the main 2025 holdout metrics. Accuracy intervals are Wilson binomial intervals, and AUC intervals are obtained from 2,000 bootstrap resamples. The logistic AUC is 0.768 with a 95% confidence interval (CI) of [0.723, 0.808], while the 13:30 price-position baseline AUC is 0.893 with a 95% CI of [0.863, 0.920]. The confidence intervals are consistent with the point estimates: the three-signal logistic model has lower AUC than the same-cutoff price-position benchmark in the 2025 holdout. Figure 1 plots the receiver operating characteristic curve for the proposed logistic model on the 2025 holdout sample, showing its probability-ranking performance across classification thresholds.




Figure 1. ROC curve on the 2025 holdout sample.

Note: The curve evaluates probability ranking by the proposed logistic model on observations not used for model estimation; the 13:30 price-position baseline AUC is reported in Table 3.

5.3 Post-13:30 forward test

As a supplementary boundary analysis, the forward test replaces the open-to-close outcome with FORWARD_TARGET, which records whether the final close is above the 13:30 close. The three original signals and the 13:30 information cutoff are unchanged. After excluding exact 13:30-to-close flat observations, the model is estimated on 4,859 primary-index observations from 2015–2024 and evaluated on 486 observations in the 2025 holdout.

Table 4 reports the 2025 forward results. The three-signal logit achieves 54.12% accuracy and an AUC of 0.481, with a bootstrap 95% confidence interval of [0.431, 0.533]. It therefore does not outperform the 55.76% majority-class accuracy. The pre-13:30 momentum-continuation benchmark achieves 54.32% accuracy and an AUC of 0.532, while the deterministic rule achieves 46.91% accuracy. Using their class outputs as scores, the majority-class and deterministic-rule AUCs are 0.500 and 0.461, respectively. These results provide no evidence of reliable probability ranking for the post-13:30 direction in the 2025 holdout.

Table 4. Post-13:30 forward results in the 2025 holdout.

Model	Accuracy	AUC
Majority class	55.76%	0.500
Deterministic rule	46.91%	0.461
Pre-13:30 momentum	54.32%	0.532
Three-signal logit	54.12%	0.481

Note: The post-13:30 target equals 1 when the final close is above the 13:30 close and 0 when it is below. Exact flat observations are excluded. The majority-class and deterministic-rule AUCs use constant or binary class scores and therefore provide only coarse ranking references; they are not directly comparable with AUCs based on continuous probability scores.

Additional diagnostics point in the same direction. In the training sample, BREAK_SIGNAL remains positive, but ZONE2_SIGNAL and ZONE5_SIGNAL are not statistically significant, and the full forward model provides little improvement over the BREAK-only specification. Rolling AUC values remain close to chance, and the SZSE 2025 result does not outperform the corresponding momentum benchmark. Taken together, the forward tests do not support a stable post-13:30 forecasting claim.

5.4 Probability and pattern diagnostics

Probability-group analysis tests whether larger predicted probabilities correspond to larger realized upward rates. In the 2025 holdout sample, the empirical upward rate increases from 16.90% in the 0–20% group to 86.55% in the 80–100% group (Table 5 and Figure 2). The ordered pattern is consistent with probability ranking, but the group estimates are not equally precise. In particular, the 20–40% group contains only 17 observations and has a wide 95% Wilson interval. The pattern should therefore be interpreted together with group counts and confidence intervals rather than as strong evidence of exact calibration.

Table 5. Probability-group validation in the 2025 holdout with 95% Wilson intervals.

Predicted probability	N	Upward rate (95% CI)
0-20%	71	16.90% (9.94–27.26)
20-40%	17	17.65% (6.19–41.03)
40-60%	203	51.72% (44.88–58.50)
60-80%	76	71.05% (60.04–80.04)
80-100%	119	86.55% (79.27–91.55)




Figure 2. Upward rate by predicted-probability group.

Note: Higher probability groups have higher observed upward rates, but uncertainty is substantial in the smaller groups.

The pattern-combination heatmap checks whether the two interval signals collapse into the same information. They do not: some Zone 2-Zone 5 combinations have clearly different observed upward rates, while mixed combinations are less decisive (Figure 3). Rare pattern pairs can produce extreme cell values near 0 or 1, so those cells should be interpreted cautiously. The heatmap is a diagnostic summary rather than an additional feature, so it does not change the three-variable model.




Figure 3. Pattern-combination heatmap.

Note: Cells report observed upward rates for Zone 2–Zone 5 pattern pairs; sparse cells may produce extreme rates. Axis abbreviations: Mono. = monotonic; Inv. = inverted; C. = complex.

