

# Can LLMs Learn Macroeconomic Narratives from Social Media?

Anonymous ACL submission

## Abstract

This study empirically tests the *Narrative Economics* hypothesis, which posits that narratives (ideas that are spread virally and affect public beliefs) can influence economic fluctuations. We introduce two curated datasets containing posts from X (formerly Twitter) which capture economy-related narratives. Employing Natural Language Processing (NLP) methods, we extract and summarize narratives from the tweets. We test their predictive power for *macroeconomic* forecasting by incorporating the tweets' or the extracted narratives' representations in downstream financial prediction tasks. Our work highlights the challenges in improving macroeconomic models with narrative data, paving the way for the research community to realistically address this important challenge. From a scientific perspective, our investigation offers valuable insights and NLP tools for narrative extraction and summarization using Large Language Models (LLMs), contributing to future research on the role of narratives in economics.<sup>1</sup>

## 1 Introduction

Narrative Economics studies how popular narratives change over time and interact with economic behavior (Shiller, 2017). A key proposition within this field is that narratives can drive economic fluctuations. This is especially intriguing at the macroeconomic level, as the theory suggests that widely shared stories can influence the collective decisions of millions of individuals. However, it presents greater challenges than for microeconomy, due to the complex interplay of numerous factors, the need for broad-covering narratives, and the inherent difficulty in isolating causal relationships.

Proving a causal link between narratives and economic changes remains challenging (Lucas and Sargent; an econometric identifiability problem (Angrist and Pischke, 2009)). Instead, we propose

applying NLP methods to represent narratives extracted from social media and test their macroeconomic predictive power. We introduce two Twitter (X) datasets crafted to capture economy-related narratives (§2), explore two NLP approaches for narrative extraction, and utilize them for prediction. First, we build prediction models directly on raw tweets, implicitly capturing the narratives within them. Second, we utilize LLMs<sup>2</sup> to explicitly extract and summarize narratives, generating a third dataset of LLM analyses. We demonstrate the presence of narratives within our datasets (§5) and the effectiveness of our approaches in representing the aggregated economic-narrative picture.

Existing evaluation strategies for narratives' effectiveness for economic predictions are inadequate, focusing on macroeconomic correlations and anomalies, or on microeconomic predictions (§8). We propose utilizing downstream macroeconomic prediction tasks, where representations of latent narratives serve as input (§4-7). Interestingly, our results reveal that the successful narrative extraction we demonstrate offer only marginal improvement in macroeconomic prediction, compared to utilizing only financial information. In Section 9 we discuss the results' implications on the validity of the Narrative Economics theory, at least when it comes to macroeconomics.

Our contributions include:

- A framework for testing the narrative hypothesis on macroeconomic predictions.
- Three curated tailored datasets: two Twitter collections and LLM-based analyses of tweets.
- Extensive analysis demonstrating the effectiveness of our methodology, revealing valuable insights while highlighting the need for new models and tasks to empirically test the Narrative Economics theory.

<sup>2</sup>We refer to a variety of models as LLMs, ranging from relatively small models (e.g., BERT (Devlin et al., 2018)) to more recent models (e.g., GPT-3.5 (OpenAI, 2021)).

<sup>1</sup>Data will be shared upon paper acceptance.

## 2 Data Resources

### 2.1 Twitter (X) Collected Datasets

We use Twitter (X) as a source of such narratives given its real-time reflection of diverse public opinions, and collect two tweets datasets capturing economy-related narratives. Unlike previous research focused on stock- or company- specific tweets (Vamossy and Skog, 2021; Chandra and Jana, 2020; Mengistie and Kumar, 2021; Sethi et al., 2020; Karami et al., 2020), our approach includes a broader time-frame, diverse global topics, and comprehensive keywords to capture a wide spectrum of perspectives.

The first dataset ranges from Twitter’s early days until the COVID-19 pandemic which disrupted the economy, while the second captures more recent trends, late 2021 onward. Both were carefully curated using targeted queries, restricted to English non-retweets with the inclusion of specific keywords, and were analyzed to ensure quality and relevance for capturing economic narratives.

**Pre-Pandemic Twitter Dataset.** We utilize Twitter API (Twitter) to collect a comprehensive dataset of 2.4 million tweets from January 2007 to December 2020, covering six topic areas with potential economic impact (economics, business, politics, current affairs, human disasters, and natural disasters) using targeted keywords detailed in App. A. For each topic, we iteratively retrieved tweets for each date within the timeframe. To prioritize viral tweets, we first retrieved the top 200 tweets based on follower count, then we randomly sampled 100 to mitigate potential bias towards highly active accounts typically associated with news outlets. This process yielded an average of roughly 400,000 tweets per topic, contributed by about 250,000 users per topic, each boasting an average follower count of 100 million, including global leaders, news outlets and other influencers. Minimal pre-processing was applied to the tweets, as common in literature (see §8 and App. A).

