

# Date-Level Financial Temporal Anchoring and Anchor-Centric Priors for News-Enhanced Overnight Stock Return Prediction

Huajie Chen

Maynooth International Engineering College, Fuzhou University, Fuzhou, 350108, China  
1520030858@qq.com

## ABSTRACT

This research explores how news contributes to predicting stock returns overnight, specifically when the timestamps have limited precision. Data from public news datasets are typically only dated or normalized so they can produce a false signal when they have been returned to the market at an exact release time. To mitigate this problem, the conservative Date-Level Financial Temporal Anchors protocol is proposed to anchor news articles to tradable, pre-open prediction anchors using the U.S. market time zone, the trading calendar, and parsed timestamps, without the intent to recover precise release times to the minute. This anchored representation is then used to provide bounded article-level weights and statistics from event type, sentiment confidence and negative asymmetry, attention shock, and sentiment consensus or dispersion to develop an Anchor-Centric mechanism. Experiments were run with price data for nine U.S. stocks from 2018 to 2023 comparing to the FNSPID financial news dataset using anchored news features under similar target definition, splits, and configuration improved the majority of correlation, ranking and ranking direction, and portfolio-type diagnostic metrics. In the main LSTM case, Anchor-Centric improved Pearson correlation, Daily IC, Rank IC, AUC, and LS Top-3 Sharpe relative to the anchored-news version. Delay-control experiments found that artificial anchor delays and news signal decay were negatively correlated. The results build a strong case that temporal misalignment negatively impacts predictive signals based on financial news, and that conservative finite anchoring with prior distribution based on financial variables improves the alignment of financial news to overnight return labels given limited timestamp precision.

**Keywords:** Financial News; Stock Return Prediction; Temporal Anchoring; Anchor-Centric Prior; Financial Temporal Alignment; Overnight Return

## 1. INTRODUCTION

Predicting stock returns is a major focus of financial data mining and quantifying market analysis. Historical prices, volume, and technical indicators are the basis of traditional stock prediction systems. The structured nature of these market structures can illustrate previous movement of stock prices, volatility, momentum, and volume. However, external shocks to the economy cannot be reflected directly in these structured market variables. External shocks such as earnings, analyst ratings, regulatory investigations, product releases, mergers, management changes, and macroeconomic policy changes will quickly change the expectations of investors in an actual financial market. Therefore, financial news is a critical source of unstructured information for stock prediction analysis <sup>[1,2]</sup>.

Recent advances in natural language processing have enabled the conversion of financial text data to numeric features. Models like BERT and FinBERT extract dense embeddings and sentiment probabilities from news articles <sup>[3,4]</sup>. Text features can be combined with price features to create news-enhanced stock prediction models. Most previous research has been concerned with either how to represent text, how to extract sentiment from text, or how to combine models. The value of the news features depends not only on how the text is represented, but also on whether the articles are associated with the prediction targets that are available in the market.

It is particularly important to consider the timing of financial news for overnight return predictions because stock returns are organized by trading days, but financial news can be released before market open, during the day, after market close, weekend, or on market holidays. Assigning news to the wrong trading-day label can result in a model learning a weaker or even an incorrect relationship between text signals and stock returns. This problem is compounded when public financial news datasets have incomplete, date-level, or normalized timestamps; in these situations, treating all recorded timestamps as the exact release time can create spurious precision with date-times and can result in questionable feature-label alignment.

This study investigates using news to improve estimating adjusted overnight stock returns before the markets open. The variable being forecasted is the previous adjusted close to the current adjusted market open. For any day being forecasted, price information available before that day and news assigned to be used to make a predication of the price of the security before the market opens. The goal is not to recover the true time each article was published at the minute level, but a protocol that creates a conservatively and repeatably anchored date for the news signal to provide better predicted alignment of the financial news signal to the overnight return labels.

The Date-Level Financial Temporal Anchoring protocol outlined in this paper provides a way to process available timestamps, convert them into the U.S. market timezone, apply a rule for pre-open anchors, and roll candidate anchors to valid trading days according to the trading calendar. Thus, news articles are mapped to tradable prediction anchors and are no longer treated simply as groups of calendar days. As a result, unsupported assertions about specific times of intraday news releases have been avoided while providing a clear information-availability rule for pre-opening predictions.

In addition to the news representation, the proposal includes an Anchor-Centric prior construction mechanism. This mechanism utilizes bounded article-level weights and anchor-level prior statistics to construct representations of financially motivated signals such as event importance, sentiment confidence, negative asymmetry, attention shock, sentiment consensus/dispersion. This is not intended to create a separate system for prediction of events, but rather to produce a more consistent aggregation of daily news to align with the intuitive understanding of finance. The Anchor-Centric design will allow for the assignment of differing levels of importance for different articles and provide a more complete description of the overall information condition for each stock-anchor pair as compared to simple averaging.

This research paper's primary research question is whether conservative financial temporal anchoring and Anchor-Centric prior construction will enhance the alignment of predictively associated financial news features and overnight return labels given the limited precision of timestamps. The experiment contains three controlled studies to investigate this question: a price-only baseline, a simple anchored-news method, and the proposed Anchor-Centric method. Each of these methods employs an identical target definition, data-split, price-features, feature standardization protocol, and LSTM main-setting as well as conducting delay-control studies to see if an artificial anchor misalignment weakens the productive anchor signal of the news.

