

Numerical Claim Detection in Finance: A New Financial Dataset, Weak-Supervision Model, and Market Analysis

Agam Shah^{♡ †}, Arnav Hiray^{♡ †}, Pratvi Shah[♣], Arkaprabha Banerjee[♣], Anushka Singh[◇]
Dheeraj Eidnani[♡], Sahasra Chava[♣], Bhaskar Chaudhury[♣], Sudheer Chava[♡]

♡ Georgia Institute of Technology

♣ DA-IICT

◇ IIT-Kharagpur

♣ Fulton Science Academy

Abstract

In this paper, we investigate the influence of claims in analyst reports and earnings calls on financial market returns, considering them as significant quarterly events for publicly traded companies. To facilitate a comprehensive analysis, we construct a new financial dataset for the claim detection task in the financial domain. We benchmark various language models on this dataset and propose a novel weak-supervision model that incorporates the knowledge of subject matter experts (SMEs) in the aggregation function, outperforming existing approaches. We also demonstrate the practical utility of our proposed model by constructing a novel measure of *optimism*. Here, we observe the dependence of earnings surprise and return on our optimism measure. Our dataset, models, and code are publicly (under CC BY 4.0 license) available on GitHub¹.

1 Introduction

Earnings conference calls are a quarterly event where the company’s top executives provide performance reports of the company over the last quarter (3 months). Between the two earnings calls analyst from various financial institutions analyze and provide earnings estimates and recommendations. For example, Jegadeesh and Kim (2010) has documented that there is a significant stock market reaction to analysts’ recommendations (ratings). Recent insights, such as those presented by McLean et al. (2020), reveal that retail investors, often perceived as unsophisticated, exhibit responsiveness to analysts’ projections, underscoring the pivotal role of analysts’ reports in informing market participants. However, analyst ratings can be biased (Michaely and Womack, 1999; Corwin et al.,

2017; Coleman et al., 2021). Therefore it is important to understand whether the ratings are backed by strong numerical financial claims in the analyst’s report. Further, the sentences with a claim have a higher density of forward-looking information. As an application, extraction of numerical ESG claims from earnings call transcripts, can help better understand whether companies do walk the talk on their environment and social responsibility claims (Chava et al., 2021). These examples underscore the necessity of numerical claim detection in the finance domain, aligning with broader research efforts to ensure the accuracy and reliability of information sources.

A key component of this paper is the identification of Numeric Financial Sentences. Specifically, Numeric Financial Sentences include a financial term, a numeric value, and either a currency or percentage symbol. Chen et al. (2020) first introduced the categorization of sentences into ‘in-claim’ and ‘out-of-claim’ specifically in the Mandarin language. Expanding on their foundation, we define an ‘in-claim’ sentence as one presenting a speculative financial forecast. Conversely, an ‘out-of-claim’ sentence presents a numerical statement about a past event, transitioning from a mere claim to a confirmed fact. For clarity, ‘in-claim’ sentences can also be termed "financial forecasts" whereas ‘out-of-claim’ can be labeled as "established financials." Every Numeric Financial Sentence that is not a speculative financial forecast (in-claim) is then identified as an ‘out-of-claim’ sentence. Figure 1 illustrates the identification of Numeric Financial Sentences as well as distinguishing between “in-claim” and “out-of-claim” sentences.

A major challenge for building or training predictive models is the scarcity of labeled data (Zhang et al., 2021; Ratner et al., 2017). Supervised learning often involves a significant amount of manual labeling of data which is often infeasible for large datasets. In such scenarios, one can leverage weak-

Correspondence to Agam Shah [ashah482@gatech.edu]

† These authors contributed equally to this work

¹<https://github.com/gtfintechlab/fin-num-claim>.

Example of In-claim and Out-of-claim sentences.

S1: "We also continued to grow our total active installed base by adding new customers."

S2: "Adjusted operating margins of over 41% were above the midpoint of guidance, as we balanced our strategic investments with prudent discretionary spend."

S3: "In q2, we achieved a record \$4.39 billion in revenue, representing 15% year-over-year growth."

S4: "Operating income is expected between \$2.1 billion and \$3.6 billion."

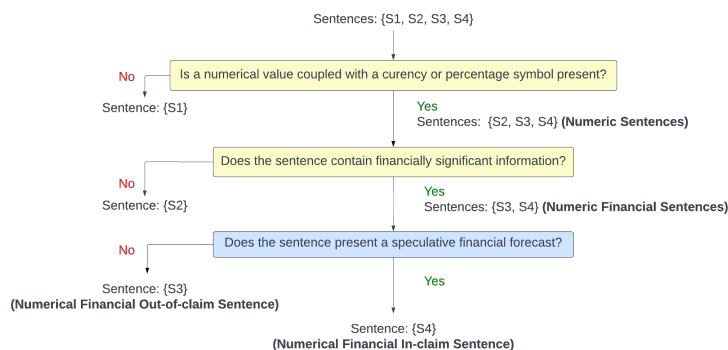


Figure 1: Example of In-claim and Out-of-claim sentences.

supervision-based learning methods (Varma and Ré, 2018) or fine-tune the pre-trained language model. Weak-supervision is a process that leverages slightly noisy or imprecise labeling functions (lfs) to label vast amounts of unlabeled data (Ratner et al., 2020; Lison et al., 2021). The strength of the weak-supervision model lies in these imperfect labels, when combined, producing improved predictive models (Lison et al., 2021; Zhang et al., 2021). However, a crucial component involves the development of effective lfs for a given raw dataset systematically rather than manual annotation (Lison et al., 2021).

The aim of our work is to derive financially significant information from the quarterly analyst reports and earnings calls by categorizing each numerical sentence as in-claim or out-of-claim. Our major contributions through this paper are the following:

- We introduce a new task of claim detection (in English) with a labeled dataset.
- We build clean, tokenized, and annotated open-source datasets based on earnings calls.
- We introduce a weak-supervision model with a novel aggregation function.
- We benchmark a wide range of language models for the claim detection task.
- We develop a novel measure of optimism and validate its usefulness in predicting various financial indicators.

2 Related Work

NLP in Finance Finance is one of the most attractive domains for the application of NLP. Araci (2019) and Liu et al. (2020) presented pre-trained language models for the Finance domain. There

are multiple datasets specifically catered for applications of NLP in finance including question answering dataset created by Chen et al. (2021) and Maia et al. (2018), and also a NER dataset constructed by Shah et al. (2023b) for the financial domain. There is a vast body of literature on undertaking sentiment analysis tasks on financial data (Maia et al., 2018; Malo et al., 2014; Day and Lee, 2016; Akhtar et al., 2017).

