

DEFINE: Decision-Making with Analogical Reasoning over Factor Profiles

Yebowen Hu,[†] Xiaoyang Wang,[‡] Wenlin Yao,[‡] Yiming Lu,[§] Daoan Zhang^{*}
Hassan Foroosh,[†] Dong Yu,[‡] Fei Liu[§]

[†]University of Central Florida [‡]Tencent AI Lab, Seattle

^{*}University of Rochester [§]Emory University

{yebowen.hu, hassan.foroosh}@ucf.edu daoan.zhang@rochester.edu
{shawnxywang, wenlinyao, dyu}@global.tencent.com fei.liu@emory.edu

Abstract

LLMs are ideal for decision-making thanks to their ability to reason over long contexts. However, challenges arise when processing speech transcripts that describe complex scenarios, as they are verbose and include repetition, hedging, and vagueness. E.g., during a company’s earnings call, an executive might project a positive revenue outlook to reassure investors, despite uncertainty regarding future earnings. It is crucial for LLMs to incorporate this uncertainty systematically when making decisions. In this paper, we introduce DEFINE, a modular framework that constructs probabilistic factor profiles from complex scenarios. It then integrates these profiles with analogical reasoning, leveraging insights from similar past experiences to guide LLMs in making critical decisions in new situations. Our framework separates the tasks of quantifying uncertainty and incorporating it into LLM decision-making. This approach is particularly useful in areas such as consulting and financial deliberation, where making decisions under uncertainty is vital.

1 Introduction

Large language models are increasingly relied on for decision-making, thanks to their powerful reasoning capabilities (OpenAI et al., 2024; Anthropic, 2025; Kavukcuoglu, 2025). While research has advanced rapidly in areas such as math, coding, and logical reasoning (Bostrom et al., 2022; Huang and Chang, 2023; Sprague et al., 2024; Mondorf and Plank, 2024; Li et al., 2025; Ren et al., 2025), there is growing interest in exploring how LLMs reason through complex, real-world environments to make high-stakes decisions, such as financial investments (Peters, 2024).

The challenges are compounded when processing long contexts (Krishna et al., 2023; Laban et al., 2024). Core issues may include recency bias, hallucinations, numerical inconsistencies, and more (Liu et al., 2023; Hu et al., 2024a,b; Gao et al., 2024).

While these models are designed to produce reasoning traces, their explanations can remain vague and even unfaithful (Chen et al., 2025). Critically, *LLMs still lack precise, quantitative insight into the key factors that lead to their final decisions, as well as mechanisms for incorporating uncertainty into their decision-making.*

In this paper, we present DEFINE, a new framework for constructing probabilistic factor profiles from speech transcripts that describe complex scenarios, and leveraging these profiles to enhance decision-making. For example, during an earnings call, a company executive might project strong revenue growth to boost investor confidence, despite substantial uncertainty surrounding these projections (Mukherjee et al., 2022). As illustrated in Figure 1, DEFINE generates a structured factor profile from each transcript, capturing not only what is explicitly stated, but also the implications of what is left unsaid. It then uses the BT model (Bradley and Terry, 1952) to identify dominant factors and evaluate how these factors collectively impact decision-making.

Our research integrates probabilistic factor profiles with analogical reasoning, a type of reasoning that identifies connections between similar situations to facilitate knowledge transfer from a familiar context to a new situation (Webb et al., 2023; Yasunaga et al., 2024). Instead of relying on text matching, we use factor profiles to retrieve analogous examples, which are historical cases with similar levels of uncertainty across key dimensions. Analogical reasoning further sets our work apart from traditional Bayesian inference models used in decision-making (Halawi et al., 2024; Lin et al., 2024; Liu et al., 2024), which usually lack direct connections to historical cases. Our key contributions are as follows.

- We introduce DEFINE, a modular framework that enhances LLM decision-making in complex sce-

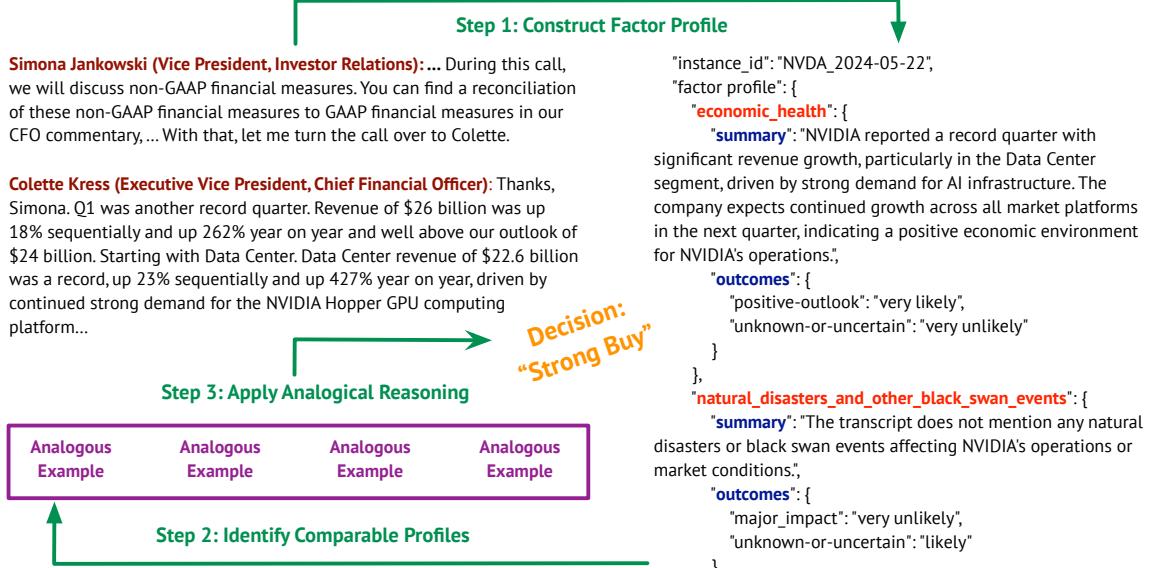


Figure 1: An excerpt from a typical earnings call transcript and its associated factor profile.

narios. It transforms lengthy speech transcripts into structured factor profiles, which contain key decision factors and their associated uncertainties. Analogical reasoning retrieves comparable profiles and draws on similar past cases to make critical decisions. DEFINE thereby adds a layer of transparency into the decision process.

- Our research demonstrates how to effectively extract decision factors from lengthy transcripts and use them to forecast post-earnings stock movements. This approach has the potential to generalize to other domains, such as consulting, financial investments, and political debates (Lehman et al., 2022), where discussions are complex and decisions carry significant consequences.

2 The DEFINE Framework

We investigate how LLMs make critical decisions about post-earnings stock movements using earnings call transcripts (Ni et al., 2024; Peters, 2024). An earnings call is a teleconference in which company executives discuss financial results with analysts and investors for a given quarter or fiscal year. As shown in Figure 5 in the Appendix, these transcripts typically consist of two parts: *prepared remarks* from the company’s executives and a subsequent *Q&A session*. During these calls, the executives provide a deep dive into the company’s financials, discuss key performance indicators, and share strategic plans for the future.

However, discussions of financials, e.g., revenue,

expenses, and profit margins, can be overwhelming. A factor profile seeks to distill these discussions into multiple variables, allowing decision-makers to focus on the most impactful factors (Eigner and Händler, 2024; Feng et al., 2024). Notably, if critical elements such as debt levels are not addressed by company executives, they can be marked as “unknown or uncertain.” This approach contrasts with traditional textual summaries of the transcript (Cho et al., 2021, 2022; Khatuya et al., 2024), which may be biased toward the topics emphasized by executives and discussed during the Q&A session.

2.1 Constructing Factor Profiles

Let X be an earnings call transcript used to predict a stock investment decision Y with five outcomes: strong buy, buy, hold, sell, and strong sell. We construct a **factor profile** for each transcript, defining a set of factors $\mathcal{F} = \{F_1, F_2, \dots, F_n\}$, where each factor F_i has potential outcomes $\{O_{i1}, O_{i2}, \dots, O_{im}\}$. The likelihood of each outcome, given X , is modeled by a probabilistic function $P(O_{ij}|X)$. These probabilities are inferred using a methodology that optimally integrates textual reasoning with quantitative analysis. Thus, each factor outcome’s probability informs the aggregation model that predicts Y .