The main contributions of this paper include the formulation of a date-level financial temporal anchoring problem to predict overnight returns based on news where timestamp precision is limited, development of a conservative anchoring protocol mapping financial news to tradable pre-open prediction anchors without claiming minute-level recovery of release time, creation of an Anchor-Centric prior construction mechanism via a combination of boundedly article-weighted news with anchor-weighted prior statistics prior to daily news aggregation, and experimental results from controlled studies using FNSPID financial news articles and nine United States stocks that collected from 2018 - 2023 demonstrate that anchored news features add predictive information beyond price only baseline models and that Anchor-Centric priors improve a variety of correlation, ranking, direction and portfolio diagnostic measures in the primary LSTM setting and controlled experiments confirm the significance of having temporal alignment among news articles for news-enhanced prediction of stock returns. The overall framework of the proposed method is illustrated in Figure 1.




Figure 1. Overall framework of the proposed method.

## 2. RELATED WORK

### 2.1 Price-based stock prediction

The use of historical market variables for forecasting future returns or movement directions is known as price-based stock prediction, with common inputs like one-day return, moving-average ratios, volatility, volume change and gap return. Commonly used traditional models include linear regression, support vector machines and tree-based methods. Deep

learning models, including recurrent neural networks, GRU, LSTM, TCN, and Transformer-based models, can process sequential dependencies of financial time series data and have been applied to this type of forecasting as well <sup>[5-9]</sup>.

The benefit of using price to predict market movements is that the data are structured and have time components. A drawback is that the price data is backward-looking. It tells how the market moved in the past but does not represent occurring events that affect investor sentiment. This provides motivation to use financial text as another form of data.

## 2.2 Financial news representation

Financial news has economic meaning expressed in natural language. Firm fundamentals, analysts opinions, legal risks, regulatory actions and partnerships are all described in financial news, including management changes and macroeconomic conditions. Early studies used dictionaries as sources for sentiment scores, or bag of word features, while recent studies use contextual language models to extract richer semantic and sentiment information.

FinBERT is widely applied to financial sentiment analysis because of its adaptation to financial language and its capacity to output positive, neutral, and negative probabilities <sup>[3]</sup>. When predicting stocks, sentiment can be established from the sum of the positive and negative probabilities, and dense embeddings can be used for semantic information, which typically consist of more than just sentiment polarity. Therefore, this paper uses both Financial BERT's sentiment features and PCA reduced dense text embedding features.

## 2.3 Text-price fusion and temporal alignment

News-enhanced stock forecasting uses structured market data and unstructured text data. The usual process is to extract the price data and text data then combine them using concatenation, sequence models, attention or gating. These systems can capture historical market data and outside news data; however, merging the features only makes sense if the news is time-based and aligned with the forecast target.

Temporal alignment is particularly important to predict overnight returns. A news article is not inherently useful because it mentions a stock, it must be dated with an anchor prediction that is the time frame for when the model can consider it. Many datasets might only have date level timestamps or normalized timestamps making this issue under-discussed; however this method attempts to overcome this practical alignment problem. The main categories of related work and their limitations are summarized in Table 1.

Table 1. Main categories of related work.

| Category                       | Typical input                        | Typical models                      | Main limitation                                         |
|--------------------------------|--------------------------------------|-------------------------------------|---------------------------------------------------------|
| Price-only prediction          | Prices, volume, technical indicators | Linear models, LSTM, GRU, TCN       | No direct external news information                     |
| Financial sentiment analysis   | Financial news text                  | BERT, FinBERT, classifiers          | Temporal mapping is often not the main focus            |
| News-enhanced stock prediction | Price features and news features     | Fusion networks, RNNs, Transformers | Fusion quality depends on temporal alignment            |
| Temporal alignment methods     | News dates and trading calendars     | Calendar rules and anchor mapping   | Date-level timestamp constraints need careful treatment |

## 3. METHOD

### 3.1 Overview

This framework includes four stages: temporal anchoring at the date-level, constructing features, building Anchor-Centric priors, and making a sequence prediction. The first stage is the collection of raw news records and stock prices in a single trading calendar. News articles are assigned to tradable pre-open prediction anchors rather than being placed in calendar date groupings. Feature values from price data and news from FinBERT will be created independently, and article-level prior weights and aggregate data will be used to aggregate the news and create the final fused stock anchor vector. Then the fused vectors will be formatted into fixed-length sequences for predicting adjusted overnight returns using an LSTM model.