**Post-2021 Twitter Dataset.** This dataset, spanning September 2021 to July 2023, was specifically constructed for LLM-driven narrative analysis (see §4.2). To fairly test the predictive power of narratives, we needed to ensure that the employed LLM (Chat Completion API with GPT-3.5 (OpenAI, 2021), data cutoff September 2021) relies solely on provided tweets and pre-existing world knowledge, preventing access to ‘future’ information. As our pre-pandemic collection concludes

in December 2020, we assembled this new tweets dataset after the LLM’s knowledge cutoff date.

Tweets were collected monthly using Twitter Advanced Search, restricted to users with at least 1,000 followers. We focused on keywords related to business and economics topics, resulting in a diverse collection of 2,881 tweets<sup>3</sup>, 90-130 per month, tweeted by 1,255 users, including politicians, CEOs, activists, and academics. Duplicate and URL removal was applied.

### 2.2 Macroeconomic Indicators

We focus on predicting three key macroeconomic indicators:

**Federal Funds Rate (FFR):** The interest rate at which depository institutions, such as banks, lend reserve balances overnight to meet reserve requirements. The FFR serves as a Federal Reserve monetary policy tool, is influenced by public perception of economic stability, and its fluctuations impact various sectors, making it widely monitored.

**S&P 500:** A stock market index measuring the performance of the 500 largest publicly traded companies in the U.S. It reflects collective investor confidence in economic growth and risk appetite and is widely regarded as a barometer of the overall health of the US stock market.

**CBOE Volatility Index (VIX):** Measures market expectations of future volatility based on S&P 500 options prices, often referred to as the ‘fear gauge’ as it tends to rise during market stress and fall during market stability.

These indicators are well-suited for evaluating the narratives’ predictive capabilities, both in their daily frequency which aligns with Twitter tempo, and their sensitivity to public opinions and behaviours. Comparably, other widely used macroeconomic indicators have a lower frequency, such as GDP or inflation (Gentzkow et al., 2019; Kalamara et al., 2022; Ellingsen et al., 2022), or are microeconomic indicators (Rahimikia et al., 2021; Yang et al., 2023b; Khedr et al., 2021; He and Gu, 2021).

Detailed information on the financial indicators and pre-processing steps are in App. B.

Appendix C summarizes all datasets and their characteristics.

<sup>3</sup>As this data collection method is more restricted than the previous, the resulting dataset is relatively smaller.

### 3 Experimental Setup

To assess the central question, whether economic narratives can provide valuable insights for future financial movements, we design a series of prediction tasks, each aims to predict future values of one of the financial indicators as a target: FFR, S&P 500, and VIX (described in § 2.2). In this section we present the tasks and the evaluation methods.

**Prediction Tasks:** We assess the predictive power of narratives across three **prediction horizons**: next-day, next-week, and next-month. Given input data covering dates  $T_1, \dots, T_t$ , the model predicts values for either  $T_{t+1}$ ,  $T_{t+7}$  or  $T_{t+30}$ . Each model variant predicts a single financial indicator at a single horizon.

**Next value:** predicts the target’s value at the specified horizon.

**Percentage change:** predicts the target’s percentage change between the specified horizon and the day before.

**Direction change:** classifies the target’s direction of change (increase or decrease) between the specified horizon and the day before.

These tasks are commonly used in macroeconomic research (Handlan, 2020; Gentzkow et al., 2019; Kalamara et al., 2022; Ahrens and McMahon, 2021; Masciandaro et al., 2021; Lee et al., 2009; Hamilton and Jorda, 2002; Kim et al., 2023; Larkin and Ryan, 2008).

**Evaluation:** We evaluate our models using Mean Squared Error (MSE) for the regression tasks, and Accuracy and  $F_1$ -Score for classification. Additionally, we compare our models against two types of baselines described below. In all experiments besides § 7.1, where the model is not being trained, we utilize an 80-20 train-test split, adhering to chronological order to preserve temporal context.

#### Financial Baselines:

Baseline	Description	Type
As-previous	Next value is the same as prev.	C/R
Inverse-previous	Next value is inverse of prev.	C
Week-majority	Majority vote of previous week	C
Train-majority	Majority vote of training data	C
Up-predictor	Always predicts an "increase"	C
Down-predictor	Always predicts a "decrease"	C
Train-mean	Mean value of training data	R

Table 1: Financial baselines. C: Classification, R: Regression

#### Counterfactual Textual Baselines:

**Random texts:** To evaluate whether the LLM actually utilizes the content of tweets, we establish a baseline which feeds it with randomly generated sentences comprised of varying random words.

**Shuffled tweets:** We assess model reliance on temporal narratives by feeding the LLM with chronologically disordered tweets, to isolate the impact of temporal narratives from confounding patterns or memorization. A well-functioning model should outperform this baseline, indicating its reliance on relevant temporal narratives for prediction.

**Synthetic ‘narratives’:** We generate synthetic narrative-like sentences expressing positive or negative cues, aligned with subsequent changes in the financial indicator. This allows us to assess the model’s ability to infer relationships between aligned narratives and the following market changes.