Our study focuses on 15 key factors grouped into three categories: *macroeconomic influences* (e.g., economic health, market sentiment), *company-specific dynamics* (e.g., mergers and major acquisitions, product launches), and *historical financial*

1. Economic Health
2. Market Sentiment and Investor Psychology
3. Political Events and Government Policies
4. Natural Disasters and Black Swan Events
5. Geopolitical Issues
6. Mergers and Major Acquisitions
7. Regulatory Changes and Legal Issues
8. Financial Health
9. Company Growth
10. Company Product Launches
11. Supply Chain
12. Technological Innovation
13. Historical Earnings Per Share (EPS)
14. Historical Revenue
15. Historical Stock Prices

Table 1: A curated set of 15 factors for forecasting stock movements following earnings.

metrics (e.g., past earnings, stock prices). These factors were iteratively selected by querying the LLM for key variables in forecasting stock movements. We limit the set to 15 factors, each with two to three possible outcomes, ensuring a balance between complexity and performance while allowing future integration of analyst-identified factors. The full list of factors is in the Appendix.

We make use of the structured output capability of GPT-4o-2024-08-06 to extract factor profiles from earnings call transcripts. Following the framework set by Liu et al. (2024), we provide the LLM with a list of factors, their potential outcomes, and associated verbalized likelihoods. For each factor, the analysis involves two steps: first, the LLM creates a concise summary specific to that factor from the transcript; second, it assigns a verbalized likelihood to each possible outcome, ranging from “very unlikely” to “very likely.” Specifically, the likelihoods of outcomes, such as EPS, revenue trends, and historical stock prices, are derived from the company’s historical financial data. An example of the factor profile is shown in Figure 1, and the prompts used are detailed in the Appendix.

To convert these categorical likelihoods into probabilities, we employ the following normalization process: let $P_{i,j}$ denote the likelihood associated with the j -th outcome for the i -th factor. Here, verbalized likelihoods are converted to numerical values using the mapping $\{\text{very unlikely}=1, \text{unlikely}=2, \text{somewhat unlikely}=3, \text{somewhat likely}=4, \text{likely}=5, \text{very likely}=6\}$. Then, the probability $P(O_{ij}|X)$ is calculated as $P(O_{ij}|X) = \frac{P_{i,j}}{\sum_k P_{i,k}}$, ensuring the sum of outcomes for each factor equals 1. Alternative techniques, such as instructing the LLM to “distribute

10 points among the outcomes”, have been explored (Yang et al., 2024), our initial evaluation reveals that using verbalized likelihoods followed by normalization improves prediction accuracy compared to these direct probability distribution methods.

2.2 Analyzing Key Factors Using the Bradley-Terry Model

The Bradley-Terry model is a probabilistic framework used for estimating the relative strengths of items based on pairwise comparisons, and the outcome of each comparison indicates which of the two items is ‘better’ in a specific context (Bradley and Terry, 1952). This model has been widely used for ranking purposes in sports tournaments, LLM preference studies, and other domains where pairwise comparison data is available (Hu et al., 2023; Zhu et al., 2024). In this model, we estimate parameters that represent the strength of each factor. These parameters are generally presented on a logistic scale, where the probability that factor A is considered more significant than factor B is modeled as:

$$P(A > B) = \frac{e^{\beta_A}}{e^{\beta_A} + e^{\beta_B}} \quad (1)$$

Here, β_A and β_B represent the strengths of factors A and B, respectively. The estimated parameters are often exponentiated, so that $p_i = e^{\beta_i}$ measures the relative strength of each factor. A higher value indicates a stronger influence. In determining which factors to prioritize in a post-earnings analysis, those with higher Bradley-Terry scores are considered more crucial.

Consider a comparative analysis of two earnings call transcripts, A and B, transcript A is more likely to lead to favorable stock movements than transcript B ($A \succ B$). We obtain such pairwise comparisons based on target labels; with ‘strong-buy’ ranked higher than ‘hold’, ‘sell’, and ‘strong-sell’; ‘buy’ outranking ‘sell’ and ‘strong-sell’; and ‘hold’ surpassing ‘strong-sell’. The comparison of A and B will involve creating a set of factor-outcome pairwise comparisons, where each outcome in transcript A is preferable to that in transcript B: $O_{\cdot,\cdot}^{(A)} \succ O_{\cdot,\cdot}^{(B)}$, suggesting that the factors associated with transcript A outperform those in transcript B.

We further consider the weight-adjusted effect of comparisons between factors. Our method compares the influence of factors from transcripts A

and B by calculating an ‘expected occurrence’, which is determined by multiplying the likelihood of these factors appearing in both transcripts, $P(O_{ij}|X^{(A)}) \times P(O_{ij}|X^{(B)})$. This approach provides a probability-based comparison, offering a more detailed evaluation than simple counting methods. These expected occurrences then feed into a Bradley-Terry model matrix W . The model helps to estimate the relative importance of each factor by assigning a coefficient p_x to each outcome O_{ij} , indicating its influence on stock investment decisions. We refine these estimates using an EM-like algorithm, which iteratively adjusts and normalizes p_x to best fit the observed data.

$$p'_x = W_x \left(\sum_{y \neq x} \frac{w_{xy} + w_{yx}}{p_x + p_y} \right)^{-1} \quad p_x = \frac{p'_x}{\sum_{y=1}^M p'_y} \quad (2)$$

3 Bayesian Decision-Making

In Bayesian decision-making, utility functions play a crucial role in navigating uncertainty (Halawi et al., 2024; Lin et al., 2024; Ye et al., 2024). A Bayesian framework updates beliefs about possible outcomes. Decisions are then made by evaluating the expected utility for each possible action, which involves calculating the utility across the updated beliefs. This method ensures that choices are made to maximize expected utility, so decisions are aligned with the decision-maker’s preferences and risk tolerance.

Concretely, to compute $P(O_{ij}|X)$, we construct a probabilistic factor profile from a given earnings call transcript, where O_{ij} represents the j -th outcome of the i -th factor. The likelihood $P(Y|O_{ij})$, which estimates how the j -th outcome influences stock investment decisions, is calculated using the Bradley-Terry model. This model provides a framework for quantifying the impact each factor outcome has on the decision-making process. Using these probabilities, the Bayesian decision-making formula integrates over all factors and their potential outcomes to determine the optimal action. The overall decision is derived by:

$$\hat{Y} = \arg \max_Y \sum_i \sum_j P(Y|O_{ij}) P(O_{ij}|X) \quad (3)$$

The parameters calculated by the Bradley-Terry model for $P(Y|O_{ij})$ help us determine how each factor influences stock movements. During our testing phase, transcripts are assigned to one of five

decision categories based on their computed scores. For example, if the ground truth indicates there are k ‘strong buy’ recommendations, the top k scoring transcripts are classified correspondingly as ‘strong buy’. This approach uses probabilistic factor profiles in conjunction with Bradley-Terry modeling to identify influential factors, providing a transparent method for understanding decision-driving elements. Moving forward, we extend beyond individual factors by examining analogous cases that directly influence decisions.

4 Analogical Reasoning

Analogical reasoning, which involves drawing parallels between similar situations (Webb et al., 2023; Ozturkler et al., 2023; Yuan et al., 2024; Sourati et al., 2024; Yasunaga et al., 2024), is an effective method for decision-making. This approach is particularly useful when analyzing how stocks react to earnings announcements by referencing past, similar events. For example, in the tech sector, stocks often show high volatility after earnings calls that introduce significant technological updates, even if the revenue and EPS meet expectations. If a tech company is rumored to discuss a new technology trend in its upcoming earnings announcement, using this method, we can infer that this company’s stock might also experience increased volatility. Investors might use this analysis to make investment decisions or hedge against potential volatility.

Accurately identifying analogous examples from earnings call transcripts is crucial. We propose a method that utilizes probabilistic factor profiles, denoted as $P(O_{ij}|X)$, where O_{ij} represents the j -th outcome of the i -th factor. To measure the similarity between profiles, we calculate the Kullback-Leibler (KL) divergence, which quantifies the information loss when one probability distribution approximates another. The KL divergence is computed as follows:

$$D_{KL}(P||Q) = \sum_{i=1}^n \sum_{j=1}^m P(O_{ij}|X) \log \frac{P(O_{ij}|X)}{Q(O_{ij}|X_c)} \quad (4)$$

Here, P represents the factor profile for the target transcript, and Q denotes the profile for a comparative transcript X_c from our training set. Transcripts with lower KL divergence values are considered more analogous, and therefore more likely to influence investor decisions similarly.