### 3.2 Prediction target and information availability

Let  $s$  represent the stock,  $t$  the trading day, and the model predicts the adjusted overnight return of stocks on trading day  $t$ . The adjusted opening price for stock  $s$  is calculated using the close-to-adjusted-close ratio applied to the opening price, as the original price file contains open, close, and adjusted close values:

$$\text{AdjustedOpen}_{s,t} = \text{Open}_{s,t} \times \frac{\text{AdjustedClose}_{s,t}}{\text{Close}_{s,t}}. \quad (1)$$

The overnight return target is defined as:

$$y_{s,t}^{\text{ON}} = \frac{\text{AdjustedOpen}_{s,t}}{\text{AdjustedClose}_{s,t-1}} - 1. \quad (2)$$

This target is the return from the adjusted close price of the previous day to the adjusted open price of the current trading day. The model is based on the pre-open prediction method, which uses price features as of the end of day  $t-1$ , to create input for price features of day  $t$ . Thus, the method shifts the pre-open price feature matrix one trading day before constructing the price feature sequences to prevent information on opening prices on the target day from contaminating the model input.

### 3.3 Date-Level Financial Temporal Anchoring

Most financial news does not have accurate minute-level timestamps on when the news is released, so the normalized timestamps that show midnight do not provide a sound reference for true release times. The timestamp field is thus used as input into a deterministic protocol that provides a reproducible date-level prediction anchor, but does not provide the accurate intraday release time.

The timestamp for every article goes from UTC to America/New\_York. If the converted local date is a trading-day and the local time is earlier than 09:30, the current pre-open prediction anchor will be assigned to that article. If the converted local time is 09:30 or later, the current candidate will shift to the next calendar day and then roll to the next valid trading day. If the converted local date falls on a weekend or market holiday, the current candidate will also roll to the next valid trading day.

For records produced on dated timestamps, the converted time will be used only for creating a deterministic anchor, and not as true intraday release session evidence. Therefore, the parsed session statistics and freshness proxies will be secondary features. The primary claim is still pre-open anchoring at the comparable date-level, with limitations around timestamp accuracy. The resulting anchor-assignment rules are summarized in Table 2. The workflow of the date-level financial temporal anchoring protocol is illustrated in Figure 2.




Figure 2. Workflow of the date-level financial temporal anchoring protocol.

Table 2. Rules for date-level financial temporal anchor assignment.

| News condition                                                | Assigned prediction anchor | Explanation                                                                                       |
|---------------------------------------------------------------|----------------------------|---------------------------------------------------------------------------------------------------|
| Converted local date is a trading day and local time < 09:30  | Current trading anchor     | The article is assumed to be available for the pre-open prediction anchor of that trading day.    |
| Converted local date is a trading day and local time >= 09:30 | Next valid trading anchor  | The article is shifted forward because it is not assumed to be available before the current open. |

Table 2. (continued) Rules for date-level financial temporal anchor assignment.

| News condition                                      | Assigned prediction anchor | Explanation                                                     |
|-----------------------------------------------------|----------------------------|-----------------------------------------------------------------|
| Converted local date is a weekend or market holiday | Next valid trading anchor  | The article is rolled to the next tradable prediction anchor.   |
| Multiple articles share the same stock and anchor   | Same stock-anchor group    | The articles are aggregated into one daily news representation. |

### 3.4 Feature construction

For every stock-anchor pair, there will be a 15-dimensional price feature vector generated from historical market data where the features are descriptive of recent return, momentum, volatility, volume behavior, moving-average deviation, and daily price range. All price features will follow the pre-open protocol and are appropriately shifted such that the target-day open price will not be included in the input. In addition, all the price feature sets will be the same for all the methods being discussed. The price feature categories and their functions are summarized in Table 3.

Table 3. Price feature categories.

| Feature category       | Example features                          | Function                                                                   |
|------------------------|-------------------------------------------|----------------------------------------------------------------------------|
| Return features        | One-day return, log return                | Describe recent price movement                                             |
| Momentum features      | Five-day return, ten-day return           | Capture short-term trend                                                   |
| Volatility features    | Five-day, ten-day, twenty-day volatility  | Measure recent risk and fluctuation                                        |
| Volume features        | Log volume, volume change, volume z-score | Describe trading activity                                                  |
| Trend features         | Moving-average ratios                     | Measure deviation from short-term and medium-term trends                   |
| Range and gap features | High-low range, previous gap return       | Describe historical range and opening-gap behavior after the feature shift |

FinBERT encodes each financial news piece with an article-level representation consisting of a dense semantic embedding, as well as features related to financial sentiment. The traditional dense semantic embedding is reduced to 128 dimensions using PCA. To prevent potential data leakage, PCA is trained on the training set and is only applied to the validation and test sets after fitting. The resulting 133-dimensional article-level news feature composition is summarized in Table 4.

The sentiment score and sentiment confidence are defined as:

$$\text{Score}_i = p_i^{\text{pos}} - p_i^{\text{neg}} \quad (3)$$

$$\text{Conf}_i = \max(p_i^{\text{pos}}, p_i^{\text{neu}}, p_i^{\text{neg}}) \quad (4)$$

Table 4. Article-level news feature composition.

| Component               | Dimension | Description                                            |
|-------------------------|-----------|--------------------------------------------------------|
| PCA text embedding      | 128       | Dense semantic representation of the article           |
| Sentiment probabilities | 3         | Positive, neutral, and negative probabilities          |
| Sentiment score         | 1         | Difference between positive and negative probabilities |
| Sentiment confidence    | 1         | Maximum probability among the three classes            |
| Total                   | 133       | Complete article-level news representation             |

Different types of financial events can show varying levels of market relevance. Each article is classified according to a known event type using a set of rules. The event type is not an independent prediction target, but rather a prior signal for news weights and statistics at the anchor level. The event categories and representative signals are listed in Table 5.