### 4 Models

Our models are categorized into 3 categories, based on the signals they use as input:

**Financial (F):** utilizes historical financial data, from the past week or month.

**Textual (T):** leverages solely textual data, either raw tweets or tweets’ analyses.

**Textual & Financial (TF):** draws upon both textual and financial data as input. We aim to effectively utilize both insights from textual narratives and historical financial patterns for enhanced prediction accuracy. Outperforming an  $F$  model with a  $TF$  model can demonstrate the added value of textual narratives in enhancing prediction capabilities.

Our model selection progresses from simpler models, commonly used in the financial literature (see §8), to more advanced architectures. This progression serves two purposes: (1) Achieving positive results with simpler models provides a stronger indication for the predictive signal of narratives; and (2) It allows us to build upon existing research in Narrative Economics, primarily rooted in finance and often utilizes relatively simple models, before exploring more advanced NLP approaches.

#### 4.1 Financial Models

The financial models include traditional *machine learning (ML) models* (Linear Regression, Lasso, Ridge, SVM, Random Forest, and Logistic Regression), *DA-RNN* (Qin et al., 2017), and *T5* (Raffel et al., 2020). Each model processes a sequence

of historical financial values, either as individual features or as a time-series. T5 receives these inputs in textual form ('increase' and 'decrease') for 'direction change' classification.

## 4.2 Textual Models

We employ increasingly complex text representations, serving as features for the remaining of the prediction model, as illustrated in Figure 1.

**Daily Sentiments:** Motivated by literature suggesting that aggregated sentiments capture the core of textual narratives and enhance economic prediction (Macaulay and Song, 2023; Yang et al., 2023b; Adams et al., 2023; Kim et al., 2023; Gurgul et al., 2023), we represent each tweet with its sentiment score utilizing VADER (Hutto and Gilbert, 2014), given its extensive adoption in the financial literature (Kalamara et al., 2022; Khedr et al., 2021; Kim et al., 2023), providing a basis for comparison with existing research. Daily sentiment scores are then computed by averaging individual tweets scores within each day, and concatenated over a week.

**Embedding-Based Representations:** We utilize pre-trained language models (BERT (Devlin et al., 2018), RoBERTa (Liu et al., 2019), and T5 (Raffel et al., 2020)) to derive embeddings for individual or concatenated tweets. *Individual tweet* embeddings (CLS token or averaged word embeddings) are aggregated daily by averaging or concatenating embeddings of same-date tweets. *Joint tweets* embeddings encode together multiple concatenated tweets from the same date, potentially capturing their collective meaning without explicit aggregation, avoiding potential information loss.

**LLM-Generated Analyses (hybrid model):** We create concise monthly analyses by feeding a month of tweets from the Post-2021 Twitter dataset (§2.1) and corresponding financial values of the target indicator to OpenAI's Chat Completion API, GPT-3.5 (OpenAI, 2021) (see App. D for LLM details and a prompt example). These analyses are either used directly for prediction (§7.1) or as an input to a subsequent T5 model (§7.2). Notably, the LLM receives both textual and financial inputs to enable analyzing relationships. Thus, this method applies only for the *TF* model type.

## 4.3 Integrating Textual and Financial Models

Having established dedicated models for processing *T* and *F* data, we now address the strategies for combining these representations in *TF* models to produce unified predictions:

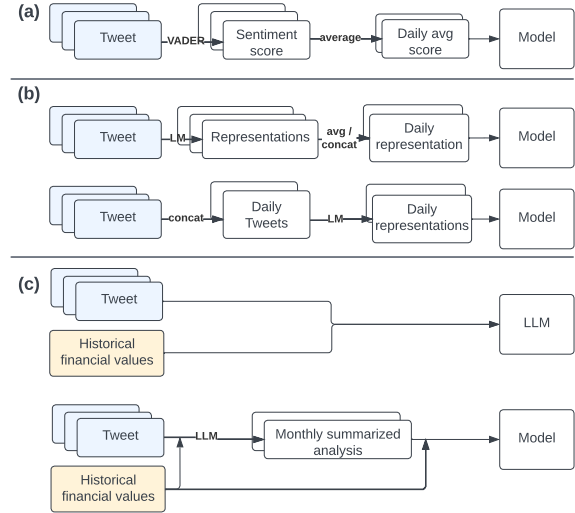


Figure 1: Textual models' pipelines to represent textual data as part of the prediction pipeline. (a) Daily sentiments. (b) Individual and joint tweets LLM's representations. (c) LLM analysis for prediction and as input to a subsequent prediction model.

**Concatenation:** The simplest approach is concatenating the *T* and *F* representations.

**DA-RNN (Qin et al., 2017):** The dual-stage attention-based RNN model, which was used in related work (Wu et al., 2018; Yu and Wu, 2019; Chung and Park, 2021), predicts the current value of a time-series based on its previous values and those of exogenous series. We feed historical financial representations (*F*) as the time series and textual representations (*T*) as the exogenous series.

**Prompt-based integration:** Both LLM-generated analyses (*T+F*) and raw financial values (*F*) are provided to a T5 model (Raffel et al., 2020) as separate segments with leading prompts instructing the model on how to use each data source (see App. E for a prompt example).