Data Statistics

Num. of Transcripts	11,950
Num. of Companies	869
Avg. #Tokens per Transcript	10,187
Avg. #QA Pairs per Transcript	10
Avg. #Trans per Company	14
Avg. #Speakers per Transcript	12
Year Range	2017–2024

Table 2: Our dataset includes 11,950 earnings call transcripts from 800+ companies.

During testing, we identify the Top-K profiles that show the least divergence from a test instance’s profile and present these as analogical examples for the LLM to consider when reasoning about stock movements. The LLM is asked to select the most analogous example from the Top-K and carefully evaluates the current test instance to make its prediction. This approach ensures that the alignment between profiles is contextually appropriate, thereby drawing meaningful comparisons across different transcripts. By focusing on factor profiles rather than full transcripts or their summaries, we emphasize key market-moving information, avoiding unnecessary details. For example, Google and Broadcom could have analogous profiles even though their discussions in earnings calls might vary widely. Using factor profiles as analogous examples also requires significantly fewer tokens within the context window than full transcripts would.

5 Data Collection

Our dataset contains 11,950 earnings call transcripts from S&P 500 and NASDAQ 500 companies, gathered from the [Motley Fool](#) over the period of 2017–2024. The Motley Fool is a well-regarded financial service website that regularly publishes earnings call transcripts from U.S. companies. We make sure to follow their terms of use carefully during data collection. We do not use audio recordings or analyze acoustic or prosodic features. Each transcript is formatted as a JSON object, including the company’s stock ticker, the date of the earnings announcement, participant names and their affiliations, executive prepared remarks, and a series of question-answer pairs from the Q&A session.

Table 2 presents the statistics of our dataset. Each transcript averages 10,187 tokens and 133 sentences. They are sourced from 869 companies, each contributing an average of 14 transcripts. We

obtain company stock prices from Yahoo Finance via the `yfinance` package and financial metrics such as revenue and earnings per share (EPS) from [Alpha Advantage](#). Our dataset spans from 2017 to 2024. It enhances previous studies which examined earnings call transcripts from 2002–2010 ([Li et al., 2020](#)); these earlier transcripts may already be used in LLM pretraining. To avoid data contamination, we established a new test set consisting of the most recent 587 transcripts from 2024, which are beyond the pretraining cut-off date for LLMs.

We seek to make stock investment decisions by analyzing earnings call transcripts and focusing on performance over the 30-day period. We establish the ground truth decision on the 30th day following each earnings announcement ([Sonkiya et al., 2021](#)): a stock drop exceeding 5% corresponds to a ‘strong sell’ decision, a decrease between 2% and 5% leads to a ‘sell’, fluctuation within -2% to +2% is labeled ‘hold’, an increase between 2% and 5% is labeled a ‘buy’, and an increase above 5% is a ‘strong buy’. In our test set, the distribution of these labels is as follows: ‘strong buy’ at 34%, ‘buy’ at 15%, ‘hold’ at 21%, ‘sell’ at 9%, and ‘strong sell’ at 21%. This distribution is generally balanced, reflecting a slightly bullish market trend in 2024.

6 Experiments

In this section, we evaluate the decision-making performance of various systems, analyze the key factors that influence stock movement predictions, and conduct an analysis of analogical reasoning.

6.1 Decision-Making with DeFine

We evaluate our system, DeFine, against various decision-making strategies: (a) **LLM+CoT+Trans**: The LLM processes the full earnings call transcript, using chain-of-thought reasoning to assign a label with interpretations. (b) **LLM+CoT+Summ** and **LLM+CoT+Factors**: Both follow a summarize-then-predict approach. **LLM+CoT+Summ** generates a textual summary, while **LLM+CoT+Factors** condenses information into a structured factor profile. Details on the prompts are in the Appendix.

Our system, DeFine, utilizes analogical reasoning by analyzing five analogous cases identified using KL-divergence as the distance metric. It examines these cases alongside the current factor profile to predict an appropriate label. In contrast, DeLLMa uses a decision theory approach and has shown strong performance in agriculture planning

System	Recall	Prec.	F ₁	Accu.	Label	Recall	Prec.	F ₁
LLM+CoT+Trans	21.56	33.66	13.52	19.59	Strong Sell	7.32	37.50	12.24
LLM+CoT+Summ	22.77	16.17	14.12	20.61	Sell	5.56	9.09	6.90
LLM+CoT+Factors	24.38	28.58	17.26	22.32	Hold	29.84	28.24	29.02
DeLLMa (Liu et al., 2024)	38.30	23.14	16.68	22.35	Buy	44.83	18.93	26.62
DEFINE (Ours)	26.15	27.67	23.73	29.64	Strong Buy	43.22	44.56	43.88

Table 3: (*Left*) We show the accuracy and macro-averaged F-scores for various systems. Our system, DEFINE, which combines factor profiles with analogical reasoning, achieves the best performance. (*Right*) DEFINE’s performance across five categories: Strong Sell, Sell, Hold, Buy, and Strong Buy.

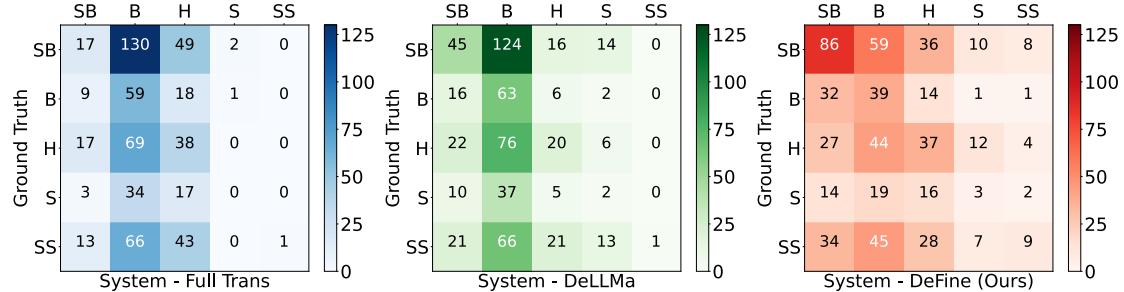


Figure 2: A comparison of confusion matrices from the LLM+CoT+Trans, DeLLMa, and DEFINE methods. While LLM+CoT+Trans and DeLLMa lean towards ‘Buy (B),’ DEFINE offers more balanced outcomes across all decision categories, showing notable improvement in ‘Strong Buy (SB),’ ‘Buy (B),’ ‘Hold (H),’ and ‘Sell (S)’ decisions.

and finance (Liu et al., 2024). For this approach, we pair each factor profile with possible labels and choose the top-ranked outcome as the final decision.

In Table 3 (left), we present the accuracy and macro-averaged F-scores for various systems, all using GPT-4o-2024-08-06. Our new system, DEFINE, which combines factor profiles with analogical reasoning, achieves the best performance. It surpasses the strong baseline system, DeLLMa, which involves ranking state-action pairs based on their preference levels as determined by the LLM. We find that LLMs generally make more accurate decisions when working with summaries rather than full transcripts; those transcripts typically contain around 10k tokens. This finding underscores the complexity of extracting and weighing key factors from lengthy transcripts, a task that remains challenging for most LLMs. In contrast, our factor profile method proves advantageous as it provides a balanced view of both macroeconomic factors and company-specific details, which are essential for rational decision-making.

We further analyze DEFINE’s performance across five categories: Strong Sell, Sell, Hold, Buy, and Strong Buy. Results are shown in Table 3 (right). DEFINE performs best at ‘Strong Buy’ recommendations and faces challenges with

‘Strong Sell’ categories. This may be due to its reliance on earnings call transcripts, which often contain optimistic remarks from executives aimed at reassuring investors, potentially skewing predictions away from ‘Strong Sell.’ Figure 2 includes a comparison of confusion matrices from the LLM+CoT+Trans, DeLLMa, and DEFINE methods. While LLM+CoT+Trans and DeLLMa predominantly lean towards ‘Buy,’ DEFINE offers more balanced outcomes across all decision categories, showing notable improvement in ‘Strong Buy,’ ‘Buy,’ ‘Hold,’ and ‘Sell’ decisions.