Table 5. Financial event categories.

| Event type             | Description                            | Example signals                  |
|------------------------|----------------------------------------|----------------------------------|
| Earnings               | Earnings reports and financial results | earnings, revenue, profit, EPS   |
| Guidance               | Forward-looking management guidance    | guidance, outlook, forecast      |
| Analyst rating         | Rating and price target changes        | upgrade, downgrade, price target |
| Merger and acquisition | Acquisition, merger, and deal events   | acquisition, merger, takeover    |
| Regulation and legal   | Lawsuits and regulatory events         | lawsuit, investigation, SEC      |
| Product and business   | Product releases and partnerships      | launch, product, partnership     |
| Management change      | Executive or board changes             | CEO, CFO, resignation            |
| Capital action         | Dividends, buybacks, and offerings     | dividend, buyback, offering      |
| Macro and policy       | Macroeconomic and policy events        | Fed, inflation, interest rate    |
| General market         | General company or market news         | market, stock, company           |

### 3.5 Anchor-Centric prior construction

In this research, financial prior methods are transparent rule-based and sentiment-based signals. They encode event significance, sentiment confidence, downside asymmetry, attention shock, and sentiment agreement. Rather than building an event prediction system; Dews intends to create daily news aggregates that are better aligned with financial intuition however still allow for reproducibility.

After mapping articles to anchors, daily aggregation assigns each article a bounded prior weight. The configuration includes four factors at the article level: freshness, event importance, sentiment confidence, and negative asymmetry. Freshness is proxied by time proximity due to timestamp precision limitations. Article-level weight is:

$$\alpha_i = \text{clip}(q_i^{\text{fresh}} \times q_i^{\text{event}} \times q_i^{\text{conf}} \times q_i^{\text{neg}}, 0.5, 1.5) \quad (5)$$

The weight clipping operation restricts weight to a limited range and prevents low-weight articles from being totally removed from the daily representation. The factors of entity relevance, novelty, and source confidence are set as optional parameters in the pipeline but were disabled for the experiments reported in the paper; thus, their values are 1.0 in the configuration of their corresponding papers.

Attention shock and sentiment consensus/dispersion are not multiplied into existing article-level weight but calculated at anchor level prior and appended to the corresponding anchor vector. Therefore, the anchor-level description and article-level weight are kept separate. The levels and roles of the reported prior factors are summarized in Table 6. The construction process of the Anchor-Centric prior module is shown in Figure 3.




Figure 3. Anchor-Centric prior construction module.

Table 6. Prior factors in the reported Anchor-Centric configuration.

| Prior factor                    | Level                     | Role in the method                                                                                             |
|---------------------------------|---------------------------|----------------------------------------------------------------------------------------------------------------|
| Freshness proxy                 | Article level             | A bounded parsed-time proximity signal between the article and its assigned anchor; not exact intraday recency |
| Event importance                | Article level             | A rule-based weight based on the financial event category                                                      |
| Sentiment confidence            | Article level             | A weight based on the confidence of the FinBERT sentiment output                                               |
| Negative asymmetry              | Article level             | A weight that emphasizes dominant negative sentiment when downside information is stronger                     |
| Attention shock                 | Anchor level              | A prior statistic that describes abnormal news attention at the stock-anchor level                             |
| Consensus / dispersion          | Anchor level              | A statistic that describes agreement or dispersion among article sentiment signals                             |
| Entity, novelty, source factors | Disabled optional factors | Implemented as optional factors but set to 1.0 in the reported experiments                                     |

### 3.6 Aggregation, fusion, and sequence prediction

For stock  $s$  and anchor  $t$ , let  $N_{\{s,t\}}$  denote the set of articles assigned to the stock-anchor pair. The weighted daily news representation is computed as:

$$n_{s,t} = \frac{\sum_{i \in N_{s,t}} \alpha_i x_i}{\sum_{i \in N_{s,t}} \alpha_i} \quad (6)$$

Where  $x_i$  is a FinBERT feature vector at the article level and  $\alpha_i$  is the article level prior weight. If there are no articles assigned to a stock-anchor pairing, then a zero news vector and news count will be returned. Thus creating the final stock-anchor feature vector will occur using the following process:

$$z_{s,t} = [p_{s,t}; n_{s,t}; a_{s,t}] \quad (7)$$

where  $p_{\{s, t\}}$  is the shifted price feature vector,  $n_{\{s,t\}}$  is the aggregated news vector, and  $a_{\{s,t\}}$  is the anchor-level prior statistic vector. The input dimensions of the three compared methods are summarized in Table 7.