Given a *TF* model, we can derive *T* or *F* models by selectively omitting or zeroing either the *F* or *T* features, respectively.

## 5 Analyzing Narratives in Textual Data

Since our datasets serve as the input sources for the downstream prediction models, we first assure the presence of latent narratives within them.

### 5.1 Tweets Narrative Analysis

We analyze the narratives captured in our two Twitter corpora (described in §2.1) using RELATIO (Ash et al., 2021), an algorithm that extracts political and economic narrative-tokens from text (see

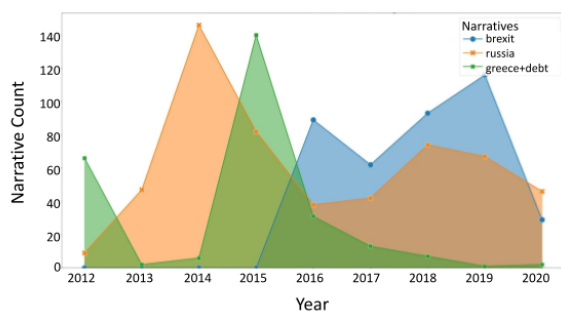


Figure 2: Temporal distribution of top three narratives from the 'economics' dataset, extracted by RELATIO (Ash et al., 2021) (see 8): UK's Brexit, Greece's financial debt, and Russia's financial crises. We can see the evolving nature of these narratives over time, where the distribution is aligned with real-life related events.

§8). Figure 2 shows the temporal distribution of the three most frequent narratives in the economy-topic dataset: UK's Brexit, Greece's financial debt, and Russia's financial crises. It demonstrates the evolving nature of these narratives over time, with peaks aligning with related real-world events such as the debt crises Greece experienced in 2015 and the UK's referendum in June 2016 about leaving the European Union. The presence of these narratives and their temporal dynamics within our dataset confirm the potential of social media data for economic understanding and forecasting. Moreover, the rapid shifts in word frequency following significant events suggest that economically relevant events are swiftly reflected in our data, making it a valuable source for short-term economic prediction. For further analyses, see App. F.1.

## 5.2 LLM-Based Narrative Analysis

Having demonstrated the presence of narratives within our raw data, we now analyze the LLM's ability to extract and summarize them. Following the methodology in § 4.2, we generated 697 LLM-based analyses, one for each date from September 2021 to July 2023. The outputs generated by the LLM can be divided into two main components: the analysis of tweets and the potential impact of them on the financial target. In each component, the analyses contain ~9 sentences on average.

A snippet of LLM-generated analysis is presented in Figure 3 (see full examples in App. F.2). This snippet illustrates the LLM's information aggregation, summarizing and distinguishing between opinions and events expressed in the tweets. Furthermore, the LLM connects the insights it found to potential future impacts on the financial

indicator, a crucial first step towards prediction.

### Analysis of Tweets:

...Some tweets express concerns about inflation, rising interest rates, and the impact on the economy and personal finances...Several tweets highlight the impact of government policies on various sectors, such as health-care, student loans, and housing. Some tweets express skepticism towards central banks and their role in the economy...A few tweets discuss the impact of global events, such as the Russian mobilization...

### Potential Effect on S&P 500:

...Concerns about inflation, rising interest rates, and economic instability expressed in the tweets may lead to increased market volatility and potential declines in the S&P 500...Global events mentioned in the tweets...may have indirect effects on the S&P 500 through their impact on global markets and investor sentiment.

Figure 3: Snippet of LLM-based analysis for 29/08/2022 to 28/09/2022. In this time period the Federal Reserve raised the interest rates in an effort to combat inflation, the US Supreme Court ruled that the Biden administration could not extend the pause on student loan payments, and more. See full analysis in App. F.2

## 6 Predicting using LLM representations

After analyzing our textual data and demonstrating the existence of narratives along with their temporal shifts aligned with world events, this section delves into employing LLMs to represent these narratives in order to use them for financial forecasting.

### 6.1 Sentiment-Based Next-Day Prediction

**Method.** We use daily sentiment from the past week as features (see § 4.2) for FFR next-day prediction. Sentiments are fed as features into several classical ML models, described in § 2.2, applied to 'next value' and 'direction change' tasks. For the 'prediction change' task the financial input is encoded either as numerical features or binary increase/decrease values to align with the label.

**Results.** Table 2 presents the results of the most performant models (full results are in App. G). In classifying 'direction change', models with financial input ( $F$  &  $TF$ ) outperform text-only models ( $T$ ), exhibiting a 5% accuracy improvement (0.94

vs. 0.89). The slight difference between  $F$  and  $TF$  models (0.936, 0.939) suggests text has little impact.<sup>4</sup> The  $T$  models achieve comparable accuracy to the  $F$  baselines (0.89, 0.81). In predicting 'next value', best  $F$ ,  $TF$  and  $T$  models yield comparable MSE to the 'train-mean' baseline (15.4, 15.6).<sup>5</sup>

**Takeaway.** Sentiments lack nuances needed for financial prediction, and classical ML models have limited capabilities. Several models failed to surpass the performance of the non-learned baselines, indicating the necessity for improved text representation and more advanced prediction models.