6.2 Influential Factors

We develop three variations of our DEFINE-BT approach, each using the Bradley-Terry model for pairwise comparisons in different contexts: DEFINE-BT-Same Sector compares companies within the same sector, DEFINE-BT-Cross Sectors examines companies across different sectors, and DEFINE-BT-Same Company analyzes a company’s current earnings call transcript against its historical ones. To ensure fairness, we maintain the same number of pairwise comparisons across all three settings, downsampling where necessary. According to the F-scores presented in Table 7, all DEFINE-BT variants outperform both the random baseline, which assigns investment decisions ran-

Factor (Outcome)	Salience
- Regulatory changes and legal issues happened (positive outlook)	0.0364
- Natural disasters and other black swan events (major impact)	0.0360
- Political events and government policies (major upheaval)	0.0349
- Geopolitical issues (escalation to conflict)	0.0345
- Supply chain (positive outlook)	0.0322
- Tech innovation (positive outlook)	0.0317
- Historical stock price change (bullish)	0.0316
- Historical EPS (bullish)	0.0315
- Financial health (positive outlook)	0.0311

Table 4: Influential factors that drive bullish investment decisions in the *Consumer Defensive* sector, e.g., food and beverage, household products, and grocery stores.

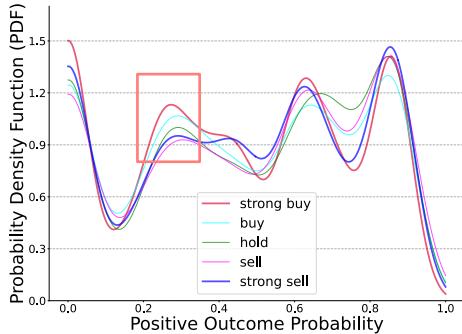


Figure 3: We analyze and plot the probability density function (PDF) of positive and negative factor outcomes for different investment decisions. Highlighted sections illustrate where the gaps between strong buy (red) and strong sell (blue) decisions are most pronounced.

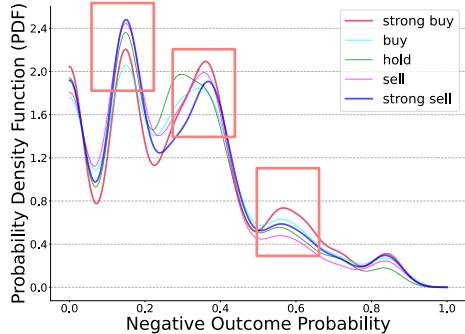
domly from five possible labels, and DeLLMa on the test set.

Among the three variants, **DEFINE-BT-Cross Sector** achieves the highest scores in both F-Score and Accuracy. This indicates that considering pairwise comparisons between earnings announcements from a diverse range of companies can enhance predictions of stock movements. Table 6 illustrates the performance of **DEFINE-BT-Cross Sector**, which was trained on one sector and tested on another. For this analysis, 100 earnings call transcripts were selected from each of the 11 financial sectors: Technology, Healthcare, Financial Services, Consumer Defensive, Energy, Industrials, Utilities, Basic Materials, Real Estate, Consumer Cyclical, and Communication Services.

Tables 4 and 5 highlight influential factors impacting investment decisions in the Consumer Defensive and Technology sectors, as identified by the Bradley-Terry model. In Consumer Defensive, which includes industries like food and beverage, household products, and grocery stores, significant drivers are natural disasters and black swan events, political events and government policies,

Factor (Outcome)	Salience
- Economic health (unknown or uncertain)	0.0362
- Market sentiment and investor psychology (unknown or uncertain)	0.0350
- Company growth (unknown or uncertain)	0.0338
- Supply chain (unknown or uncertain)	0.0326
- Geopolitical issues (escalation to conflict)	0.0322
- Historical revenue (decline)	0.0319
- Historical stock price change (bullish)	0.0318
- Tech innovation (unknown or uncertain)	0.0315
- Natural disasters and other black swan events (major impact)	0.0315
- Political events and government policies (major upheaval)	0.0313

Table 5: Factors that drive bullish investment decisions in the *Technology* sector, including industry leaders such as Apple, Microsoft, Amazon, Google, and Meta.



and geopolitical issues. These challenging macroeconomic circumstances often lead to buy-in decisions from investors. In contrast, the Technology sector, with industry leaders such as Apple, Microsoft, Amazon, Google, Meta, and Nvidia, shows that decisions to invest often hinge on unclear or uncertain factors. Technology stocks have seen considerable growth from 2017–2024. This pattern suggests that investment models may favor purchases in these companies despite encountering negative issues in earnings announcements.

In Figure 3, we analyze the probability of positive and negative factor outcomes, represented as a continuous random variable, and plot its probability density function (PDF) for various investment decisions. Highlighted sections illustrate where the gaps between strong buy (red) and strong sell (blue) decisions are most pronounced. Our analysis indicates that buy decisions often occur when the probability of positive outcomes is relatively low (about 0.2-0.3) and the likelihood of negative outcomes is moderate to high (ranging from 0.3 to 0.65), but not overly negative. Conversely, sell decisions tend to occur when negative outcome probabilities are min-

	Tech	FS	Health	CC	Ind	CS	CD	Energy	RE	BM	Util
Technology (Tech)	15.40	17.99	17.39	12.10	13.15	18.19	25.85	27.09	13.82	22.86	26.67
Financial Services (FS)	15.96	17.96	26.84	7.99	10.21	26.22	15.45	4.80	13.37	21.37	0.00
Healthcare (Health)	16.73	19.80	17.89	21.85	28.46	10.86	20.23	18.73	3.64	43.45	73.33
Consumer Cyclical (CC)	18.14	11.02	19.38	19.39	15.86	9.49	17.70	17.40	12.19	22.22	36.67
Industrials (Ind)	17.02	11.14	14.37	11.24	18.81	15.93	19.48	25.11	3.20	24.44	0.00
Communication Services (CS)	18.61	14.68	18.87	14.03	19.47	33.70	10.71	16.99	11.87	10.26	13.33
Consumer Defensive (CD)	24.91	21.71	19.15	19.89	21.38	2.67	23.09	12.50	9.72	29.52	50.00
Energy	19.49	16.50	23.62	14.25	19.03	8.90	19.03	8.98	12.10	28.79	0.00
Real Estate (RE)	22.86	15.74	16.76	14.08	12.34	4.00	15.61	11.28	12.34	43.18	0.00
Basic Materials (BM)	20.67	13.69	15.26	18.18	29.64	9.52	21.19	17.19	17.10	16.67	37.50
Utilities (Util)	17.82	27.75	23.15	26.25	12.61	25.49	20.63	5.70	18.40	14.29	53.33

Table 6: The performance of DEFINE-BT was evaluated by training it on one financial sector and testing it on another using 100 earnings call transcripts from each of the 11 sectors.

	F ₁	Accu.
Random Baseline	18.00	19.11
DeLLMa ((Liu et al., 2024))	16.68	22.53
DEFINE-BT-Same Sector	20.11	22.15
DEFINE-BT-Same Company	20.42	23.68
DEFINE-BT-Cross Sectors	24.45	27.43

Table 7: DEFINE-BT-Cross Sector achieves the highest scores, suggesting that considering pairwise comparisons from a diverse range of companies can enhance the predictions of stock movements.

imal (about 0.1-0.2). These observations suggest that rational investment decisions can sometimes appear counterintuitive: essentially, selling high and buying low. We find that a thorough analysis of various factors is advantageous. Our approach incorporates not just the known issues but also the uncertain or hidden factors, thereby enhancing the decision-making process.

6.3 Insights into Analogical Reasoning

Analogical reasoning utilizes a select number of analogous examples, denoted as K , to inform decision-making in LLMs. In Figure 4, we adjust K from 3 to 9 and observe its impact on the F-Score. In these experiments, we use the majority vote from the K examples as the final prediction. We find that $K = 4$ achieves the highest performance, potentially due to some tie-induced randomness compared to odd numbers. Typically, odd numbers for K are preferred for majority voting to avoid ties, with $K = 3, 5, 7$ showing similar effectiveness. For our system, DEFINE, we have opted for $K = 5$ to strike a balance between providing enough analogous examples and maintaining a manageable context length for the LLM.