Works of Li et al. (2020) and Sawhney et al. (2020) were centered around predicting volatility using earnings call transcripts in the domain of risk management. Chava et al. (2022) measure the firm level inflation exposure by fine-tuning RoBERTa (Liu et al., 2019), while Li et al. (2021) leveraged word-embeddings to measure the corporate culture. Moreover, Nguyen et al. (2021) and Hu and Ma (2021) used multimodal machine learning for credit rating prediction and measurement of persuasiveness respectively. Shah et al. (2023a) investigated the impact of monetary policy communication on financial markets. Cao et al. (2020) critically examined the evolution of corporate disclosure in recent years, influenced by the rising application of NLP in Finance. Our research focuses on identifying numerical financial claims from a vast set of English analyst reports and earnings calls using a weak-supervision model. This differs from Chen et al. (2020), which targets numeric claim detection in a smaller Chinese language dataset.

Weak-Supervision In order to reduce the complexities associated with manual labeling, several standard techniques such as semi-supervised learning (Chapelle et al., 2009), transfer learning (Pan and Yang, 2010), and active learning (Settles, 2009) have been employed. However, many researchers (Meng et al., 2018; Kartchner et al., 2020) and practitioners also employ weak-supervision-based models to further reduce the computational costs while

retaining the accuracy of the labeled data. Weak-supervision models were primarily developed in a bid to replace standard labeling techniques with models which can leverage slightly noisy or imprecise sources to label vast amounts of data (Ratner et al., 2020). Techniques such as distant supervision (Mintz et al., 2009) and crowd-sourced labels (Yuen et al., 2011) are often associated with weak-supervision-based models, however, they tend to have limited coverage and accuracy (Lison et al., 2021). In the case where we have noisy labels from multiple sources available, there have been efforts made to use majority vote, weighted majority vote (Ratner et al., 2020), and other label-models (Yu et al., 2022; Zhang et al., 2022).

3 Dataset

We collect two categories of text and financial market datasets. Analyst reports are procured from a proprietary source while earnings call transcripts are collected in a manner that allows us to make the resulting dataset open-source.

3.1 Analyst Reports

The raw dataset consists of quarterly analyst reports (in English) for a large number of public firms in the U.S. These analyst reports were collected from Zacks Equity Research and were available to us through the Nexis Uni license².

The text documents are first split into sentences using multiple regex-based rules. This segmentation process utilizes a comprehensive set of regular expression (regex) rules to accurately identify sentence boundaries, accounting for a variety of English language nuances, including abbreviations, titles, websites, and numerical expressions, to ensure precise sentence delineation. We employ regex-based rules as they typically are significantly faster with similar accuracy compared to standard libraries in sentence tokenization. Next, sentences containing quantitative data - specifically sentences with a numeric value AND either a currency symbol as a prefix or percentage symbol as a postfix are extracted, as they have numerical relevance (Chen et al., 2019). This numerical condition filter reduced the number of sentences by 66.7%.

The next step in the pipeline uses a whitelisting technique to retain only sentences with finan-

²Nexis Uni license doesn't authorize republication of full or partial text. To solve this problem, we also collect and construct a dataset from earnings calls which can be made public under CC BY 4.0 license.

cially significant information, achieved by cross-referencing each sentence with a financial dictionary containing a comprehensive list of financial market terms and related literature. The financial dictionary used in this study, developed by Shah et al. (2022), contains over 8,200 financially significant terms. Sentences are cross-referenced with this dictionary to verify financial significance; if no words match, the sentence is marked as irrelevant. This filtering reduced the dataset by an additional 17.2%. The dataset contains 8,583,093 total sentences, 2,857,567 numeric sentences, and 2,364,977 numeric-financial sentences after filtration. This two-tier filtering method enriched the data by retaining only 27.5% of the sentences from the original data.

3.2 Earnings Call Transcripts

To make our work more impactful, we also collect earnings call transcripts for NASDAQ 100 companies from their investor relation page. We were able to write individual scripts for 78 out of 100 NASDAQ companies. As all the companies in this list are public companies, their data can be accessed and shared publicly which allows us to open-source the resulting dataset. Collecting data till March of 2023 results in a total of 1,085 earnings call transcripts. The biggest advantage of writing separate scripts for each company is that it allows us to keep adding more transcripts every quarter increasing the size of the dataset shared over time. We apply text processing (tokenization, numerical filter, financial dictionary filter) on earnings call transcripts similar to what is used for analyst reports.

3.3 Comparison with Related Dataset

In this section we compare our proposed datasets with NumClaim (Chen et al., 2020), an expert-annotated dataset in the Chinese language. Our dataset of raw analyst reports in the English Language from 1,530 major companies over the period of 2017-20 is significantly larger than NumClaim or other associated datasets. Our open-sourced dataset from collected earnings call transcripts is also larger than the NumClaim dataset. The detailed comparison of our datasets with NumClaim is provided in Table 1.

3.4 Financial Market Data

Stock Price and Earnings Surprise Data We collect stock price data from Polygon.io³ starting

³<https://polygon.io/stocks>

Dataset	Analyst Reports	Earnings Calls	NumClaim (Chen et al., 2020)
Language	English	English	Chinese
Year	2017-20	2017-23	NA
Sector Information	Yes	Yes	No
# Stocks	1,530	78	NA
# Files	87,536	1,085	NA
# Words	167,301,873	11,641,673	42,594
# Numeric Sentences	2,857,567	48,686	5,144
# Numeric Financial Sentences	2,364,977	41,013	NA
# Numeric Financial In-Claim Sentences	336,252	5362	1,233

Table 1: Comparison of our datasets with NumClaim (Chen et al., 2020) dataset.

January 1st, 2017. We collect the actual earnings per share (EPS) and forecasted median EPS from the I/B/E/S dataset⁴.

Sector Data For each firm in our dataset, we collect sector information by collecting GSECTOR classification from the annual fundamental COMPUSTAT database. GSECTOR maps each company to one of the twelve sectors.

3.5 Sampling and Manual Annotation

From the complete raw dataset of 87,536 analyst reports and 1,085 earnings call transcripts, we sample data and annotate sentences. The sampled dataset consisted of 96 analyst reports consisting of two files per sector per year, accounting for about 2,681 unique financial-numeric sentences. We also sample 12 earnings call transcripts randomly consisting of two files per year, consisting of 498 financial-numeric sentences. This set was manually annotated and assigned ‘in-claim’ or ‘out-of-claim’ labels by two of the authors with a foundational background in finance (one of them is now an analyst at a top investment bank) and domain expertise developed through examples provided by a co-author. This co-author is a financial expert with a Master’s degree in Quantitative Finance, currently pursuing a PhD under the guidance of the Chair Professor of Finance, and has contributed to work at leading finance journals and conferences. The annotator agreement was 99.21% and 95.78% for analyst reports and earnings call transcripts respectively. Any disagreement between the two annotators was resolved with the help of the financial expert mentioned earlier. The dataset (Train, Val, Test) is split as follows: Analyst Reports (1,715, 429, 537) and Earnings Calls (318, 80, 100).

⁴<https://www.investopedia.com/terms/i/ibes.asp>

4 Experiments

4.1 Models

In this section, we provide details of the four categories of models we have used. Initially, we provide detail on the proposed weak-supervision model with the customized aggregation function. In order to provide a comprehensive benchmark for the claim detection task and comparison with proposed weak-supervision model, we add Bi-LSTM, six BERT architecture-based PLMs, and three generative LLMs.