Type	Model	Accuracy
F baselines	As-previous	0.812
F	Random Forest Numeric	<b>0.936</b>
TF	Random Forest Numeric	<b>0.939</b>
T	Logistic Regression	0.885
T	SVM	0.885

Type	Model	MSE
F baseline	Train-mean	15.661
F	SVM	15.416
TF	SVM	15.416
T	SVM	15.36

Table 2: Results of most performant models for predicting FFR using sentiments. **1.** In 'direction change', financial features are encoded either numerically.  $TF$  and  $F$  perform nearly identically, outperforming  $T$  models and  $F$  baselines. **2.** In 'next value', best models are comparable to each other and to the  $F$  baseline.

## 6.2 Embeddings for Time-Series Prediction

**Method.** To better capture the richness and complexity of information concealed within the tweets, we turn to the two 'Embedding-Based Representations' described in § 4.2. Additionally, we step beyond ML models and instead utilize the time-series deep learning model DA-RNN (Qin et al., 2017), designed to capture temporal dynamics and complex relationships within data (see § 4.3).

Through rigorous evaluation, we explore various model configurations, target variables (FFR and VIX), tasks ('percentage change', 'direction change' and both together), prediction horizons, LLM architectures, aggregation methods, and the daily number of tweets. Lastly, we assess the models' reliance on temporal context and relevant narratives, using the 'random texts' and 'shuffled tweets'

<sup>4</sup>Feature importance confirms this, with the previous day's financial feature dominating in  $TF$  models (not shown).

<sup>5</sup>Feature importance reveals low scores for both sentiments and financial features (not shown).

counterfactual baselines described in § 3. Due to space constraints, we present results for a single task and setting- VIX 'next value' prediction with 1- or 7- days horizons. Additional experiments showed a recurring pattern to the presented results.

**Results.** Figure 4 presents the models' performance. For a 1-day horizon, the 'as-previous'  $F$  baseline outperforms all other models (3.079 MSE). This suggests that the input information might not be beneficial for such a short-term prediction. For a 7-day horizon, both  $TF$  models (13.148, 13.147) initially appear to outperform the  $F$  model (13.463) and  $F$  baseline (16.172), implying a potential influence of the textual content. However, the 'random texts'  $TF$  baseline outperforms (13.056) all other models, indicating that the good performance of the  $TF$  models is not likely due to the content of the tweets. We hypothesize that the presence of text improves performance, even when random, due to either spurious correlations or random noise that helps the model generalize, similarly to regularization techniques. The difficulty in capturing and representing the aggregated tweets information in a way that is meaningful for financial prediction might be a contributing factor. Additionally, challenges may lie in using historical financial data to predict future values of an indicator characterized by frequent random movements and fluctuations.

**Takeaway.** The models struggle to learn from the tweets for the macroeconomic predictions, suggesting that implicitly representing and aggregating latent narratives within LLMs remains challenging.

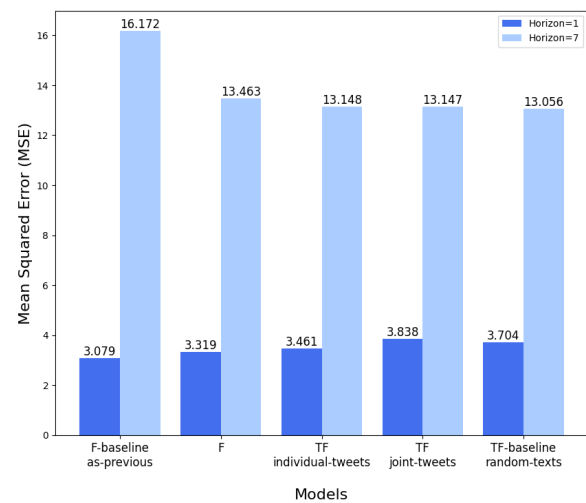


Figure 4: VIX 'next value' prediction for 1/7-days horizons. The  $F$  and  $TF$  baselines outperform all models in 1- and 7-day horizons, respectively, suggesting all models struggle to learn from tweets for the prediction.

## 7 Predicting using LLM Analyses

Our former attempts to predict financial indicators directly from raw tweets proved insufficient, potentially due to the difficulty of implicitly utilizing narratives and aggregating information from diverse tweets. To address this, we apply generative LLMs, providing them tweets and financial data from the same time period, to summarize insights and analyze potential effects on future market movements. We initially explored using the LLM directly as a predictor (§7.1), but due to limited success, we repurposed its analyses as refined inputs for a subsequent prediction model (§7.2).

### 7.1 LLMs for End-to-End prediction

We predict the average weekly VIX (or S&P 500) value based on a monthly window of tweets and corresponding VIX (or S&P 500) values. Utilizing the web chat version of GPT (OpenAI, 2022) for continuous reasoning, we tested few- and zero-shot settings, and multi-step reasoning with Chain-of-Thought (CoT) prompting (Wei et al., 2022). See App. H for full description.