5.5 Cross-index transfer and sensitivity analysis

Cross-index transfer validation tests whether the fixed model estimated from the SSE Composite and CSI 300 can be applied to the SZSE Component Index without re-estimation. Across yearly evaluations, accuracy ranges from 65.98% to 75.62%, while AUC ranges from 0.7095 to 0.8321 (Figure 4). These results provide evidence of transfer across indices within the Chinese A-share market, but they should not be interpreted as independent-market external validation. Because the 2015–2024 SZSE observations overlap in calendar time with the primary estimation period, those years evaluate cross-index transfer rather than temporal out-of-sample performance. The SZSE 2025 result provides the clearest joint test of temporal and cross-index transfer.




Figure 4. Annual cross-index transfer accuracy and AUC of the fixed logistic model on the SZSE Component Index.

Note: The SZSE Component Index is not used in model fitting, parameter selection, or prominence selection. Results for 2015–2024 evaluate cross-index transfer under overlapping calendar periods; only the 2025 result is also outside the primary model’s training window.

Table 6. Prominence sensitivity.

Ratio	Accuracy	AUC
0.10	68.72%	0.777
0.15	68.93%	0.766
0.20	69.55%	0.768
0.25	70.37%	0.766
0.30	70.16%	0.767

Prominence-sensitivity analysis shows that the primary results do not depend heavily on one peak-detection threshold. Repeating feature extraction and model estimation for prominence ratios from 0.10 to 0.30 changes accuracy and AUC only modestly (Table 6). The reported accuracy and AUC are therefore not materially sensitive to the selected prominence ratio within this range.

5.6 Rolling out-of-sample validation

Expanding-window validation tests whether the primary 2025 holdout result reflects a single-year artifact. The model is trained on 2015 through the year before each test year and evaluated on the next calendar year from 2020 to 2025. Figure 5 reports rolling out-of-sample accuracy, whereas Figure 6 reports rolling out-of-sample AUC. Logit accuracy ranges from 69.21% to 75.83%, while logit AUC ranges from 0.768 to 0.831. The 13:30 price-position baseline remains stronger in every rolling test year, with accuracy from 81.61% to 86.19% and AUC from 0.893 to 0.944. Taken together, the rolling tests reduce the concern that the primary 2025 holdout result is driven by a single calendar year, while confirming the stronger performance of the same-cutoff benchmark across test years.




Figure 5. Rolling out-of-sample accuracy, 2020–2025.




Figure 6. Rolling out-of-sample AUC, 2020–2025.

6. DISCUSSION

The main estimates are consistent with the hypothesized directional relationships, but they should not be interpreted as causal evidence. `BREAK_SIGNAL` has the largest coefficient and odds ratio, suggesting that a Zone 5 breakout of the morning range carries more directional information than the sign of either 30-minute window alone within the proposed three-signal specification. `ZONE2_SIGNAL` and `ZONE5_SIGNAL` are smaller but still statistically significant, indicating that the original heuristic contains measurable interval-pattern information for the final open-to-close sign. The logistic model provides a transparent probabilistic representation of these signals while preserving the original rule structure.

The stronger 13:30 price-position benchmark narrows how the primary results should be interpreted. In the 2025 holdout and rolling tests, this same-cutoff benchmark outperforms the three-signal logistic model, indicating that much of the final open-to-close direction is already summarized by the index's location relative to the daily open at 13:30. The proposed framework should therefore be viewed as complementary to, rather than competing with, the same-cutoff price-position benchmark. This comparison also suggests that intraday directional performance should be interpreted relative to the information state available at the decision time. Once a substantial portion of the day's direction has already been incorporated into the price state, the marginal contribution of additional heuristic features may naturally become smaller. The weaker performance relative to this benchmark is itself informative because it indicates that the heuristic components capture part of the market information already reflected in the observed price state, rather than providing an independent superior predictor.

The post-13:30 analysis provides a useful boundary check. Once the target is redefined as the remaining-session move from 13:30 to the close, the three-signal framework no longer shows stable ranking ability and does not outperform simple same-cutoff alternatives. This distinction supports interpreting the framework as an open-to-close diagnostic representation of the intraday price state rather than as a reliable forecast of the remaining session.

This interpretation is consistent with a market-microstructure view of intraday price formation. Rather than being incorporated only at the market close, information is progressively reflected through trading activity, liquidity conditions, trading frictions, and time-of-day effects during the session. Recent evidence distinguishes intraday from overnight return formation ^[21], shows that trading frictions can affect price discovery and information acquisition ^[22], and links intraday predictability to gradual information incorporation and liquidity-related factors ^[23]. In this context, the proposed signals provide an interpretable decomposition of the intraday process through which the observed 13:30 price state develops, distinguishing early-session direction, post-lunch continuation, and breakout confirmation.