Weak-Supervision Model For implementing a weak-supervision model we use the Snorkel library (Ratner et al., 2017), leveraging its inherent pipeline structure for generating labels for each data segment and then passing the outputs through the customized aggregation function.

Labeling functions used in our model include rule-based pattern matching combined with part-of-speech (POS) tag constraints for some phrases. We create seventeen labeling functions for the categorization of results and also make use of multiple other labeling functions in order to divide sentences representing assertions or written in the past tense. These labeling functions are listed in Table 5. More details on the construction of the labeling function can be found in Appendix B.

Aggregation Function The output of the labeling functions needs to be aggregated to decide the final label of the sentence. Unlike other models, we use independent, weighted labeling functions with weights based on the level of confidence assigned by Subject Matter Experts (SMEs). Our labeling function can produce four distinct types of output: -1 for a high confidence out-of-claim sentence, 0 for abstention from making a claim, 1 for a low confidence in making a claim, and 2 for a high confidence in making a claim. This system

allows us to further differentiate in-claim sentences into two levels of confidence. The pseudo-code in Algorithm 1 illustrates our aggregation function.

Algorithm 1 Aggregation Function

```
if any of the labeling functions' output is -1 then
    label ← "out-of-claim"
else if the max of the labeling functions' output is 2 then
    label ← "in-claim"
else
    label ← majority vote output
end if
```

Traditional majority vote takes decisions based on votes from all the labeling functions, meaning assigning equal weights. The weighted majority vote aggregation function, such as Snorkel, learns the weight for each labeling function from the data itself. In our case, Subject Matter Experts decide that some labeling functions are higher in the hierarchy than others. This means that we look at their labels first before looking at the output of other labeling functions. If those higher-valued labeling functions refrain from voting (by giving an abstain label, value=0), we look at the output of other labeling functions. Otherwise, we take labels based on the majority vote.

To facilitate a comprehensive comparison of our weak-supervision model against various other model categories, we additionally leverage Generative Large Language Models (LLMs) in both zero-shot and few-shot settings, and conduct fine-tuning on Bi-LSTM as well as other Pre-trained Language Models (PLMs). Detailed information regarding the implementation of these models is delineated in the Appendix C.

4.2 Results

In this section, we present the results obtained using the above models and provide a detailed analysis of the outcomes.

Weak-Supervision Model The performance in Table 2, highlights how well our Weak-Supervision based model performs when compared with manually annotated data. In order to make sure that there is no contamination issue between the labeling functions and annotated data, we perform a robustness check in Appendix A. We also perform ablation on the number of labeling functions in Appendix D.

We consider majority voting and Snorkel’s aggregation function (Ratner et al., 2017) as baseline aggregation functions for comparative ablation

analysis. The accuracy of baseline aggregation functions along with our aggregation function is reported in Table 3. For all three models, the same set of labeling functions is used and they only differ in the aggregation part.⁵ The result highlights the importance of the construction of a customized aggregation function for a weak-supervision model where a small set of labeling functions are complete and less noisy.

Generative LLMs There are a few observations regarding the performance of Generative LLMs. First, we see that utilizing a more detailed prompt leads to large improvements in performance across all three models. Secondly, Falcon and Llama have a large increase in performance as well when using six-shot prompting. However, ChatGPT did not have as large of an improvement when utilizing few-shot prompting. While the reasoning behind this is uncertain, it is clear that prompt engineering (particularly creating detailed prompts) can lead to substantial improvement. Zero-shot ChatGPT fails to outperform both weak-supervision and fine-tuned PLMs. It still achieves impressive performance without having access to any labeled data. Of the variations of prompting attempted, Llama with six-shot prompting yielded the best results. This seems to suggest that through the use of prompt engineering, open-source models may be able to close the gap with closed LLMs.

Bi-LSTM The Bi-LSTM model outperforms the weak-supervision model on analyst reports data but doesn’t outperform on earnings call data. The potential reason can be the larger fine-tuning dataset available for analyst reports. It doesn’t outperform the model based on BERT on any of the four configurations.

PLMs The fine-tuned models utilizing the BERT architecture demonstrate superior performance compared to other model classes, emphasizing the significant value gained from annotated data. Intriguingly, the model that achieves the highest performance within a particular train-test dataset category does not necessarily exhibit the best performance on transfer learning datasets. This finding underscores the importance of separate data annotation. Notably, the RoBERTa model emerges as the top performer within the same train-test data category.

⁵We do not perform any post-processing on the output to convert abstain label to one of the labels.

Panel A: Models Without Further Training		
Model	Analyst Reports (AR)	Earnings Calls (EC)
Weak-Supervision	0.9272 (0.0116)	0.9382 (0.0213)
Falcon-7B (0-shot)	0.4167 (0.0075)	0.3884 (0.0624)
Llama-2-70B (0-shot)	0.7278 (0.0079)	0.5407 (0.0267)
ChatGPT-3.5 (0-shot)	0.9191 (0.0144)	0.7569 (0.0023)
Falcon-7B (6-shots)	0.3410 (0.0109)	0.3021 (0.0343)
Llama-2-70B (6-shots)	0.9169 (0.0049)	0.7972 (0.0228)
ChatGPT-3.5 (6-shots)	0.8943 (0.0033)	0.7334 (0.0198)

Panel B: Fine-Tuned Models				
Train/Test	AR/AR	EC/AR	AR/EC	EC/EC
Bi-LSTM	0.9309 (0.0235)	0.8244 (0.0332)	0.8961 (0.0236)	0.8892 (0.0375)
BERT-base-uncased	0.9532 (0.0192)	0.9269 (0.0150)	0.9251 (0.0113)	0.9376 (0.0205)
FinBERT-base	0.9617 (0.0076)	0.9381 (0.0112)	0.9209 (0.0257)	0.9279 (0.0135)
FLANG-BERT-base	0.9611 (0.0137)	0.9270 (0.0109)	0.9119 (0.0257)	0.9363 (0.0089)
RoBERTa-base	0.9615 (0.0091)	0.9319 (0.0131)	0.8906 (0.0301)	0.9563 (0.0036)
BERT-large-uncased	0.9539 (0.0111)	0.9183 (0.0063)	0.9197 (0.0349)	0.9416 (0.0349)
RoBERTa-large	0.9642 (0.0069)	0.9381 (0.0138)	0.8975 (0.0244)	0.9427 (0.0153)

Table 2: In the table, A/B indicates that the model is fine-tuned on dataset A and tested on dataset B. All values are F1 scores. An average of 3 seeds was used for all models. The standard deviation of F1 scores is in parentheses.

Aggr. Function	AR	EC
Majority Vote	0.4274 (0.0208)	0.5313 (0.0427)
Snorkel’s WMV	0.4269 (0.0204)	0.5309 (0.0372)
Ours	0.9272 (0.0116)	0.9382 (0.0213)

Table 3: Performance comparison of our aggregation function with baseline aggregation functions. All values are F1 scores. An average of 3 seeds was used for all models. The standard deviation of F1 scores is reported in parentheses.