Moreover, we examine how the most analogous examples influence DEFINE’s predictions. Our

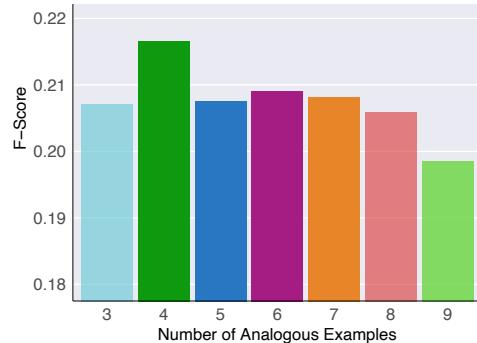


Figure 4: Analogical reasoning in LLMs uses analogous examples for decision-making. We vary K from 3 to 9 to assess its impact on the F-Score.

study finds that in 69% of cases, the LLM’s predictions match the labels from the most analogous examples. In the other 31% of cases, the LLM chooses to make its own predictions. E.g., when the analogous example is labeled “Strong Buy,” DEFINE concurs with “Strong Buy” in 63% of cases. It opts for “Buy” in 26% and “Hold” in 11% of the cases. Conversely, when the example is “Strong Sell,” DEFINE agrees with “Strong Sell” 50% of the time, chooses “Sell” in 25% of cases, and “Hold” in 12.5%. These results indicate that while DEFINE effectively utilizes analogous historical data to inform its predictions, it also critically evaluates the current factor profiles, demonstrating a balanced approach in its decision-making abilities.

7 Related Work

Analogical Reasoning . This type of reasoning identifies connections between similar, though not identical, situations to transfer knowledge from a known context to a new one (Webb et al., 2023; Ozturkler et al., 2023; Yu et al., 2024; Yuan et al., 2024; Sourati et al., 2024; Yasunaga et al., 2024). It helps decision makers draw parallels between

current situations and past experiences, effectively leveraging historical insights. Analogical reasoning plays a crucial role in various fields, e.g., doctors apply knowledge from one disease to diagnose another, and lawyers use past rulings to argue new cases (Lehman et al., 2022; Charmet et al., 2022; Cao et al., 2024a). This ability to recognize and use similarities in different situations is important for decision-making.

While zero-shot analogical reasoning is a desired capability for LLMs, recent studies show they lack the robustness and generality of human analogy-making, as evidenced by counterexamples in tasks such as letter string analogies (Hodel and West, 2024; Lewis and Mitchell, 2024). Musker et al. (2024) test both humans and LLMs on tasks that require transferring semantic structure and content between domains. Yasunaga et al. (2024) introduce analogical prompting, where LLMs self-generate relevant examples using prompts such as “*# Recall relevant problems and solutions:*” before solving the original problem; Qin et al. (2024) find that the accuracy of self-generated examples is key to eliciting such capability. Unlike previous research, our study employs probabilistic factor profiles to model analogical reasoning, grounding our approach in solid mathematical principles.

Decision-Making under Uncertainty. The use of LLMs in decision-making has surged due to their remarkable ability to reason over complex scenarios (Halawi et al., 2024; Lin et al., 2024; Ye et al., 2024; Band et al., 2024). However, the challenge of balancing a multitude of often conflicting factors in decision making remains understudied. For example, Falck et al. (2024) investigate whether adding more data points in in-context learning reduces uncertainty, as typically expected in Bayesian learning, and find evidence against this theory. The DELLMa framework (Liu et al., 2024) incorporates uncertainty into LLM decision-making using Bayesian networks and has been tested on tasks such as agriculture planning and finance. Feng et al. (2024) employ LLM entailment to map factors to context and utilize trained Bayesian models for probability estimation. Our work builds on these initiatives by integrating analogical reasoning with factor profiles to enhance the accuracy and transparency of LLM decision-making.

Financial Forecasting. Recent advancements in LLMs have revolutionized traditional financial tasks (Keith and Stent, 2019; Sawhney et al., 2020,

2021; Chuang and Yang, 2022; Ang and Lim, 2022; Sang and Bao, 2022; Medya et al., 2022; Wang et al., 2023; Koa et al., 2024; Srivastava et al., 2024). Notably, Chen et al. (2022) introduce FinQA, a dataset constructed from financial statements for assessing LLMs’ multi-step numerical reasoning. Moreover, TAT-QA (Zhu et al., 2021) tackles QA over tabular and textual data; FiNER (Loukas et al., 2022) focuses on numerical entity recognition; DocFinQA (Reddy et al., 2024) is a dataset designed for long-document financial QA; RiskLabs (Cao et al., 2024b) employs LLMs for financial risk assessments. Nie et al. (2024) provide a comprehensive survey on the use of LLMs across various financial domains. Our study focuses on analyzing earnings transcripts to understand how LLMs handle the ambiguities inherent in spoken language, thus providing insight into their decision-making under uncertainty. The research findings have broader applications including medical consultations, negotiations, and political debates.

8 Conclusion

We propose DEFINE, a new framework for decision-making in complex scenarios, such as those encountered in corporate earnings calls. By combining probabilistic factor profiles with analogical reasoning, this framework not only captures the uncertainties embedded in earnings call transcripts but also allows the LLM to apply previous insights to new challenges more efficiently. Our approach surpasses strong baseline models and enhances the practical utility of LLMs by identifying analogous examples. The DEFINE framework offers a promising avenue for navigating complex data and supporting decision-making processes.

Acknowledgements

We are grateful to the reviewers for their insightful feedback, which has helped improve our paper. This research has been partially supported by the NSF CAREER award, #2303655.

9 Limitations

The DEFINE framework, as detailed in our paper, offers a promising approach to enhancing decision-making through its innovative use of probabilistic factor profiles and analogical reasoning. Developed under carefully controlled experimental conditions, its potential is noteworthy. However, it’s important

to acknowledge that its performance might vary in the complexity of real-world settings due to factors such as data quality and specific contextual nuances. We encourage users to consider these variables when adapting the DEFINE approach to ensure its optimal application and to mitigate potential discrepancies between expected and actual results.

References

Gary Ang and Ee-Peng Lim. 2022. [Guided attention multimodal multitask financial forecasting with inter-company relationships and global and local news](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6313–6326, Dublin, Ireland. Association for Computational Linguistics.

Anthropic. 2025. [Claude 3.7 sonnet and claude code](#). Technical report. Accessed: 2025-05-13.

Neil Band, Xuechen Li, Tengyu Ma, and Tatsunori Hashimoto. 2024. [Linguistic calibration of long-form generations](#). *Preprint*, arXiv:2404.00474.

Kaj Bostrom, Zayne Sprague, Swarat Chaudhuri, and Greg Durrett. 2022. [Natural language deduction through search over statement compositions](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 4871–4883, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Ralph Allan Bradley and Milton E. Terry. 1952. [Rank analysis of incomplete block designs: I. the method of paired comparisons](#). *Biometrika*, 39(3/4):324–345.

Lang Cao, Zifeng Wang, Cao Xiao, and Jimeng Sun. 2024a. [PILOT: Legal case outcome prediction with case law](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 609–621, Mexico City, Mexico. Association for Computational Linguistics.

Yupeng Cao, Zhi Chen, Qingyun Pei, Fabrizio Dimino, Lorenzo Ausiello, Prashant Kumar, K. P. Subbalakshmi, and Papa Momar Ndiaye. 2024b. [Risklabs: Predicting financial risk using large language model based on multi-sources data](#). *Preprint*, arXiv:2404.07452.

Thibault Charmet, Inès Cherichi, Matthieu Allain, Urszula Czerwinska, Amaury Fouret, Benoît Sagot, and Rachel Bawden. 2022. [Complex labelling and similarity prediction in legal texts: Automatic analysis of France’s court of cassation rulings](#). In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 4754–4766, Marseille, France. European Language Resources Association.

Yanda Chen, Joe Benton, Ansh Radhakrishnan, Jonathan Uesato, Carson Denison, John Schulman, Arushi Somani, Peter Hase, Misha Wagner, Fabien Roger, Vlad Mikulik, Samuel R. Bowman, Jan Leike, Jared Kaplan, and Ethan Perez. 2025. [Reasoning models don’t always say what they think](#). *Preprint*, arXiv:2505.05410.

Zhiyu Chen, Wenhui Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, and William Yang Wang. 2022. [Finqa: A dataset of numerical reasoning over financial data](#). *Preprint*, arXiv:2109.00122.