Table 7. Input dimension breakdown of compared methods.

| Component                                  | Price-only | Simple Anchor | Anchor-Centric |
|--------------------------------------------|------------|---------------|----------------|
| Price features                             | 15         | 15            | 15             |
| PCA text embedding                         | 0          | 128           | 128            |
| FinBERT sentiment features                 | 0          | 5             | 5              |
| Daily news statistics                      | 0          | 5             | 5              |
| Sentiment-score extrema                    | 0          | 2             | 2              |
| Event-count features                       | 0          | 10            | 10             |
| Parsed-session count features              | 0          | 4             | 4              |
| Additional Anchor-Centric prior statistics | 0          | 0             | 21             |
| Final input dimension                      | 15         | 169           | 190            |

Simple Anchor is an anchored-news baseline that produces only aggregate descriptive statistics and does not use any article-level prior weighted down or additional Anchor-Centric prior statistical vectors. This definition keeps the comparison focused on the incremental nature of Anchor-Centric prior construction.

For a stock  $s$  and target day  $t$ , the prediction model uses a fixed-length lookback window. The input sequence consists of the fused feature vectors for the  $L$  trading anchors which end at tick time  $t$ :

$$X_{s,t} = [z_{s,t-L+1}, z_{s,t-L+2}, \dots, z_{s,t}] \quad (8)$$

The price in  $z_{s,t}$  is based on market history until  $t-1$ , and the news contains articles linked to  $t$ 's pre-open anchor. Therefore,  $z_{s,t}$  will be treated as a pre-open feature vector instead of a post-open observation, and it allows the model to learn:

$$\hat{y}_{s,t}^{ON} = f(X_{s,t}) \quad (9)$$

The primary model is a two-layer LSTM and feeds the last hidden representation into a fully connected dense layer to predict the overnight return. The main LSTM model settings are summarized in Table 8. The model is trained via mean squared error:

$$L = \frac{1}{M} \sum_{j=1}^M (\hat{y}_j^{ON} - y_j^{ON})^2 \quad (10)$$

Table 8. Main LSTM model settings.

| Parameter             | Value              |
|-----------------------|--------------------|
| Lookback window       | 20 trading days    |
| Number of LSTM layers | 2                  |
| Hidden size           | 128                |
| Dropout               | 0.2                |
| Learning rate         | 0.001              |
| Epochs                | 20                 |
| Batch size            | 64                 |
| Random seeds          | 42, 43, 44, 45, 46 |

## 4. EXPERIMENTS AND RESULTS

### 4.1 Dataset and experimental protocol

FNSPID financial news, as well as the associated stock price data, are used for the experiment as data sources <sup>[10]</sup>. There are nine United States stocks included in the stock pool: AMD, CRM, CVX, FDX, INTC, KO, MRK, ORCL, and QCOM. The time period for the data starts on 2018-01-02 and ends on 2023-12-28. Records of the financial news and the stock price data are processed into stock-anchor samples using the parsed-time pre-open anchoring protocol.

The data has many news timestamps with limited precision or normalized. Therefore, the date-level financial time anchoring methodology is used in this paper, which does not have precise intraday news release modeling. The data split uses global dates as the basis for splitting. The end of the training period is at 2021-12-31; the validation period is for 2022, and the test period is for 2023. All feature standardization is fitted using the training set only and applied to the validation set and test set subsequently. The resulting dataset statistics are summarized in Table 9.

Table 9. Dataset statistics.

| Item                               | Value  |
|------------------------------------|--------|
| Raw news articles                  | 47,460 |
| Successfully aligned news articles | 47,446 |
| Daily stock-news samples           | 11,696 |

Table 9. (continued) Dataset statistics.

| Item                              | Value  |
|-----------------------------------|--------|
| Daily stock panel rows            | 13,572 |
| Non-empty overnight return labels | 13,563 |
| Training sequence samples         | 8,712  |
| Validation sequence samples       | 2,259  |
| Test sequence samples             | 2,241  |

The supervised sequence samples are fewer than the number of rows in the panel because there is no overnight label for the first day of trades of each stock, the rolling price features will only work with historical observations and LSTM inputs must consist of a finite 20 days feature window. The date ranges of the training, validation, and test sets are listed in Table 10.

Table 10. Data split.

| Split          | Period           |
|----------------|------------------|
| Training set   | Until 2021-12-31 |
| Validation set | 2022             |
| Test set       | 2023             |

#### 4.2 Compared variants and evaluation metrics

The three price-only, simple anchor, and Anchor-Centric method variants are examined in this study. The price-only uses just the 15-dimensional price feature vector that is shifted. The simple anchor uses a pre-open temporal anchoring protocol, descriptive news aggregation statistics, and does not use article-level prior weightings or the extra Anchor-Centric prior statistics that were used. The Anchor-Centric method uses price features, anchored news features, article-level prior weightings, and anchor-level financial prior statistics. All three methods use the same target definition, splits, standardization method, and main LSTM configuration.

Empirical claims can only be made for these controlled variants, as well as delay-control settings. A pre-open natural-date baseline has yet to be rerun, therefore we cannot claim the proposed anchor is globally optimal over all calendar-date aggregation rules. The three controlled method variants are summarized in Table 11.