Latency and Financial Applicability In finance, latency is crucial as investors aim to surpass competitors. Figure 2 shows just how stark the differences is in latency. Our weak-supervision (WS) model stands out for its low latency, offering significant advantages in the fast-moving financial markets. Despite challenges in measuring latency for API-based, closed-source models like ChatGPT, our analysis on Falcon-7B and Llama-70B highlights the WS model’s superior speed and efficiency. This model’s performance is key in finance, where processing speed can be decisive in transaction success. Furthermore, even if generative LLMs do overcome the hurdle of latency, large ethical challenges in finance as identified by (Khan and Umer,




Figure 2: This bar chart compares the latency (log scale) of various models relative to the weak-supervision model.

2024) still persist. We also discuss carbon emission comparison of models in Appendix E.

5 Market Analysis

5.1 Experiment Setup

Construction of the Optimism Measure We use our weak-supervision model to label all the financial numeric sentences in the analyst reports

and earnings calls as in-claim or out-of-claim. We then filter the sentences and only keep in-claim sentences to evaluate predictions.

We further label each in-claim sentence as ‘positive’, ‘negative’, or ‘neutral’ using the [fine-tuned](#) sentiment analysis model specifically for the financial domain. The model is fine-tuned for financial sentiment analysis using the pre-trained FinBERT ([Araci, 2019](#)). We then use labeled sentences in each document to generate a document-level measure of analyst optimism for document i using the following formula:

$$\text{Optimism}_i = 100 \times \frac{\text{Pos. In-claim}_i - \text{Neg. In-claim}_i}{\text{Total Sentences}_i} \quad (1)$$

where Pos. In-claim_i and Neg. In-claim_i are the number of positive and negative in-claim sentences respectively in document i after the filter, and Total Sentences_i is the total number of sentences in the document.

Empirical Specification We use the following empirical specification for market analysis.

$$Y_{i,t} = \alpha + \beta \times \text{Optimism}_{i,t} + \epsilon_{i,t} \quad (2)$$

Here $Y_{i,t}$ is the outcome variable of interest for firm i at time t , α is a constant term, and $\epsilon_{i,t}$ is an error term. The coefficient (β) will help us understand the influence of $\text{Optimism}_{i,t}$ on the outcome variable ($Y_{i,t}$).

Outcome (Y)	Constant (α)	Beta (β)
Earn. Surp.	0.1744 ***	-1.9883 ***
CAR [+2, +30]	0.9548 ***	-34.5749 ***
CAR [+2, +60]	0.8559 **	-54.335 ***

Table 4: Market analysis result based on the empirical regression. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

5.2 Post Earnings Prediction

We examine the relation between optimism in analyst reports for a company in a specific quarter and its effect on earnings. Using earnings-based metrics, we perform a regression as per Eq 2 using earnings call transcripts and analyst report data. For quarters with multiple reports on one stock, we aggregate sentences and claims to compute Optimism_i .

Earnings Surprise (%) The Earning Surprise (%) is calculated by subtracting the median EPS (in the last 90 days) from the actual EPS. The difference is scaled by the stock price at the end of the quarter and multiplied by 100. This method aligns with [Chava et al. \(2022\)](#).

The Earnings Surprise (%) is set as the outcome variable ($Y_{i,t}$). The results in Table 4 show a significant link between optimism and the Earnings Surprise (%). A negative β coefficient indicates that with every unit rise in optimism in analyst reports, the Earnings Surprise (%) drops. This implies that heightened optimism in reports often leads to the actual EPS underperforming expectations. This "false optimism" aligns with previous studies like ([Coleman et al., 2021](#)), highlighting analysts’ tendency to overestimate firm performance.

Cumulative Abnormal Returns We further aim to explore the influence of optimism in analyst reports on the magnitude of cumulative abnormal return (CAR) post-earnings. CAR for a firm represents the total daily abnormal stock return in the period after a specific event, in our context, the firm’s earnings conference call.

We analyze two CAR time frames. CAR[+2, +30] is the cumulative abnormal for the [+2,+30] trading day window post-earnings call, as determined by [Chava et al. \(2022\)](#). The same methodology is used to calculate CAR[+2, +60] as well.

Table 4 shows that greater optimism in analyst reports corresponds with a larger decline in CAR. This emphasizes the ‘false optimism’ trend in reports, where increased optimism leads to greater discrepancies from actual outcomes, leading to a larger negative cumulative abnormal return.

The prevailing notion in finance literature is that analysts are overly optimistic. While [Francis and Philbrick \(1993\)](#) and [Barber et al. \(2007\)](#) believe this bias helps maintain good ties with corporate insiders, [Michaely and Womack \(1999\)](#) sees it as a means for personal financial gains. Recently, [Brown et al. \(2022\)](#) found that analysts favor firms with attributes like high debt or fluctuating earnings. This suggests such firms might exaggerate earnings, potentially through manipulation. Our market analysis aligning with these theories reinforces our method’s accuracy and the financial relevance of our study. Furthermore, [Bhojraj et al. \(2009\)](#) shows that simply exceeding or failing to meet analyst expectations under certain conditions can lead to unique post-earnings characteristics for




Figure 3: Normalized Confusion Matrix illustrating the percentage of trades categorized by negative or positive adjusted optimism and their corresponding CAR[+2,+60] outcomes. Each cell represents the percentage of total trades that fall within each category.

a company.

5.3 Predictive Power of Optimism

To highlight a usage of Optimism for making trading predictions, we employ a simple “trading strategy”. We utilize analyst reports from 2017-2019 as a training set to identify the average positive bias in the "optimism" measure. To adjust for the bias in our test set, the 2020 analyst reports, we subtract the mean bias from the optimism score for each company, correcting for the inherent positive bias. The division of the dataset into training and testing phases is crucial to avoid look-ahead bias in calculating mean optimism. After adjusting the optimism measure in the test dataset, we implement a straightforward investment strategy: short selling companies with a positive adjusted optimism score and buying shares of companies with a negative adjusted optimism score. This approach is based on the rationale of investing in companies with **overly** pessimistic sentiment and divesting from those with **overly** optimistic sentiment. We use Earnings Surprise, CAR[+2, +30], and CAR[+2,+60] to determine the success or failure of our hypothetical trades.

The confusion matrix corresponding to the results of CAR[+2,+60] are visualized in Figure 3, while Earnings Surprise and CAR[+2,+30] are shown in Appendix G. The confusion matrix shows that such a rule-based strategy achieves an approximate 81% accuracy in correctly predicting the direction of stock movement. Additionally, the high accuracy lasting up to 60 days indicates that using

optimism can effectively predict stock movements for more than just a few days, demonstrating a valuable preliminary application of such identification for the financial field.