The ordered probability groups further support this diagnostic interpretation. Although the three-signal model does not outperform the 13:30 price-position benchmark, its probability bands generally rank intraday states by directional strength rather than only produce a hard directional call. The wide interval for the small 20–40% group nevertheless cautions against interpreting the grouping as precise calibration evidence. This emphasis on interpretable signal decomposition is consistent with broader financial modeling literature that stresses transparency and auditability in predictive settings ^[24].

The Chinese A-share market gives the heuristic a specific institutional meaning: the midday trading break makes the 13:01–13:30 window the first half-hour after trading resumes. However, this institutional setting also limits direct transfer to markets with different trading hours, liquidity regimes, or participant structures.

Several limitations define the scope of the evidence. The main results support real-time open-to-close directional prediction and an interpretable signal structure, while the supplementary test does not establish reliable prediction of the post-13:30 direction. Neither result establishes trading profitability or rejects market efficiency. Directional accuracy may not translate into economic value after commissions, bid-ask spreads, execution delays, slippage, position sizing, market impact, and risk constraints. The study also does not include order-book depth, volume-based features, or liquidity measures, and the cross-index transfer evidence is limited to another Chinese equity index. Applying the framework to other markets would require redefining local session windows, recalculating breakout ranges, and repeating the leakage, cross-index transfer, and sensitivity checks.

7. CONCLUSION

This study examined whether a practitioner-style intraday heuristic can be transformed into a reproducible statistical framework under a fixed real-time information boundary. By converting the Zone 2–Zone 5 judgment into three predefined signals, the primary analysis uses information available at 13:30 each trading day to predict whether the final close will be above or below the daily open. This open-to-close classification task treats discretionary trading intuition as an object of empirical testing rather than as an informal rule of thumb. A supplementary target instead isolates the direction from 13:30 to the close, providing a forward test of the remaining session.

The findings suggest that different stages of the trading session contribute unevenly to end-of-day direction. In the primary analysis, BREAK_SIGNAL carries the strongest directional content, while Zone 2 and Zone 5 directions provide smaller but statistically significant incremental information. The stronger 13:30 price-position benchmark shows that much of the final open-to-close direction is already reflected in the index's location relative to the daily open at the decision time, but the three-signal framework still provides an interpretable decomposition of that intraday state. The supplementary forward test does not support reliable prediction of the post-13:30 direction and therefore defines the boundary, rather than replacing the conclusion, of the primary analysis.

More broadly, the study illustrates how qualitative market judgment can be converted into transparent statistical representations under explicit information boundaries. Once formalized, such heuristics can become reproducible research objects that later studies can compare, refine, or reject. In this sense, the paper contributes not only to intraday directional prediction, but also to the methodological problem of how discretionary market judgment can be formalized without losing interpretability.

The analysis remains limited in important ways. It does not establish trading profitability, causal mechanisms, or market-wide generality, and it does not incorporate transaction costs, order-book depth, volume-based variables, or execution constraints. Future research may extend this formalization framework in three directions. First, future work may examine

whether practitioner-defined signals provide incremental information beyond richer same-cutoff market-state variables. Second, it may adapt the framework to markets with different trading mechanisms and institutional structures. Third, it may apply the same methodology to other practitioner-defined heuristics, including support-and-resistance judgments, failed breakouts, and volume-confirmation rules. Such extensions may help establish whether discretionary market knowledge can be systematically transformed into reproducible empirical objects while maintaining explicit information boundaries, interpretability, and rigorous out-of-sample validation.

Systematic signal engineering and disciplined validation can build a reproducible bridge between discretionary trading intuition and empirical financial research. Practitioner knowledge need not remain anecdotal when its definitions are fixed, its information boundary is explicit, and its claims are tested against strong alternatives.

DATA AVAILABILITY

We obtained the minute-level data from MyQuant, subject to the provider's terms. The author will provide the algorithm specification, analysis code, software-version specification, and instructions for reproducing both outcome definitions upon reasonable request.