Table 11. Compared method variants.

| Method         | Input information                          | Description                                                                                                                                     |
|----------------|--------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------|
| Price-only     | Historical price features                  | Uses only shifted technical price features                                                                                                      |
| Simple Anchor  | Price features and anchored news           | Uses the same pre-open temporal anchor with basic descriptive news aggregation, but without Anchor-Centric weighting and extra prior statistics |
| Anchor-Centric | Price, anchored news, and financial priors | Adds bounded article-level prior weights and anchor-level prior statistics on top of anchored news                                              |

Metrics of evaluation include regression error, correlations, cross-sectional ranking, directional discrimination, and conventional portfolio style diagnostics. RMSE measures prediction errors, Pearson correlation measures linear associations, daily IC and Rank IC measure cross-sectional predictive correlations and rankings, and the AUC measures directional discrimination. LS Top-3 Sharpe, Top-Bottom Spreads, LS Maximum Drawdowns, Turnover, and Large Move Hit Rates provide evidence of portfolio-style diagnostics. Because only nine stocks are in the pool, and transaction costs have not been applied, portfolio metrics are treated as Rank Diagnostics and not as evidence of actual trading profitability.

#### 4.3 Main predictive results

The test set has main predictive results summarized in Table 12. The reported values are means and standard deviations over five random seeds (42–46). The two news-enhanced variants show observed increases across Pearson correlation,

Daily IC, Rank IC, and LS Top-3 Sharpe over Price-only suggesting that anchored financial news has incremental predictive value with respect to the given specifications. With respect to the main LSTM setting, the Anchor-Centric gives more improvements to Pearson, Daily IC, Rank IC, AUC, and LS Top-3 Sharpe than Simple Anchor. The gain for AUC is small and the RMSE are similar across methods. Hence, much of the evidence will be interpreted using correlation, ranking and portfolio-style diagnostic metrics instead of using RMSE alone. A visual comparison of the main predictive metrics is provided in Figure 4.

Table 12. Main predictive results on the test set.

| Method         | RMSE ↓          | Pearson ↑     | Daily IC ↑    | Rank IC ↑     | AUC ↑         | LS Top-3 Sharpe ↑ |
|----------------|-----------------|---------------|---------------|---------------|---------------|-------------------|
| Price-only     | 0.0123 ± 0.0001 | 0.019 ± 0.020 | 0.014 ± 0.014 | 0.029 ± 0.013 | 0.518 ± 0.007 | -0.45 ± 0.39      |
| Simple Anchor  | 0.0124 ± 0.0001 | 0.081 ± 0.023 | 0.063 ± 0.020 | 0.052 ± 0.017 | 0.520 ± 0.009 | 1.95 ± 0.78       |
| Anchor-Centric | 0.0122 ± 0.0001 | 0.124 ± 0.042 | 0.093 ± 0.030 | 0.078 ± 0.018 | 0.529 ± 0.003 | 2.40 ± 1.08       |




Figure 4. Main predictive metrics. AUC is included as a directional metric and should be interpreted together with correlation and ranking metrics.

#### 4.4 Portfolio-style diagnostic results

Portfolio-style diagnostics evaluate the overall usefulness of predicted scores for cross-sectional ordering. Based on the results of the controlled experimental setting, Anchor-Centric's prior framework improves ranking quality with the highest Top-Bottom Spread, LS Top-3 Sharpe, and Large-Move Hit Rate. However, since there was also an increase in Turnover for the Anchor-Centric framework, the portfolio-style results should be considered additional diagnostic assessment for ranking purposes rather than concessionary evidence of deployable trading profit. The corresponding portfolio-style diagnostic values are reported in Table 13. The portfolio-style diagnostic results are visualized in Figure 5.

Table 13. Portfolio-style diagnostic results.

| Method         | Top-Bottom Spread | LS Top-3 Sharpe | LS Max Drawdown | Turnover      | Large-Move Hit Rate |
|----------------|-------------------|-----------------|-----------------|---------------|---------------------|
| Price-only     | 0.0001 ± 0.0003   | -0.45 ± 0.39    | -0.178 ± 0.048  | 0.390 ± 0.067 | 0.499 ± 0.016       |
| Simple Anchor  | 0.0016 ± 0.0005   | 1.95 ± 0.78     | -0.087 ± 0.020  | 0.523 ± 0.045 | 0.561 ± 0.013       |
| Anchor-Centric | 0.0020 ± 0.0007   | 2.40 ± 1.08     | -0.094 ± 0.029  | 0.591 ± 0.062 | 0.575 ± 0.017       |




Figure 5. Portfolio-style diagnostics. Top-Bottom Spread is multiplied by 1000 only for visualization.

#### 4.5 Prior component ablation

This ablation study investigates how one prior component affects a target, split, and model configuration, by testing each prior with Simple Anchor first and then aggregating all active article weights and anchor-level prior statistics to create the Anchor-Centric variant. Ultimately, the findings show that all of the prior components contribute towards improving different sets of metrics: Consensus/dispersion produced the largest single-module benefit of all the prior components, while attention shock served as a valuable auxiliary diagnostic prior, but did not independently improve all metrics.