6 Conclusion

Our work presents claim based labeled dataset in the English language alongside presenting a weak-supervision model with an accuracy of 93%. Developed customized aggregation function outperforms baseline aggregation functions. We benchmark various language models and compare the performance with the weak-supervision model. We show the application of claim detection by generating a measure of optimism from the weak-supervision model. We also validate the measure by studying its applicability in predicting earnings surprise, abnormal returns, and earnings optimism. We release our models, code, and benchmark data (for earnings call transcripts only) on Hugging Face and GitHub. We also note that the trained model for claim detection can be used on other financial texts.

Limitations

By acknowledging the following limitations, we pave the way for future research to address these areas and further enhance the understanding and applicability of our approach.

- *Limited Scope of Text Data:* Our analysis is restricted to analyst reports and earnings calls, excluding other potentially valuable text datasets such as related news articles and investor presentations. Incorporating these additional sources of information could provide a more comprehensive understanding of pre-earnings drifts.
- *Exclusion of Audio and Video Features:* Our measure construction does not utilize audio or video features from earnings calls, which may contain supplementary information.
- *Omission of Alternative Weak-Supervision Models:* We do not explore multiple end models, such as the confidence-based sampling with contrastive loss proposed in the COSINE framework by Yu et al. (2020). Incorporating such alternative weak-supervision models could offer additional insights and improve the robustness of our approach.

Ethics Statement

Our work adheres to ethical considerations, although we acknowledge certain biases and limitations in our study. We do not identify any potential risks stemming from our research; however, we recognize the presence of geographic and gender biases in our analysis.

- *Geographic Bias*: Our study focuses solely on publicly listed companies in the United States of America, which introduces a geographic bias. The findings may not be fully representative of global firms and markets.
- *Gender Bias*: We acknowledge the gender bias present in our study due to the predominant representation of male analysts, CEOs, and CFOs.
- *Data Ethics*: The data used in our study, derived from publicly available sources, does not raise ethical concerns. All raw data is obtained from public companies that are obligated to disclose information under the guidance of the SEC and are subject to public scrutiny.
- *Language Model Ethics*: The language models employed (with proper citation) in our research are publicly available and fall under license categories that permit their use for our intended purposes. While most models employed are publicly available, it is important to note that ChatGPT's prompt answers will not be made public due to licensing conditions. We acknowledge the environmental impact of large pre-training of language models and mitigate this by limiting our work to fine-tuning existing models.
- *Annotation Ethics*: All annotations were performed by the authors, ensuring that no additional ethical concerns arise from the annotation process.
- *Hyperparameter Reporting*: In the interest of clarity and readability, we refrain from reporting the best hyperparameters found through grid search in the main paper. Instead, we will make all grid search results, including hyperparameter information, publicly available on GitHub. This transparency allows interested readers to access detailed information on our experimental setup.

- *Publicly Available Data*: We specify the datasets that will be made publicly available and indicate the applicable licenses under which they will be shared.

By acknowledging these ethical considerations and limitations, we strive to maintain transparency and promote responsible research practices.

References

- Md Shad Akhtar, Abhishek Kumar, Deepanway Ghosal, Asif Ekbal, and Pushpak Bhattacharyya. 2017. A multilayer perceptron based ensemble technique for fine-grained financial sentiment analysis. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, pages 540–546.
- Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Merouane Debbah, Etienne Goffinet, Daniel Hestlow, Julien Launay, Quentin Malartic, Badreddine Noune, Baptiste Pannier, and Guilherme Penedo. 2023. Falcon-40B: an open large language model with state-of-the-art performance.
- Dogu Araci. 2019. Finbert: Financial sentiment analysis with pre-trained language models. *arXiv preprint arXiv:1908.10063*.
- Brad M Barber, Reuven Lehavy, and Brett Trueman. 2007. Comparing the stock recommendation performance of investment banks and independent research firms. *Journal of financial economics*, 85(2):490–517.
- Sanjeev Bhojraj, Paul Hribar, Marc Picconi, and John McInnis. 2009. Making sense of cents: An examination of firms that marginally miss or beat analyst forecasts. In *The Journal of Finance*, volume 64 (5), pages 2361–2388.
- Anna Bergman Brown, Guoyu Lin, and Aner Zhou. 2022. Analysts' forecast optimism: The effects of managers' incentives on analysts' forecasts. *Journal of Behavioral and Experimental Finance*, 35:100708.
- Sean Cao, Wei Jiang, Baozhong Yang, and Alan L Zhang. 2020. How to talk when a machine is listening: Corporate disclosure in the age of ai. Technical report, National Bureau of Economic Research.
- Olivier Chapelle, Bernhard Scholkopf, and Alexander Zien. 2009. Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3):542–542.
- Sudheer Chava, Wendi Du, and Baridhi Malakar. 2021. Do managers walk the talk on environmental and social issues? Available at SSRN 3900814.
- Sudheer Chava, Wendi Du, Agam Shah, and Linghang Zeng. 2022. Measuring firm-level inflation exposure: A deep learning approach. Available at SSRN 4228332.