Overall, the LLM consistently produced meaningful narrative analyses and comprehension of financial implications, although not being instructed to do so, but exhibited inconsistencies in integrating these insights for prediction. For example, it occasionally refused to provide predictions due to safety guardrails. In other cases it mirrored input ranges neglecting the potential impact of successfully analyzed narratives. When presented with 'synthetic narratives' (§ 3), it recognized the change direction but struggled to quantify the magnitude of it. Addressing these limitations could unlock the full potential of LLMs for financial forecasting.

### 7.2 Two-Stage prediction with LLM Analyses

**Method.** The previous experiment revealed the LLM's ability to create insightful analyses of tweets-financial data (see §5.2 and F.2 for LLM-outputs analysis). Here we use these analyses as inputs for a dedicated prediction model (see §4.2), to predict the S&P 500 'direction change'. Unlike the embedding-based approach (§ 6.2) which struggled to aggregate diverse narratives, here the LLM produces concise analyses. Compared to the LLM predictor approach (§ 7.1), here we implement a separate fine-tuned model for the downstream prediction task. Additionally, we compare the models to the 'synthetic narratives' *T* baseline (see § 3).

**Results.** Table 3 shows that there is no significant difference between the best *TF* and *F* models, with a performance gap of ~2% on a limited test set of ~90 samples.<sup>6</sup> The McNemar's test (Dror et al., 2018) reveals no statistically significant difference (p-value=0.48). We regard the small gap as a negative result and understand if this causes curiosity of the readers. Notably, in comparison to previous experiments, here the models surpass all baselines.

**Takeaway.** While *TF* and *F* models outperform all others, the gap between them is insignificant.

Type	Model	Accuracy	F <sub>1</sub> -Score
F-baselines	Train-majority	0.424	0.0
	Week-majority	0.484	0.598
	As-previous	0.484	0.552
	Inverse-previous	0.517	0.511
	Up-predictor	0.576	0.731
	Down-predictor	0.424	0.0
F	T5 Base	<b>0.604</b>	0.723
F	T5 Large	0.593	<b>0.727</b>
TF	TF Base	0.626	0.738
TF	T5 Large	<b>0.627</b>	<b>0.742</b>
T	T5 Large	0.587	0.726
T-baseline	Synthetic narratives	0.489	0.254

Table 3: S&P 500 'direction change'. *TF* and *F* models outperform all other models, with an insignificant gap of ~2% favoring the *TF* model.

## 8 Related Work

**Financial Market Prediction.** Financial prediction is a longstanding research area, with approaches ranging from classical quantitative methods (Arthur, 1995; Andersen et al., 1999; Wolfers and Zitzewitz, 2004; Hamilton and Jorda, 2002; Athey and Imbens, 2019) to traditional ML models using historical financial data (Kalamara et al., 2022; Gentzkow et al., 2019; Masciandaro et al., 2021). However, these methods often struggle to capture market complexities due to their limitations in modeling non-linear relationships.

Others leverage RNNs, CNNs, and attention mechanisms (Yang et al., 2023b; Handlan, 2020; Lee et al., 2009; He and Gu, 2021; Qin et al., 2017). Very large LLMs have recently started blooming in this domain, for learning from financial texts and time-series data (Garza and Mergenthaler-Canseco, 2023; Wu et al., 2023; Yang et al., 2023a; Yu et al., 2023; Xie et al., 2023) and for stock price pre-

<sup>6</sup>As a reminder, we can only use the second Twitter dataset, of tweets that were posted after the LLM's training cutoff date, and our financial indicators are of daily frequency, therefore the small dataset for this type of experiments.

diction (Swamy et al., 2023; Chen et al., 2023; Lopez-Lira and Tang, 2023; Koa et al., 2024). Yet, their often closed-source nature and tendency to hallucinate limit their application (Bybee, 2023).

Features used for prediction include a wide range of financial information, from stocks and macroeconomic indicators to news, Fed announcements and companies reports (Ahrens et al., 2023; Larkin and Ryan, 2008; Yang et al., 2023b; Rahimikia et al., 2021; Kim et al., 2023; Gentzkow et al., 2019).

**Social Media and Finance.** Social media provides a wealth of user-generated content for analyzing investors opinions and market dynamics. Researchers commonly collect data by filtering hashtags or accounts, often preprocessing it by removing emojis and links (Vamossy and Skog, 2021; Chandra and Jana, 2020; Mengistie and Kumar, 2021; Sethi et al., 2020; Karami et al., 2020).

Studies have explored various text representations, from closed forms such as sentiments, emotions and topics (Wang et al., 2023; Mengistie and Kumar, 2021; Wazery et al., 2018; Khan and Malviya, 2020; Kaur and Ranjan, 2020; Sethi et al., 2020; Soumya and Pramod, 2020), to embeddings from pre-trained LLMs (Sethi et al., 2020; Soumya and Pramod, 2020; Kalamara et al., 2022; Gentzkow et al., 2019; Chanda, 2021; Ye et al., 2020) or specific-domain fine-tuned LLMs (Kim et al., 2023; Chu et al., 2023; Gurgul et al., 2023).