REFERENCES

- [1] E.F. Fama, Efficient capital markets: A review of theory and empirical work, *Journal of Finance* 25(2) (1970) 383–417. <https://doi.org/10.2307/2325486>
- [2] D. Kahneman, A. Tversky, Prospect theory: An analysis of decision under risk, *Econometrica* 47(2) (1979) 263–291. <https://doi.org/10.2307/1914185>
- [3] N. Barberis, A. Shleifer, R. Vishny, A model of investor sentiment, *Journal of Financial Economics* 49(3) (1998) 307–343. [https://doi.org/10.1016/S0304-405X\(98\)00027-0](https://doi.org/10.1016/S0304-405X(98)00027-0)
- [4] W. Brock, J. Lakonishok, B. LeBaron, Simple technical trading rules and the stochastic properties of stock returns, *Journal of Finance* 47(5) (1992) 1731–1764. <https://doi.org/10.1111/j.1540-6261.1992.tb04681.x>
- [5] R. Sullivan, A. Timmermann, H. White, Data-snooping, technical trading rule performance, and the bootstrap, *Journal of Finance* 54(5) (1999) 1647–1691. <https://doi.org/10.1111/0022-1082.00163>
- [6] A.W. Lo, H. Mamaysky, J. Wang, Foundations of technical analysis: Computational algorithms, statistical inference, and empirical implementation, *Journal of Finance* 55(4) (2000) 1705–1765. <https://doi.org/10.1111/0022-1082.00265>
- [7] M. O'Hara, *Market Microstructure Theory*, Blackwell, Oxford, 1995
- [8] J. Hasbrouck, Measuring the information content of stock trades, *Journal of Finance* 46(1) (1991) 179–207. <https://doi.org/10.1111/j.1540-6261.1991.tb03749.x>
- [9] T.G. Andersen, T. Bollerslev, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance* 4(2–3) (1997) 115–158. [https://doi.org/10.1016/S0927-5398\(97\)00004-2](https://doi.org/10.1016/S0927-5398(97)00004-2)
- [10] S.L. Heston, R.A. Korajczyk, R. Sadka, Intraday patterns in the cross-section of stock returns, *Journal of Finance* 65(4) (2010) 1369–1407. <https://doi.org/10.1111/j.1540-6261.2010.01573.x>
- [11] P.C. Jain, G.H. Joh, The dependence between hourly prices and trading volume, *Journal of Financial and Quantitative Analysis* 23(3) (1988) 269–283. <https://doi.org/10.2307/2331067>
- [12] H. Zhu, Z.Q. Jiang, S.P. Li, W.X. Zhou, Profitability of simple technical trading rules of Chinese stock exchange indexes, *Physica A: Statistical Mechanics and its Applications* 439 (2015) 75–84. <https://doi.org/10.1016/j.physa.2015.07.032>
- [13] C.H. Park, S.H. Irwin, What do we know about the profitability of technical analysis?, *Journal of Economic Surveys* 21(4) (2007) 786–826. <https://doi.org/10.1111/j.1467-6419.2007.00519.x>
- [14] D.W. Hosmer, S. Lemeshow, R.X. Sturdivant, *Applied Logistic Regression*, third ed., Wiley, Hoboken, 2013. ISBN: 978-0-470-58247-3
- [15] M.T. Leung, H. Daouk, A.S. Chen, Forecasting stock indices: A comparison of classification and level estimation models, *International Journal of Forecasting* 16(2) (2000) 173–190. [https://doi.org/10.1016/S0169-2070\(99\)00048-5](https://doi.org/10.1016/S0169-2070(99)00048-5)
- [16] L. Blume, D. Easley, M. O'Hara, Market statistics and technical analysis: The role of volume, *Journal of Finance* 49(1) (1994) 153–181. <https://doi.org/10.1111/j.1540-6261.1994.tb04424.x>

- [17] D. Easley, M. O'Hara, Time and the process of security price adjustment, *Journal of Finance* 47(2) (1992) 577–605. <https://doi.org/10.1111/j.1540-6261.1992.tb04402.x>
- [18] T. Chordia, R. Roll, A. Subrahmanyam, Order imbalance, liquidity, and market returns, *Journal of Financial Economics* 65(1) (2002) 111–130. [https://doi.org/10.1016/S0304-405X\(02\)00136-8](https://doi.org/10.1016/S0304-405X(02)00136-8)
- [19] P. Virtanen, et al., SciPy 1.0: Fundamental algorithms for scientific computing in Python, *Nature Methods* 17 (2020) 261–272. <https://doi.org/10.1038/s41592-019-0686-2>
- [20] F. Pedregosa, et al., Scikit-learn: Machine learning in Python, *Journal of Machine Learning Research* 12(85) (2011) 2825–2830. <https://jmlr.org/papers/v12/pedregosa11a.html>
- [21] V. Bogousslavsky, The cross-section of intraday and overnight returns, *Journal of Financial Economics* 141(1) (2021) 172–194. <https://doi.org/10.1016/j.jfineco.2020.07.020>
- [22] C.M.C. Lee, E.M. Watts, Tick Size Tolls: Can a Trading Slowdown Improve Earnings News Discovery?, *The Accounting Review* 96(3) (2021) 373–401. <https://doi.org/10.2308/tar-2018-0689>
- [23] S. Aleti, T. Bollerslev, M. Siggaard, Intraday Market Return Predictability Culled from the Factor Zoo, *Management Science* 71(9) (2025) 7731–7751. <https://doi.org/10.1287/mnsc.2023.01657>
- [24] J. Černevičienė, A. Kabašinskas, Explainable artificial intelligence (XAI) in finance: A systematic literature review, *Artificial Intelligence Review* 57, Article 216 (2024). <https://doi.org/10.1007/s10462-024-10854-8>