Sangwoo Cho, Franck Dernoncourt, Tim Ganter, Trung Bui, Nedim Lipka, Walter Chang, Hailin Jin, Jonathan Brandt, Hassan Foroosh, and Fei Liu. 2021. [StreamHover: Livestream transcript summarization and annotation](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6457–6474, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.

Sangwoo Cho, Kaiqiang Song, Xiaoyang Wang, Fei Liu, and Dong Yu. 2022. [Toward unifying text segmentation and long document summarization](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 106–118, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Chengyu Chuang and Yi Yang. 2022. [Buy tesla, sell ford: Assessing implicit stock market preference in pre-trained language models](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 100–105, Dublin, Ireland. Association for Computational Linguistics.

Eva Eigner and Thorsten Händler. 2024. [Determinants of llm-assisted decision-making](#). *Preprint*, arXiv:2402.17385.

Fabian Falck, Ziyu Wang, and Chris Holmes. 2024. [Is in-context learning in large language models bayesian? a martingale perspective](#). *Preprint*, arXiv:2406.00793.

Yu Feng, Ben Zhou, Weidong Lin, and Dan Roth. 2024. [Bird: A trustworthy bayesian inference framework for large language models](#). *Preprint*, arXiv:2404.12494.

Muhan Gao, TaiMing Lu, Kuai Yu, Adam Byerly, and Daniel Khashabi. 2024. [Insights into LLM long-context failures: When transformers know but don’t tell](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7611–7625, Miami, Florida, USA. Association for Computational Linguistics.

Danny Halawi, Fred Zhang, Chen Yueh-Han, and Jacob Steinhardt. 2024. [Approaching human-level forecasting with language models](#). *Preprint*, arXiv:2402.18563.

Damian Hodel and Jevin West. 2024. [Response: Emergent analogical reasoning in large language models](#). *Preprint*, arXiv:2308.16118.

Yebowen Hu, Kaiqiang Song, Sangwoo Cho, Xiaoyang Wang, Hassan Foroosh, and Fei Liu. 2023. [Decipherpref: Analyzing influential factors in human preference judgments via gpt-4](#). *Preprint*, arXiv:2305.14702.

Yebowen Hu, Kaiqiang Song, Sangwoo Cho, Xiaoyang Wang, Hassan Foroosh, Dong Yu, and Fei Liu. 2024a. [SportsMetrics: Blending text and numerical data to understand information fusion in LLMs](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 267–278, Bangkok, Thailand. Association for Computational Linguistics.

Yebowen Hu, Kaiqiang Song, Sangwoo Cho, Xiaoyang Wang, Wenlin Yao, Hassan Foroosh, Dong Yu, and Fei Liu. 2024b. [When reasoning meets information aggregation: A case study with sports narratives](#). *Preprint*, arXiv:2406.12084.

Jie Huang and Kevin Chen-Chuan Chang. 2023. [Towards reasoning in large language models: A survey](#). *Preprint*, arXiv:2212.10403.

Koray Kavukcuoglu. 2025. [Gemini 2.5: Our most intelligent ai model](#). Technical report. Accessed: 2025-05-13.

Katherine Keith and Amanda Stent. 2019. [Modeling financial analysts’ decision making via the pragmatics and semantics of earnings calls](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 493–503, Florence, Italy. Association for Computational Linguistics.

Subhendu Khatuya, Koushiki Sinha, Niloy Ganguly, Saptarshi Ghosh, and Pawan Goyal. 2024. [Instruction-guided bullet point summarization of long financial earnings call transcripts](#). *Preprint*, arXiv:2405.06669.

Kelvin J.L. Koa, Yunshan Ma, Ritchie Ng, and Tat-Seng Chua. 2024. [Learning to generate explainable stock predictions using self-reflective large language models](#). In *Proceedings of the ACM Web Conference 2024*, volume 12706 of *WWW ’24*, page 4304–4315. ACM.

Kalpesh Krishna, Erin Bransom, Bailey Kuehl, Mohit Iyyer, Pradeep Dasigi, Arman Cohan, and Kyle Lo. 2023. [LongEval: Guidelines for human evaluation of faithfulness in long-form summarization](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1650–1669, Dubrovnik, Croatia. Association for Computational Linguistics.

Philippe Laban, Alexander R. Fabbri, Caiming Xiong, and Chien-Sheng Wu. 2024. [Summary of a haystack: A challenge to long-context llms and rag systems](#). *Preprint*, arXiv:2407.01370.

Eric Lehman, Vladislav Lialin, Katelyn Edelwina Legaspi, Anne Janelle Sy, Patricia Therese Pile, Nicole Rose Alberto, Richard Raymund Ragasa, Corinna Victoria Puyat, Marianne Katharina Taliño, Isabelle Rose Alberto, Pia Gabrielle Alfonso, Dana Moukheiber, Byron Wallace, Anna Rumshisky, Jennifer Liang, Preethi Raghavan, Leo Anthony Celi, and Peter Szolovits. 2022. [Learning to ask like a physician](#). In *Proceedings of the 4th Clinical Natural Language Processing Workshop*, pages 74–86, Seattle, WA. Association for Computational Linguistics.

Martha Lewis and Melanie Mitchell. 2024. [Using counterfactual tasks to evaluate the generality of analogical reasoning in large language models](#). *Preprint*, arXiv:2402.08955.

Dacheng Li, Shiyi Cao, Tyler Griggs, Shu Liu, Xiangxi Mo, Eric Tang, Sumanth Hegde, Kourosh Hakhamaneshi, Shishir G. Patil, Matei Zaharia, Joseph E. Gonzalez, and Ion Stoica. 2025. [Llms can easily learn to reason from demonstrations structure, not content, is what matters!](#) *Preprint*, arXiv:2502.07374.

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, CIKM ’20*, page 3063–3070, New York, NY, USA. Association for Computing Machinery.

Zhen Lin, Shubhendu Trivedi, and Jimeng Sun. 2024. [Generating with confidence: Uncertainty quantification for black-box large language models](#). *Preprint*, arXiv:2305.19187.

Nelson F. Liu, Kevin Lin, John Hewitt, Ashwin Paranjape, Michele Bevilacqua, Fabio Petroni, and Percy Liang. 2023. [Lost in the middle: How language models use long contexts](#). *Preprint*, arXiv:2307.03172.

Ollie Liu, Deqing Fu, Dani Yogatama, and Willie Neiswanger. 2024. [Dellma: A framework for decision making under uncertainty with large language models](#). *Preprint*, arXiv:2402.02392.

Lefteris Loukas, Manos Fergadiotis, Ilias Chalkidis, Irini Spyropoulou, Prodromos Malakasiotis, Ion Androutsopoulos, and Georgios Paliouras. 2022. [Finer: Financial numeric entity recognition for xbrl tagging](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. Association for Computational Linguistics.

Yiming Lu, Yebowen Hu, Hassan Foroosh, Wei Jin, and Fei Liu. 2025. [STRUX: An LLM for decision-making with structured explanations](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 131–141, Albuquerque, New Mexico. Association for Computational Linguistics.

Sourav Medya, Mohammad Rasoolinejad, Yang Yang, and Brian Uzzi. 2022. *An exploratory study of stock price movements from earnings calls*. *Preprint*, arXiv:2203.12460.

Philipp Mondorf and Barbara Plank. 2024. *Beyond accuracy: Evaluating the reasoning behavior of large language models – a survey*. *Preprint*, arXiv:2404.01869.

Rajdeep Mukherjee, Abhinav Bohra, Akash Banerjee, Soumya Sharma, Manjunath Hegde, Afreen Shaikh, Shivani Shrivastava, Koustuv Dasgupta, Niloy Ganguly, Saptarshi Ghosh, and Pawan Goyal. 2022. *Ect-sum: A new benchmark dataset for bullet point summarization of long earnings call transcripts*. *Preprint*, arXiv:2210.12467.

Sam Musker, Alex Duchnowski, Raphaël Millière, and Ellie Pavlick. 2024. *Semantic structure-mapping in llm and human analogical reasoning*. *Preprint*, arXiv:2406.13803.

Haowei Ni, Shuchen Meng, Xupeng Chen, Ziqing Zhao, Andi Chen, Panfeng Li, Shiyao Zhang, Qifu Yin, Yuanqing Wang, and Yuxi Chan. 2024. *Harnessing earnings reports for stock predictions: A qloran-enhanced llm approach*. In *2024 6th International Conference on Data-driven Optimization of Complex Systems (DOCS)*, page 909–915. IEEE.