- Chung-Chi Chen, Hen-Hsen Huang, and Hsin-Hsi Chen. 2019. Numeral attachment with auxiliary tasks. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1161–1164.
- Chung-Chi Chen, Hen-Hsen Huang, and Hsin-Hsi Chen. 2020. Numclaim: Investor’s fine-grained claim detection. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 1973–1976.
- Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, et al. 2021. Finqa: A dataset of numerical reasoning over financial data. *arXiv preprint arXiv:2109.00122*.
- Braiden Coleman, Kenneth J Merkle, and Joseph Pacelli. 2021. Human versus machine: A comparison of robo-analyst and traditional research analyst investment recommendations. *The Accounting Review, Forthcoming*.
- Shane A Corwin, Stephannie A Larocque, and Mike A Stegemoller. 2017. Investment banking relationships and analyst affiliation bias: The impact of the global settlement on sanctioned and non-sanctioned banks. *Journal of Financial Economics*, 124(3):614–631.
- Min-Yuh Day and Chia-Chou Lee. 2016. Deep learning for financial sentiment analysis on finance news providers. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 1127–1134. IEEE.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Jennifer Francis and Donna Philbrick. 1993. Analysts’ decisions as products of a multi-task environment. *Journal of Accounting Research*, 31(2):216–230.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Allen Hu and Song Ma. 2021. Persuading investors: A video-based study. Technical report, National Bureau of Economic Research.
- Narasimhan Jegadeesh and Woojin Kim. 2010. Do analysts herd? an analysis of recommendations and market reactions. *The Review of Financial Studies*, 23(2):901–937.
- David Kartchner, Wendi Ren, David Nakajima An, Chao Zhang, and Cassie S Mitchell. 2020. Regal: Rule-generative active learning for model-in-the-loop weak supervision. *Advances in neural information processing systems*.
- Muhammad Salar Khan and Hamza Umer. 2024. [Chatgpt in finance: Applications, challenges, and solutions](#). *Heliyon*, 10(2):e24890.
- Loïc Lanelongue, Jason Grealey, and Michael Inouye. 2021. Green algorithms: Quantifying the carbon footprint of computation. *Advanced Science*, 8(12):2100707.
- Jiazheng Li, Linyi Yang, Barry Smyth, and Ruihai Dong. 2020. Maec: A multimodal aligned earnings conference call dataset for financial risk prediction. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*, pages 3063–3070.
- Kai Li, Feng Mai, Rui Shen, and Xinyan Yan. 2021. Measuring corporate culture using machine learning. *The Review of Financial Studies*, 34(7):3265–3315.
- Pierre Lison, Jeremy Barnes, and Aliaksandr Hubin. 2021. skweak: Weak supervision made easy for nlp. *arXiv preprint arXiv:2104.09683*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, M. Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *ArXiv*, abs/1907.11692.
- Zhuang Liu, Degen Huang, Kaiyu Huang, Zhuang Li, and Jun Zhao. 2020. Finbert: A pre-trained financial language representation model for financial text mining. In *IJCAI*, pages 4513–4519.
- Macedo Maia, Siegfried Handschuh, André Freitas, Brian Davis, Ross McDermott, Manel Zarrouk, and Alexandra Balahur. 2018. [Www’18 open challenge: Financial opinion mining and question answering](#). In *Companion Proceedings of the The Web Conference 2018, WWW ’18*, page 1941–1942, Republic and Canton of Geneva, CHE. International World Wide Web Conferences Steering Committee.
- Pekka Malo, Ankur Sinha, Pekka Korhonen, Jyrki Wallenius, and Pyry Takala. 2014. Good debt or bad debt: Detecting semantic orientations in economic texts. *Journal of the Association for Information Science and Technology*, 65(4):782–796.
- R David McLean, Jeffrey Pontiff, and Christopher Reilly. 2020. Retail investors and analysts. Technical report, Working Paper, Boston College.
- Yu Meng, Jiaming Shen, Chao Zhang, and Jiawei Han. 2018. Weakly-supervised neural text classification. In *proceedings of the 27th ACM International Conference on information and knowledge management*, pages 983–992.
- Roni Michaely and Kent L Womack. 1999. Conflict of interest and the credibility of underwriter analyst recommendations. *The Review of Financial Studies*, 12(4):653–686.
- Mike Mintz, Steven Bills, Rion Snow, and Dan Jurafsky. 2009. Distant supervision for relation extraction without labeled data. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural*

- Language Processing of the AFNLP*, pages 1003–1011.
- Cuong V Nguyen, Sanjiv R Das, John He, Shenghua Yue, Vinay Hanumaiah, Xavier Ragot, and Li Zhang. 2021. Multimodal machine learning for credit modeling. In *2021 IEEE 45th Annual Computers, Software, and Applications Conference (COMPSAC)*, pages 1754–1759. IEEE.
- Sinno Jialin Pan and Qiang Yang. 2010. [A survey on transfer learning](#). *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. 2019. [Pytorch: An imperative style, high-performance deep learning library](#). In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems* 32, pages 8024–8035. Curran Associates, Inc.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *EMNLP*, volume 14, pages 1532–1543.
- Alexander Ratner, Stephen H Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. 2017. Snorkel: Rapid training data creation with weak supervision. In *Proceedings of the VLDB Endowment. International Conference on Very Large Data Bases*, volume 11 (3), page 269. NIH Public Access.
- Alexander Ratner, Stephen H Bach, Henry Ehrenberg, Jason Fries, Sen Wu, and Christopher Ré. 2020. Snorkel: Rapid training data creation with weak supervision. *The VLDB Journal*, 29(2):709–730.
- Anna Rogers, Niranjan Balasubramanian, Leon Derczynski, Jesse Dodge, Alexander Koller, Sasha Luccioni, Maarten Sap, Roy Schwartz, Noah A. Smith, and Emma Strubell. 2023. [Closed ai models make bad baselines](#).
- Ramit Sawhney, Piyush Khanna, Arshiya Aggarwal, Taru Jain, Puneet Mathur, and Rajiv Shah. 2020. Voltage: volatility forecasting via text-audio fusion with graph convolution networks for earnings calls. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8001–8013.
- Mike Schuster and Kuldip K Paliwal. 1997. Bidirectional recurrent neural networks. *IEEE transactions on Signal Processing*, 45(11):2673–2681.
- Burr Settles. 2009. Active learning literature survey. *University of Wisconsin-Madison Department of Computer Sciences*.
- Agam Shah, Suvan Paturi, and Sudheer Chava. 2023a. Trillion dollar words: A new financial dataset, task & market analysis. *arXiv preprint arXiv:2305.07972*.
- Agam Shah, Ruchit Vithani, Abhinav Gullapalli, and Sudheer Chava. 2023b. Finer: Financial named entity recognition dataset and weak-supervision model. *arXiv preprint arXiv:2302.11157*.
- Raj Shah, Kunal Chawla, Dheeraj Eidnani, Agam Shah, Wendi Du, Sudheer Chava, Natraj Raman, Charese Smiley, Jiaao Chen, and Diyi Yang. 2022. [When FLUE meets FLANG: Benchmarks and large pre-trained language model for financial domain](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 2322–2335, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.
- Paroma Varma and Christopher Ré. 2018. Snuba: Automating weak supervision to label training data. In *Proceedings of the VLDB Endowment. International Conference on Very Large Data Bases*, volume 12 (3), page 223. NIH Public Access.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Yi Yang, Mark Christopher Siy Uy, and Allen Huang. 2020. [Finbert: A pretrained language model for financial communications](#). *CoRR*, abs/2006.08097.
- Yue Yu, Lingkai Kong, Jieyu Zhang, Rongzhi Zhang, and Chao Zhang. 2022. Actune: Uncertainty-based active self-training for active fine-tuning of pretrained language models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1422–1436.
- Yue Yu, Simiao Zuo, Haoming Jiang, Wendi Ren, Tuo Zhao, and Chao Zhang. 2020. Fine-tuning pretrained language model with weak supervision: A

contrastive-regularized self-training approach. *arXiv preprint arXiv:2010.07835*.

Man-Ching Yuen, Irwin King, and Kwong-Sak Leung. 2011. A survey of crowdsourcing systems. In *2011 IEEE third international conference on privacy, security, risk and trust and 2011 IEEE third international conference on social computing*, pages 766–773. IEEE.

Jieyu Zhang, Yue Yu, Yinghao Li, Yujing Wang, Yaming Yang, Mao Yang, and Alexander Ratner. 2021. Wrench: A comprehensive benchmark for weak supervision. *arXiv preprint arXiv:2109.11377*.

Rongzhi Zhang, Yue Yu, Pranav Shetty, Le Song, and Chao Zhang. 2022. Prboost: Prompt-based rule discovery and boosting for interactive weakly-supervised learning. *arXiv preprint arXiv:2203.09735*.