Using full Anchor-Centric design produced best values across all metrics of Pearson, Daily IC, Rank IC, AUC, and LS Top-3 Sharpe for reported LSTM ablation; supporting the value of using complementary financial priors before aggregating daily news. Limitation of the claim is that this paper does not state that every prior factor will provide useful information; will improve given metric or model architecture. The complete prior-component ablation results are reported in Table 14.

Table 14. Ablation results of Anchor-Centric prior components.

| Variant                  | Pearson | Daily IC | Rank IC | AUC   | LS Top-3 Sharpe |
|--------------------------|---------|----------|---------|-------|-----------------|
| Simple Anchor            | 0.081   | 0.063    | 0.052   | 0.520 | 1.95            |
| + Event Importance       | 0.098   | 0.072    | 0.062   | 0.526 | 1.95            |
| + Sentiment Confidence   | 0.098   | 0.070    | 0.061   | 0.526 | 2.05            |
| + Negative Asymmetry     | 0.097   | 0.070    | 0.061   | 0.526 | 1.98            |
| + Attention Shock        | 0.095   | 0.070    | 0.059   | 0.521 | 1.90            |
| + Consensus / Dispersion | 0.108   | 0.081    | 0.069   | 0.523 | 2.17            |
| Anchor-Centric           | 0.124   | 0.093    | 0.078   | 0.529 | 2.40            |

#### 4.6 Delay-control experiment

The delay-control experiment determines how much the news signal responds to anchor assignment. In this experiment the Anchor-Centric news anchor has been artificially delayed by either one trading day or five trading days. The predictive signal is based upon the alignment between news and return labels; therefore, a delay will weaken the strength of the predictive signal. The results support this hypothesis. The daily IC was reduced from 0.093 to 0.006 after one-day and to 0.009 after five-day delays. The LS Top-3 Sharpe was also greatly reduced.

These findings provide evidence in support of the relevance of temporal alignment; however they should be interpreted as a test for sensitivity, they do not prove that the proposed anchor is globally optimal among all anchoring rules available. The corresponding delay-control results are reported in Table 15. The effect of artificial anchor delays on the predictive signal is illustrated in Figure 6.

Table 15. Delay-control experiment.

| Setting              | Daily IC | AUC   | LS Top-3 Sharpe |
|----------------------|----------|-------|-----------------|
| Original anchor      | 0.093    | 0.529 | 2.40            |
| Delay 1 trading day  | 0.006    | 0.495 | -0.44           |
| Delay 5 trading days | 0.009    | 0.502 | -0.04           |




Figure 6. Delay-control experiment. Artificial delays weaken the predictive signal.

#### 4.7 Model-family check

To limit reliance on one LSTM model, other models such as Ridge, LightGBM, MLP, GRU, LSTM, TCN, and Transformer-small are compared by additional tests. Daily IC of these models is shown in Table 16 and news-enhanced versions typically outperform price only. Anchor-Centric performs best with Ridge, GRU, LSTM and Transformer-small; Simple Anchor outperforms with LightGBM, MLP and TCN. The original difference between the LSTM values in Tables 12 and 16 resulted from averaging over five random seeds (42–46) and three random seeds (42–44), respectively, rather than from differences in model architecture, data partitioning, or optimization settings. For consistency, the LSTM row in Table 16 reproduces the five-seed results reported in Table 12.

The above demonstrates the benefit of using anchored news features across model families although the analysis does not suggest that Anchor-Centric is superior among all model families. Future research will need to address the relationship between financial prior features and model structure.

Table 16. Model-family check based on Daily IC.

| Model    | Daily IC (Price-only / Simple Anchor / Anchor-Centric) | Best variant   |
|----------|--------------------------------------------------------|----------------|
| Ridge    | 0.021 / 0.050 / 0.056                                  | Anchor-Centric |
| LightGBM | 0.015 / 0.171 / 0.165                                  | Simple Anchor  |
| GRU      | 0.031 / 0.064 / 0.076                                  | Anchor-Centric |

Table 16. (continued) Model-family check based on Daily IC.

| Model             | Daily IC (Price-only / Simple Anchor / Anchor-Centric) | Best variant   |
|-------------------|--------------------------------------------------------|----------------|
| LSTM              | 0.014 / 0.063 / 0.093                                  | Anchor-Centric |
| MLP               | -0.001 / 0.012 / 0.001                                 | Simple Anchor  |
| TCN               | 0.010 / 0.016 / 0.003                                  | Simple Anchor  |
| Transformer-small | 0.007 / -0.013 / 0.008                                 | Anchor-Centric |

Note: The LSTM row reproduces the five-seed main-experiment results reported in Table 12 (seeds 42–46). Other model-family results are reported as three-seed means from the model-robustness experiments (seeds 42–44).

## 5. DISCUSSION

### 5.1 What the results support

Two controlled conclusions support these findings. First, the proposed pre-open anchoring protocol finds that financial news features provide additional incremental information over price-only features. Price-only models represent momentum, reversal, volatility, and volume effects, but they do not have a direct representation of both firm’s events and market’s events. Moving from Price Only to Simple Anchor results in useful predictive information contained in anchored news features based on this findings.