Yuqi Nie, Yaxuan Kong, Xiaowen Dong, John M. Mulvey, H. Vincent Poor, Qingsong Wen, and Stefan Zohren. 2024. *A survey of large language models for financial applications: Progress, prospects and challenges*. *Preprint*, arXiv:2406.11903.

OpenAI, ;, Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, Ahmed El-Kishky, Aiden Low, Alec Helyar, Aleksander Madry, Alex Beutel, Alex Carney, Alex Iftimie, Alex Karpenko, Alex Tachard Passos, Alexander Neitz, Alexander Prokofiev, Alexander Wei, Allison Tam, and 244 others. 2024. *Openai o1 system card*. *Preprint*, arXiv:2412.16720.

Batu Ozturkler, Nikolay Malkin, Zhen Wang, and Nebojsa Jojic. 2023. *Thinksum: Probabilistic reasoning over sets using large language models*. *Preprint*, arXiv:2210.01293.

Gary Peters. 2024. *Hedge funds’ use of artificial intelligence and machine learning technologies*. Technical report, U.S. Senate Committee on Homeland Security and Governmental Affairs. Accessed: 2025-05-12.

Chengwei Qin, Wenhan Xia, Tan Wang, Fangkai Jiao, Yuchen Hu, Bosheng Ding, Ruirui Chen, and Shafiq Joty. 2024. *Relevant or random: Can llms truly perform analogical reasoning?* *Preprint*, arXiv:2404.12728.

Varshini Reddy, Rik Koncel-Kedziorski, Viet Dac Lai, Michael Krundick, Charles Lovering, and Chris Tanner. 2024. *Docfinqa: A long-context financial reasoning dataset*. *Preprint*, arXiv:2401.06915.

Z. Z. Ren, Zhihong Shao, Junxiao Song, Huajian Xin, Haocheng Wang, Wanjia Zhao, Liyue Zhang, Zhe Fu, Qihao Zhu, Dejian Yang, Z. F. Wu, Zhibin Gou, Shiron Ma, Hongxuan Tang, Yuxuan Liu, Wenjun Gao, Daya Guo, and Chong Ruan. 2025. *Deepseek-prover-v2: Advancing formal mathematical reasoning via reinforcement learning for subgoal decomposition*. *Preprint*, arXiv:2504.21801.

Yunxin Sang and Yang Bao. 2022. *DialogueGAT: A graph attention network for financial risk prediction by modeling the dialogues in earnings conference calls*. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1623–1633, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Ramit Sawhney, Mihir Goyal, Prakhar Goel, Puneet Mathur, and Rajiv Ratn Shah. 2021. *Multimodal multi-speaker merger & acquisition financial modeling: A new task, dataset, and neural baselines*. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6751–6762, Online. Association for Computational Linguistics.

Ramit Sawhney, Piyush Khanna, Arshiya Aggarwal, Taru Jain, Puneet Mathur, and Rajiv Ratn 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, Online. Association for Computational Linguistics.

Priyank Sonkiya, Vikas Bajpai, and Anukriti Bansal. 2021. *Stock price prediction using bert and gan*. *Preprint*, arXiv:2107.09055.

Zhivar Sourati, Filip Ilievski, Pia Sommerauer, and Yifan Jiang. 2024. *Arn: Analogical reasoning on narratives*. *Preprint*, arXiv:2310.00996.

Zayne Sprague, Xi Ye, Kaj Bostrom, Swarat Chaudhuri, and Greg Durrett. 2024. *Musr: Testing the limits of chain-of-thought with multistep soft reasoning*. *Preprint*, arXiv:2310.16049.

Pragya Srivastava, Manuj Malik, Vivek Gupta, Tanuja Ganu, and Dan Roth. 2024. *Evaluating llms’ mathematical reasoning in financial document question answering*. *Preprint*, arXiv:2402.11194.

Liping Wang, Jiawei Li, Lifan Zhao, Zhizhuo Kou, Xiaohan Wang, Xinyi Zhu, Hao Wang, Yanyan Shen, and Lei Chen. 2023. *Methods for acquiring and incorporating knowledge into stock price prediction: A survey*. *Preprint*, arXiv:2308.04947.

Taylor Webb, Keith J. Holyoak, and Hongjing Lu. 2023. *Emergent analogical reasoning in large language models*. *Preprint*, arXiv:2212.09196.

Joshua C. Yang, Damian Dailisan, Marcin Korecki, Catarina I. Hausladen, and Dirk Helbing. 2024. [Llm voting: Human choices and ai collective decision making](#). *Preprint*, arXiv:2402.01766.

Michihiro Yasunaga, Xinyun Chen, Yujia Li, Panupong Pasupat, Jure Leskovec, Percy Liang, Ed H. Chi, and Denny Zhou. 2024. [Large language models as analogical reasoners](#). *Preprint*, arXiv:2310.01714.

Yining Ye, Xin Cong, Shizuo Tian, Yujia Qin, Chong Liu, Yankai Lin, Zhiyuan Liu, and Maosong Sun. 2024. [Rational decision-making agent with internalized utility judgment](#). *Preprint*, arXiv:2308.12519.

Junchi Yu, Ran He, and Rex Ying. 2024. [Thought propagation: An analogical approach to complex reasoning with large language models](#). *Preprint*, arXiv:2310.03965.

Siyu Yuan, Jiangjie Chen, Changzhi Sun, Jiaqing Liang, Yanghua Xiao, and Deqing Yang. 2024. [Analogykb: Unlocking analogical reasoning of language models with a million-scale knowledge base](#). *Preprint*, arXiv:2305.05994.

Banghua Zhu, Jiantao Jiao, and Michael I. Jordan. 2024. [Principled reinforcement learning with human feedback from pairwise or k-wise comparisons](#). *Preprint*, arXiv:2301.11270.

Fengbin Zhu, Wenqiang Lei, Youcheng Huang, Chao Wang, Shuo Zhang, Jiancheng Lv, Fuli Feng, and Tat-Seng Chua. 2021. [Tat-qa: A question answering benchmark on a hybrid of tabular and textual content in finance](#). *Preprint*, arXiv:2105.07624.

A Influential Factors

- **Macroeconomic Influences.** These encompass broad economic factors that affect the entire market or large segments of it. This includes the overall economic health, market sentiment, political events, natural disasters and geopolitical issues (Liu et al., 2024). Each factor leads to two potential outcomes; for instance, natural disasters might cause a ‘Major Impact’ by disrupting economies and global supply chains, and directly affecting market performance; the ‘Unknown or Uncertain’ outcome reflects the unpredictability of such events.
- **Company-Specific Dynamics.** These factors are linked to the internal operations and strategic decisions of individual companies, such as mergers and acquisitions, regulatory changes, financial health, company growth potential, product launches, and issues within the supply chain. Each factor can result in one of

two potential outcomes. For example, a ‘Positive Outlook’ on regulatory changes can open up new business opportunities, whereas ‘Unknown or Uncertain’ could signify regulatory uncertainties that lead to financial challenges.

- **Historical Financial Metrics.** Important metrics include historical earnings per share (EPS), revenue trends, and past stock price movements. Each factor can result in three outcomes: ‘Bullish’, where metrics like earnings per share, revenue, and stock prices consistently rise, indicating strong financial health; ‘Stable’, characterized by steady movements; ‘Bearish,’ marked by declining financial figures, possibly leading investors to be pessimistic about the company’s future performance.

Transcript Excerpt from Delta Air Lines (DAL) Earnings Call

[Prepared Remarks:]

>> Operator

Good morning, everyone, and welcome to the Delta Air Lines September-quarter 2021 financial results conference call. My name is Jen, and I will be your coordinator. [Operator instructions] As a reminder, today's call is being recorded. I would now like to turn the conference over to Ms. Julie Stewart, vice president of investor relations...

>> Julie Stewart -- Vice President of Investor Relations

Thank you, Jen. Good morning, everyone, and thanks for joining us for our September-quarter 2021 earnings call. Joining us from Atlanta today are CEO, Ed Bastian; our president, Glen Hauenstein; our CFO, Dan Janki. And Ed will open the call with an overview of Delta's performance and strategy.