A Robustness Check

From a data engineering perspective, there can be concern about the model design and gold data construction as the authors who designed the weak-supervision model have annotated the data. This can lead to exaggerated performance on the data, which may taint the test set. To ensure that there is no contamination issue in the weak-supervision model and it is generalizable, we get the same test dataset annotated separately by four annotators with master’s degrees in Quantitative Finance. These annotators were hired by the department as Graduate Assistants based on merit and were paid a \$20 per hour salary for their work which is more than double the federal minimum wage and higher than the highest minimum wage (\$15.74 in Washington, D.C) in the USA. The rates are standard and in compliance with ethical standards. These annotators had no information about the rules/patterns used in our weak-supervision model. Each sample in the test dataset is annotated by two annotators, and we drop the observations where there is a disagreement among annotators.⁶ The F1 score of the weak-supervision model on a dataset annotated by non-authors is 0.9281 which is close to a score of 0.9272 on the author-annotated dataset. We also recalculate the F1 score of the model based on the author-annotated labels after dropping observations dropped in a non-author annotated dataset. The model gives a higher mean F1 score of 0.9360 which is expected as ambiguous sentences are dropped. Overall these results show the robustness of our model on the dataset annotated

⁶There is 98.59% agreement between two annotators.

by annotators who don’t have knowledge of the rules used in the weak-supervision model. From here onwards, the performance is always calculated on a gold dataset created by authors.

B Labeling Functions Methodology

The following illustrates the methodology adopted by us while choosing the rules to define the weak-supervision mode. All rules were acknowledged post detailed analysis of sample documents distributed over sector and time :

1. Certain phrases such as "reasons to buy", "reasons to sell" or the presence of words which are indicative of past tense such as "was", "were" are characteristic of out-of-claim sentences, since they indicated either facts or events which happened in the past. Examples are given in the set 1 of Table 5.
2. Phrases often provided definitive information about a given sentence in a document and in most cases they had a fairly consistent linguistic composition. Examples are given in the set 2 of Table 5.
3. In a bid to capture the effect of a few other verb forms indicative of a probabilistic event, we also chose to look at its lemmatized form to reduce inflectional usage and use the base token for a more holistic evaluation over multiple usage formats. Examples are given in the set 3 of Table 5.
4. POS tags were also derived for "project" as a word wherever present. This was done to segregate its usage as a verb. Its usage as a verb was usually observed to be adopted while making claims or predictions. Examples are given in the set 4 of Table 5.
5. The alternate adoption of phrase matching was to identify in-claim sentences. This mostly consisted of a verb form indicative of a probabilistic event (eg: likely, intends) coupled with a preposition (usually "to" or "at"). Based on the ambiguity of the resulting phrase they were either categorised as a high-confidence claim or a low-confidence one. Examples are given in the set 5 of Table 5.

Set	Used to detect	Output	Type	Keyword or phrase
1	High Confidence out-of-claim (Past Tense or Assertions)	-1/0	Phrase Matching	reasons to buy:, reasons to sell:, was, were, declares quarterly dividend, last earnings report, recorded
2	Low Confidence in-claim	1/0	Phrase Matching	earnings guidance to, touted to, entitle to
3	High Confidence in-claim	2/0	Lemmatized Word matching	expect, anticipate, predict, forecast, envision, contemplate
4	High Confidence in-claim	2/0	POS Tag for word "project"	VBN, VB, VBD, VBG, VBP, VBZ
5	High Confidence in-claim	2/0	Phrase Matching	to be, likely to, on track to, intends to, aims to, to incur, pegged at

Table 5: Labeling Functions used in weak-supervision model. SpaCy Lemmatizer has been used for labeling functions involving lemmatized word matching.

C Additional Models

C.1 Generative LLMs

To understand the capabilities of current state-of-the-art (SOTA) generative LLMs’ in a zero-shot and few-shot manner, we add ChatGPT⁷ performance benchmark in our study. We use the "gpt-3.5-turbo-0613" model with 200 max tokens for output, and a 0.0 temperature value. The ChatGPT API was accessed on Feb 2nd, 2024. In a recent article, Rogers et al. (2023) made a case for why closed models like ChatGPT make bad baselines. In order to understand where SOTA open-source LLMs stand in comparison to ChatGPT and fine-tuned models, we also test the Falcon-7B-Instruct (Almazrouei et al., 2023) and "Llama-2-70B-chat" (Touvron et al., 2023) models. The prompt templates are provided in Table 6. All our prompting was done in consistency with reputable resources, such as the "Prompt Engineering Guide"⁸. We also test the model with zero-shot and six-shot. The six-shot prompting consists of 3 ‘in-claim’ examples and 3 ‘out-of-claim’ examples.

C.2 Bi-LSTM

In the realm of text classification problems, Long Short-Term Memory (LSTM) was a popular recurrent neural network architecture (Hochreiter and Schmidhuber, 1997). An enhanced approach to LSTM is the Bidirectional LSTM (Bi-LSTM), which processes input in both directions (Schuster and Paliwal, 1997). In order to assess the efficacy of Recurrent Neural Networks (RNNs) in claim detection, we employ the Bi-LSTM model on the datasets we have developed. Instead of training it from scratch, we initialize the embedding layer of

the Bi-LSTM using 300-dimensional GloVe embeddings trained using Common Crawl (Pennington et al., 2014). Here we perform the task of sequence classification while minimizing the cross-entropy loss. We employ a grid search approach to identify the optimal hyperparameters for each model, considering four different learning rates (1e-4, 1e-5, 1e-6, 1e-7) and four different batch sizes (32, 16, 8, 4). In our training process, we employ a maximum of 100 epochs, incorporating early stopping criteria. In cases where the validation F1 score does not exhibit an improvement of greater than or equal to 1e-2 over the subsequent 7 epochs, we designate the previously saved best model as the final fine-tuned model.

C.3 PLMs

In order to establish a performance benchmark, our study encompasses a range of transformer-based (Vaswani et al., 2017) models of varying sizes. For the small models, we employ BERT (Devlin et al., 2018), FinBERT (Yang et al., 2020), FLANG-BERT (Shah et al., 2022), and RoBERTa (Liu et al., 2019). Within the category of large models, we incorporate BERT-large (Devlin et al., 2018) and RoBERTa-large (Liu et al., 2019). To avoid over-fitting on financial text, we refrain from conducting any pre-training on these models prior to fine-tuning. Here we perform the task of sequence classification while minimizing the cross-entropy loss. For PLMs, we employ grid-search, fine-tuning, and early stopping similar to what we used for Bi-LSTM. The experiments are conducted using PyTorch (Paszke et al., 2019) on an NVIDIA RTX A6000 GPU. Each model is initialized with the pre-trained version from the Transformers library provided by Huggingface (Wolf et al., 2020).