Second, moving from Simple Anchor to Anchor-Centric indicates the need for transparent financial priors when aggregating daily news. Because a simple average considers all articles in a stock-anchor group to have equal value, the assumption is often too strong. Factors such as event importance, confidence in sentiment, negative asymmetry, attention shock, and agreement of sentiment provide light context for identifying informative anchor-day news from less informative anchor-day news signals.

### 5.2 Why temporal anchoring matters

The delay control experiment indicates that the news signal is influenced by anchor assignments. When anchors are delayed, the strength of the relationship between news features and overnight return labels is reduced; this supports the key point that news-enhanced predictions are not just a textual representation issue, but also an issue of temporal alignment.

The proposed protocol for Date-Level Financial Temporal Anchoring is intended for use with datasets whose exact release times are either not known or unreliable. In these circumstances, a deterministic parsed-time anchor is better defensible than any claim of exact minute level release-time recovery. The protocol is focused on mapping news to tradeable prediction anchors with the validity of the conclusions being restricted to date-based pre-open prediction.

### 5.3 Boundaries and limitations

The study has several limitations. The method of analysis is not its own news release minute measurement. It should be called date-level financial temporal anchoring, not intraday news modeling. The experimental method includes the nine-stock pool used in the analysis. The inclusiveness of the statistical results is controlled experimental evidence and does not represent the same type of evidence as factor validation of a large-scale study. The portfolio-style diagnostics of the results do not include transaction costs, slippage, liquidity constraints, or risk limits as should real trading profitability evidence.

The study assessed predictive alignment but does not demonstrate the causality of the effect of news on the price of the stock. The use of auxiliary proxies for parsed-session features and freshness estimates is limited as many timestamps have limited precision. The primary experimental analysis failed to incorporate a newly rerun pre-open natural-date baseline into the analysis, limiting the empirical conclusions to those determined from the reported controlled versions of the results and artificial delay-control tests conducted in the study.

## 6. CONCLUSION

The paper studies overnight stock return prediction of news articles under limited timestamp precision and proposes a conservative date or financial temporal anchoring protocol to map articles to tradable pre-open prediction anchors via parsed timestamp, U.S. market timezone and trading calendar. The protocol avoids any claims of a true minute level of

release time recovery. The authors then propose an anchor-centric mechanism to infer article level prior weights based on bounded article-level prior weighting and anchor-level prior statistics to use as daily news aggregation.

Tests on FNSPID financial reports and nine US stock types show that news-enhanced versions using a proposed pre-open anchoring protocol outperform prices alone, both concerning correlation, rankings, directionality, and portfolio style. In a typical LSTM application, Anchor-Centric moderately to greatly improves Pearson correlation, Daily IC, Rank IC, AUC, and LS Top-3 Sharpe over Simple Anchor. Delay-control test results support the conclusion that deliberate anchor delays reduce the predictive signal. These results support the finding that misalignment over time can degrade the predictive signal based on news, while using conservative anchoring via financial prior construction can improve alignment between financial news and overnight return labels where timestamp precision is limited.

Future work will focus on adding additional natural-date and alternative anchoring baselines before extending to a broader universe of stocks, using more precise timestamps whenever possible, more accurate event typings, and measuring portfolio performance with consideration of transaction costs and stricter risk constraints.

## ACKNOWLEDGMENTS

The author thanks the open-source dataset and model communities for providing financial news data resources and pretrained language models used in this study.

## REFERENCES

- [1] Xu, Y., Cohen, S.B. (2018) Stock movement prediction from tweets and historical prices. In: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics. Melbourne. pp. 1970-1979
- [2] Fama, E.F. (1970) Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25: 383-417
- [3] Araci, D. (2019) FinBERT: Financial sentiment analysis with pre-trained language models. arXiv:1908.10063
- [4] Devlin, J., Chang, M.W., Lee, K., et al. (2019) BERT: Pre-training of deep bidirectional transformers for language understanding. In: Proceedings of NAACL-HLT 2019. Minneapolis. pp. 4171-4186
- [5] Hochreiter, S., Schmidhuber, J. (1997) Long short-term memory. *Neural Computation*, 9: 1735-1780
- [6] Bao, W., Yue, J., Rao, Y. (2017) A deep learning framework for financial time series using stacked autoencoders and long-short term memory. *PLOS ONE*, 12: e0180944
- [7] Vaswani, A., Shazeer, N., Parmar, N., et al. (2017) Attention is all you need. In: *Advances in Neural Information Processing Systems* 30. Long Beach
- [8] Ke, G., Meng, Q., Finley, T., et al. (2017) LightGBM: A highly efficient gradient boosting decision tree. In: *Advances in Neural Information Processing Systems* 30. Long Beach
- [9] Cho, K., Van Merriënboer, B., Gulcehre, C., et al. (2014) Learning phrase representations using RNN encoder-decoder for statistical machine translation. In: Proceedings of EMNLP 2014. Doha. pp. 1724-1734
- [10] Dong, Z., Fan, X., Peng, Z. (2024) FNSPID: A comprehensive financial news dataset in time series. arXiv:2402.06698