Common approaches to aggregating these representations, such as averaging, extracting topics, or creating indices (Kalamara et al., 2022; Ellingsen et al., 2022; Mezo and Ferrari, 2021; Adams et al., 2023; Kim et al., 2023; Khedr et al., 2021) often struggle to capture the diversity of opinions, topics and narratives. We propose generating more holistic LLM analyses, aggregating common narratives while distinguishing differing viewpoints.

Social media data has been used to learn relationships with financial markets, often finding correlations and similarities (Chandra and Jana, 2020; Karami et al., 2020; Nyman et al., 2021; Gholampour and van Wincoop, 2019; Ahrens and McMahon, 2021; Masciandaro et al., 2021; Macaulay and Song, 2023). While this data shows promise for financial prediction, existing work often relies on simplified representations, focuses on microeconomic variables or events-predicting (Gurgul et al., 2023; Adams et al., 2023; Wang et al., 2023), while macroeconomic forecasting is underexplored.

**Learning Narratives.** Emerging research explores Narrative Economics (Ellingsen et al., 2022;

Nyman et al., 2021; Ahrens and McMahon, 2021; Macaulay and Song, 2023; Feder et al., 2022). For example, RELATIO (Ash et al., 2021) is an algorithm extracting political and economic narratives from texts by identifying entity groups and mapping their relationships. While useful for analysis, RELATIO outputs discrete tokens, making them less suitable for our LLM-based prediction tasks.

**In summary,** while significant progress has been made in leveraging social media data and NLP models for economic applications, the impact of narratives on macroeconomic movements remains unclear. This work builds upon existing research, but distinguishes itself by combining advanced NLP, for extracting and aggregating a broad scope of economy-related narratives, and integrating them with financial data to assess their potential to improve short-term macroeconomic predictions, an intersection not yet addressed in the literature.

## 9 Discussion and implications of Results

This research explores the Narrative Economics hypothesis, extracting and analyzing economy-related narratives from social media and testing their utility for downstream macroeconomic prediction, with tweet representations serving as input.

While demonstrating the presence of narratives within our curated datasets and establishing NLP building blocks for narrative extraction, evaluating their macroeconomic impact remains a challenge. Our models incorporating narrative data showed limited improvement over those using solely financial data. They failed to consistently outperform our baselines or financial models, and any observed improvements were marginal and statistically insignificant and we regard it as a negative result.

Addressing the question possessed in the title—**Can LLMs Learn Macroeconomic Narratives from Social Media?**—our findings suggest so, as demonstrated by our analyses (§5). Yet, the practical utility of this learned knowledge for economic applications remains an open question. This gap between successful extraction of narratives and limited improvement in macroeconomic prediction raises a question regarding the extent to which these narratives, on their own, can truly drive and predict economic fluctuations, at least at the macroeconomic level. This study serves as a foundation for further exploration, highlighting the need for new macroeconomic models or tasks designed to assess the extracted narratives’ influence on the economy.



## 10 Limitations

This research has potential limitations while facing several challenges. First, focusing on short-term prediction horizons (nowcasting) presents a significant challenge due to the inherent complexity and randomness of economic markets. The Efficient Market Hypothesis suggests that the predictive power of nowcasting is limited, as public information is instantly reflected in asset prices. However, Narrative Economics proposes that narratives can affect peoples' decisions and therefore help us predict and understand economic fluctuations (Shiller, 2017).

We focused on a limited number of economic targets (FFR, S&P 500, and VIX) influenced and shifted by diverse, external and unobserved sources. Utilizing other financial targets or other tasks, such as anomaly detection or profit prediction, might have lead to stronger evidences of the impact of narratives on the economy.

Identifying the temporal lag between the emergence of a narrative and its potential influence on financial targets presents another hurdle. Although being comprehensive, our experiments only examined limited lags and prediction horizons.

While our datasets were carefully curated to capture potential narratives, identifying them definitively is challenging, especially when aggregating multiple narratives for a holistic economic picture. The definition of 'narrative' is broad and subjective, and narratives are typically only recognizable as such in retrospect. Combined with the inherent noise and susceptible to misinformation in social media, extracting reliable narratives with certainty is a complex task.

Lastly, we are limited to publicly accessible LLMs with a known cutoff date, to avoid potential 'future' world knowledge. Utilizing other models might lead to better results.

It is important to acknowledge that this research deals with predicting human and economic behavior, which carries potential risks of misuse. This technology could be applied in harmful or unfair ways, and therefore should be developed and used with caution and awareness of its ethical implications.

## References

Travis Adams, Andrea Ajello, Diego Silva, and Francisco Vazquez-Grande. 2023. More than words: Twit-

ter chatter and financial market sentiment. *arXiv preprint arXiv:2305.16164*.

Maximilian Ahrens, Deniz Erdemlioglu, Michael McMahon, Christopher J Neely, and Xiye Yang. 2023. Mind your language: Market responses to central bank speeches. *Available at SSRN 4471242*.