⁷<https://chat.openai.com/>

⁸<https://www.promptingguide.ai/>

Prompt Name	Description
Zero-shot	Discard all the previous instructions. Behave like you are an expert sentence classifier. Classify the following sentence into either 'INCLAIM' or 'OUTOFCLAIM'. 'INCLAIM' refers to predictions or expectations about financial outcomes. 'OUTOFCLAIM' refers to sentences that provide numerical information or established facts about past financial events. For each classification, 'INCLAIM' can be thought of as 'financial forecasts', and 'OUTOFCLAIM' as 'established financials'. Now, for the following sentence provide the label in the first line and provide a short explanation in the second line. The sentence: {sentence}
Few-shot	Discard all the previous instructions. Behave like you are an expert sentence classifier. Classify the following sentence into either 'INCLAIM' or 'OUTOFCLAIM'. 'INCLAIM' refers to predictions or expectations about financial outcomes. 'OUTOFCLAIM' refers to sentences that provide numerical information or established facts about past financial events. For each classification, 'INCLAIM' can be thought of as 'financial forecasts', and 'OUTOFCLAIM' as 'established financials'. Here are a few examples: Example 1: free cash flow of \$2.3 billion was up 10.5%, benefiting from the positive year-over-year change in net working capital due to covid at both nbcu and sky, half of which resulted from the timing of when sports rights payments were made versus when sports actually aired and half of which resulted from a slower ramp in content production. // The sentence is OUTOFCLAIM Example 2: we've also used our scale of more than 15,000 combined stores to drive merchandise cost savings exceeding \$70 million. // The sentence is OUTOFCLAIM Example 3: consolidated total capital was \$2.9 billion for the quarter. // The sentence is OUTOFCLAIM Example 4: third, as a result of the continued strength of the u.s. dollar, we are now factoring in an incremental fx headwind of \$175 million across q3 and q4 revenue. // The sentence is INCLAIM Example 5: though early, we are planning our business based on the expectation of cy '23 wfe declining approximately 20% based on increasing global macroeconomic concerns and recent public statements from several customers, particularly in memory, and the impact of the new u.s. government regulations on native china investment. // The sentence is INCLAIM Example 6: we expect revenue growth to be in the range of 5.5% to 6.5% year on year. // The sentence is INCLAIM Now, for the following sentence provide the label in the first line and provide a short explanation in the second line. The sentence: {sentence}

Table 6: Prompts used for zero-shot and few-shot inference.

D Ablation: Number of Labeling Functions

Figure 4, shows how the accuracy of the model changes depending on the number of labeling functions. For this plot, we initially computed the con-

tribution of each labeling function (Table 5, High confidence and Low Confidence in-claim) towards the detection of in-claim sentences and then considered the addition of new labeling function at each step to ensure the steepest ascent to saturation. At

each step, in addition to one new labeling function, all labeling functions present in Table 5 for Past Tense and Assertions, were also used. They either abstain or classify sentences as out-of-claim and help improve the classification of out-of-claim sentences. From the plot, we can notice that after around thirteen labeling functions, the addition of new labeling functions does not produce any change in the accuracy. In fact, increasing labeling functions thereafter leads to a minor decrease in accuracy. This suggests that we can effectively capture the required trends for classification in this setting with thirteen labeling functions.




Figure 4: Accuracy v/s Number of labelling functions. Note: This is accuracy, not F1 score.

E Environmental Impact

Our investigation extends beyond just performance metrics, embracing a conscientious approach towards the environmental implications of AI usage. To ensure a standardized and rigorous assessment of CO₂e, we drew upon the methodology outlined by Lannelongue et al. (2021) and utilized the Green Algorithms calculator⁹. The value of CO₂e are reported in Figure 5. This dual focus on minimizing latency and CO₂e without compromising performance highlights our commitment to advancing sustainable and efficient AI technologies in sectors where both are of paramount importance, such as finance. The CO₂ emissions (CO₂e) associated with the inference phase of these models are particularly telling, with our WS model not only leading in latency but also in sustainability, registering the lowest CO₂e among all models reviewed. This underscores the viability of employing AI in environments where both speed and environmental responsibility are valued. In contrast, models such

⁹<https://calculator.green-algorithms.org/>

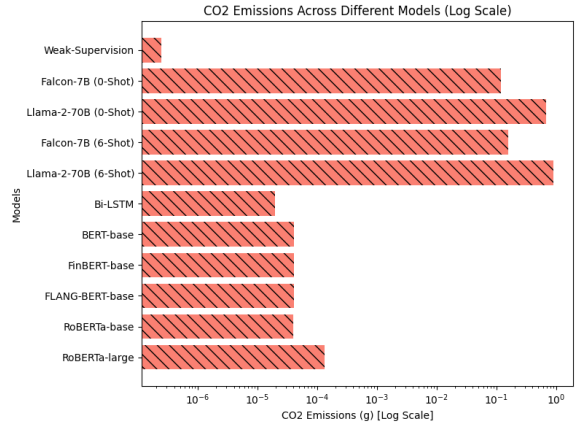


Figure 5: This bar chart compares the CO₂ emissions (log scale) of various models relative to the weak-supervision model.

as Llama-70B, despite their performance coming close to our model, incur significantly higher (more than a million times larger) CO₂e due to their reliance on extensive GPU resources.

F Ablation Study: Market Analysis

To understand the influence of “in-claim” sentences on market sentiment, we introduce the optimism measure in section 5, outlining its implications. In this section, we carry out an ablation study to better understand the impact of “in-claim” sentences. As such, we compute the optimism score for four sentence subsets: Unfiltered, Numerical, Numerical Financial, and Numerical Financial “In-claim” sentences for each file. For example, the optimism score for a subset of Numerical sentences for document i is given by:

$$\text{Optimism (Numerical)}_i = 100 \times \frac{\text{Pos. Numerical}_i - \text{Neg. Numerical}_i}{\text{Total Sentences}_i}$$

We standard normalize these scores for uniform comparison by deducting their mean and dividing by the standard deviation. As the beta coefficient lacks full context, to factor in the size of the sentence subset, we adjusted each coefficient by the average sentence count, terming it as the adjusted beta. This illustrates the information density in each filtered sentence set. When examining the Earnings Surprise (%) columns of Table 7 the Adjusted Beta for Earnings Surprise increases, implying that a mere average of 3.7 “in-claim” sentences holds crucial information. This highlights the high information density of our filtered sentences. While

Sentence Type/Subset	Average Sentences	ES (%)	CAR [+2,+30]	CAR [+2, +60]
		Adj. β	Adj. β	Adj. β
<i>Unfiltered</i>	98	-0.054***	-0.02**	-.03***
<i>Numeric</i>	26	-0.28***	-.06***	-.09***
<i>Numeric Financial</i>	21.6	-0.29***	-.07***	-.11***
<i>Numeric Financial In-claim</i>	3.7	-1.51***	-.26***	-.41***

Table 7: Ablation on market analysis, highlighting the importance and information density of “in-claim” sentences. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

we aren’t dismissing the importance of other sentences, our analysis reveals that the ones we’ve extracted are the most informative on a per-sentence basis.

G Predictive Power of Optimism (Earnings Surprise and CAR[+2,+30])



Figure 6: Percentage of trades categorized by negative or positive adjusted optimism and their corresponding Earnings Surprise outcomes.

Figure 6 and 7 show the results of making trades based on a positive or negative adjusted optimism in terms of the respective performance of the company.




Figure 7: Percentage of trades categorized by negative or positive adjusted optimism and their corresponding CAR[+2,+30] outcomes.