Glen will provide an update on the revenue environment and our brand momentum, and Dan will discuss cost fleet and our balance sheet. Similar to last quarter's call, we've scheduled today's call for 90 minutes to make sure we have plenty of time for questions. [Operator instructions] After the analyst Q&A, we will move to our media questions, after which, Ed will provide a brief closing statement. Today's discussion contains forward-looking statements that represent our beliefs or expectations about future events.

All forward-looking statements involve risks and uncertainties that could cause the actual results to differ materially from the forward-looking statements. Some of the factors that may cause such differences are described in Delta's SEC filings. We also discuss non-GAAP financial measures, and all results exclude special items unless otherwise noted. You can find a reconciliation of our non-GAAP measures on the Investor Relations page at ir.delta.com. And with that, I'll turn the call over to Ed.

>> Ed Bastian -- Chief Executive Officer

Well, thank you, Julie, and good morning, everyone. Appreciate you joining us this morning. The September quarter marked another important milestone in our recovery. We achieved our first quarterly profit since the start of the pandemic with a pre-tax result of \$216 million and a pre-tax margin of nearly 3% despite still missing one-third of our revenue base compared to the same period in 2019... [omitted.]

[Questions & Answers:]

>> Operator

Thank you. And we'll go first to Jamie Baker with J.P. Morgan.

>> Jamie Baker -- J.P. Morgan -- Analyst

Hey. Good morning, everybody. First question goes potentially to Glen and Dan. So pre-COVID, I had asked Paul about the amount of time that it would typically take Delta to recalibrate the higher fuel prices.

I'm not staring at the transcript, but his estimate at the time was four to six months, which was an improvement from historic levels. So my question, I guess, for Glen is whether the booking curve is steep enough right now that you might actually be able to recapture the top line more quickly than that. And similarly, for Dan, whether there's anything we should be taking on the cost or operations side that could accelerate the process. I'm basically just trying to understand whether four to six months is still the right estimate for us to be using.

>> Glen Hauenstein -- President

Well, I would just comment, I think we're a bit in uncharted territory here as the recovery continues. And while I think it might be difficult in the very short run, despite the fact that the booking curve has moved in a bit, that I would estimate that, that four to six months is about right because we believe that demand and capacity will fall back into a very good equilibrium by next spring which would put you inside that window...

Figure 5: Transcript Excerpt from Delta Air Lines (DAL) Earnings Call. Adapted from (Lu et al., 2025).

Building a Factor Profile from the Earnings Call Transcript

System Message

You are a financial analyst specializing in earnings call transcripts. You will receive the complete transcript of an earnings call, which includes both the prepared remarks and the Q&A session. Your job is to identify the key factors from the transcript and assign probabilities to the potential outcomes of these factors.

User Message

Your task is to conduct a comprehensive analysis of the earnings call transcript below. Be sure to accurately capture the important factors and estimate the likelihood of each factor resulting in specific outcomes.

Earnings Call Transcript for Company {Company}

Prepared Remarks

Speaker: {Speech...}

...

Questions and Answers

Analyst: {Speech...}

...

Please analyze the above earnings call transcript, focusing on the following key factors:
{Enumerate factors, descriptions, and outcomes}

1. Economic health: Economic health refers to the overall stability and performance of the economy, reflected in factors like growth, employment, inflation, and market confidence.
Outcomes: {positive-outlook, unknown-or-uncertain}

2. Market sentiment and investor psychology: Market sentiment reflects the overall mood or attitude of investors toward a particular market, influenced by news, economic data, and global events. Investor psychology refers to the emotions and cognitive biases that drive decisions, often leading to behaviors like fear-driven selling or greed-fueled buying. Outcomes:
{optimistic, unknown-or-uncertain}

...

Please take the time to thoroughly understand the transcript. For each key factor, provide a detailed summary based on the given transcript. Then, review all associated outcomes and assess the likelihood of each outcome. The likelihood should be strictly selected from the following options: {very likely, likely, somewhat likely, somewhat unlikely, unlikely, very unlikely}. Format your response in JSON.

Example Output:

{JSON output example}

Your Output:

Figure 6: Building a Factor Profile from the Earnings Call Transcript

Applying Analogical Reasoning to Investment Decisions

System Message

You're a financial analyst who specializes in giving investors buy or sell recommendations by thoroughly analyzing earnings call transcripts.

User Message

Here are several example company profiles. Each profile highlights key factors from an earnings call transcript and probabilities for potential outcomes based on those factors. Each profile represents a specific company and is based on its historical earnings call data. Your job is to pick the most analogous example and use its strategy to solve the initial problem.

Example Company Profile 1:

{Factor Profile 1}

Analyst recommendation: {Action 1}

Example Company Profile 2:

{Factor Profile 2}

Analyst recommendation: {Action 2}

Example Company Profile 3:

{Factor Profile 3}

Analyst recommendation: {Action 3}

Example Company Profile 4:

{Factor Profile 4}

Analyst recommendation: {Action 4}

Example Company Profile 5:

{Factor Profile 5}

Analyst recommendation: {Action 5}

**** Initial Problem ****

Based on your analysis of the earnings call for {Company Name} held on {Announcement Date}, decide on the most likely analyst recommendation for the next 30 days from these options:

- Action 1: strong buy: The stock price will increase by more than 5%
- Action 2: buy: The stock price will increase by 2% to 5%
- Action 3: hold: The stock price is expected to remain stable, fluctuating between -2% to 2%
- Action 4: sell: The stock price will decrease by 2% to 5%
- Action 5: strong sell: The stock price will decrease by more than 5%

Below is the company profile summarized from {Company Name}'s earnings call on {Announcement Date} and the historical price trend probabilities judged by an analyst:

{Factor Profile Constructed Using an Earnings Call Transcript}

**** Solve the Initial Problem ****

Please respond with the analyst recommendation for this stock in JSON format, including these keys: ('idx', 'recommendation', 'justification'). 'idx' is the index of the most analogous example profile, and 'recommendation' should be one of the actions mentioned above for 30 days of trading, and 'justification' should clearly explain your recommendation using the strategy you learned from the selected example company profile.

Figure 7: Applying Analogical Reasoning to Investment Decisions.

Chain-of-Thought Prompt for Investment Decisions

System Message

You're a financial analyst specializing in giving investors buy or sell recommendations by thoroughly analyzing earnings call transcripts.

User Message

Based on your analysis of the earnings call for {Company Name} held on {Announcement Date}, decide on the most likely analyst recommendation for the next 30 days from these options:

- Action 1: strong buy: The stock price will increase by more than 5%
- Action 2: buy: The stock price will increase by 2% to 5%
- Action 3: hold: The stock price is expected to remain stable, fluctuating between -2% to 2%
- Action 4: sell: The stock price will decrease by 2% to 5%
- Action 5: strong sell: The stock price will decrease by more than 5%

Below is the {Factor Profile, Transcript or Summary} from {Company Name}'s earnings call on {Announcement Date}:}

{Factor Profile, Transcripts or Summary}

Please think step by step and respond with the analyst recommendation for this stock in JSON format, including these keys: ('thoughts', 'recommendation', 'justification'). 'Thoughts' should be your detailed reasoning steps, 'recommendation' should be one of the actions mentioned above for 30 days trading, 'Justification' should clearly explain your recommendation using the strategy you learned from the selected example company profile.

Figure 8: Chain-of-Thought Prompt for Soliciting Investment Decisions.

Prompt for Analyzing Trends in Historical Financial Data

System Message

You are a financial analyst specializing in historical data analysis, including stock prices, earnings per share (EPS), and revenue. Your goal is to assess the likelihood of different market trends based on past data.

User Message

The potential outcomes to consider are: {bullish, stable, and bearish}. For each outcome, please assign a likelihood level from the following options: {very likely, likely, somewhat likely, somewhat unlikely, unlikely, very unlikely}.

Below, you will be provided with a historical data table: {Data Name}:{Description}
{Historical Data Table}

Date	Close Price
2023-07-31	195.22
2023-08-01	195.46
{...until the date of the earnings announcement.}	

Please analyze this historical data and provide the likelihood of each outcome in JSON format.

```
# Example Output:  
{"historical EPS":{"bullish": "very likely", "stable": "somewhat likely", "bearish": "unlikely"}  
# Your Output:
```

Figure 9: Prompt for Analyzing Trends in Historical Financial Data.