Maximilian Ahrens and Michael McMahon. 2021. Extracting economic signals from central bank speeches. In *Proceedings of the Third Workshop on Economics and Natural Language Processing*. Association for Computational Linguistics.

Torben G Andersen, Tim Bollerslev, and Steve Lange. 1999. Forecasting financial market volatility: Sample frequency vis-a-vis forecast horizon. *Journal of empirical finance*, 6(5):457–477.

Joshua D Angrist and Jörn-Steffen Pischke. 2009. *Mostly harmless econometrics: An empiricist's companion*. Princeton university press.

W Brian Arthur. 1995. Complexity in economic and financial markets. *Complexity*, 1(1):20–25.

Elliott Ash, Germain Gauthier, and Philine Widmer. 2021. Relatio: Text semantics capture political and economic narratives. *arXiv preprint arXiv:2108.01720*.

Susan Athey and Guido W Imbens. 2019. Machine learning methods that economists should know about. *Annual Review of Economics*, 11:685–725.

Leland Bybee. 2023. Surveying generative ai's economic expectations. *arXiv preprint arXiv:2305.02823*.

Ashis Kumar Chanda. 2021. Efficacy of bert embeddings on predicting disaster from twitter data. *arXiv preprint arXiv:2108.10698*.

Yogesh Chandra and Antoreep Jana. 2020. Sentiment analysis using machine learning and deep learning. In *2020 7th international conference on computing for sustainable global development (INDIACom)*, pages 1–4. IEEE.

Zihan Chen, Lei Nico Zheng, Cheng Lu, Jialu Yuan, and Di Zhu. 2023. Chatgpt informed graph neural network for stock movement prediction. *arXiv preprint arXiv:2306.03763*.

Eric Chu, Jacob Andreas, Stephen Ansolabehere, and Deb Roy. 2023. Language models trained on media diets can predict public opinion. *arXiv preprint arXiv:2303.16779*.

Chaeshick Chung and Sukjin Park. 2021. Dual-stage attention-based recurrent neural networks for market microstructure. In *2021 IEEE Asia-Pacific Conference on Computer Science and Data Engineering (CSDE)*, pages 1–6.

740	Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. <i>arXiv preprint arXiv:1810.04805</i> .	Amir Karami, Vishal Shah, Reza Vaezi, and Amit Bansal. 2020. Twitter speaks: A case of national disaster situational awareness. <i>Journal of Information Science</i> , 46(3):313–324.	795
741			796
742			797
743			798
744	Rotem Dror, Gili Baumer, Segev Shlomov, and Roi Reichart. 2018. <i>The hitchhiker’s guide to testing statistical significance in natural language processing</i> . In <i>Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)</i> , pages 1383–1392. Association for Computational Linguistics.	Rupinder Kaur and Sandeep Ranjan. 2020. Sentiment analysis of 21 dayscovid-19 indian lockdown tweets. <i>International Journal of Advance Research in Science and Engineering</i> , 9(7):37–44.	799
745			800
746			801
747			802
748			
749		Mudassir Khan and Aadarsh Malviya. 2020. Big data approach for sentiment analysis of twitter data using hadoop framework and deep learning. In <i>2020 International Conference on Emerging Trends in Information Technology and Engineering (ic-ETITE)</i> , pages 1–5. IEEE.	803
750			804
751	Jon Ellingsen, Vegard H Larsen, and Leif Anders Thorsrud. 2022. News media versus fred-md for macroeconomic forecasting. <i>Journal of Applied Econometrics</i> , 37(1):63–81.		805
752			806
753			807
754			808
755	Amir Feder, Katherine A Keith, Emaad Manzoor, Reid Pryzant, Dhanya Sridhar, Zach Wood-Doughty, Jacob Eisenstein, Justin Grimmer, Roi Reichart, Margaret E Roberts, et al. 2022. Causal inference in natural language processing: Estimation, prediction, interpretation and beyond. <i>Transactions of the Association for Computational Linguistics</i> , 10:1138–1158.	Ahmed M Khedr, Ifra Arif, Magdi El-Bannany, Saadat M Alhashmi, and Meenu Sreedharan. 2021. Cryptocurrency price prediction using traditional statistical and machine-learning techniques: A survey. <i>Intelligent Systems in Accounting, Finance and Management</i> , 28(1):3–34.	809
756			810
757			811
758			812
759			813
760			814
761		Jihwan Kim, Hui-Sang Kim, and Sun-Yong Choi. 2023. Forecasting the s&p 500 index using mathematical-based sentiment analysis and deep learning models: a finbert transformer model and lstm. <i>Axioms</i> , 12(9):835.	815
762	Fred. Daily Economic Datasets. <a href="https://fred.stlouisfed.org/tags/series?t=daily/">https://fred.stlouisfed.org/tags/series?t=daily/</a> .		816
763			817
764	Azul Garza and Max Mergenthaler-Canseco. 2023. Timegpt-1. <i>arXiv preprint arXiv:2310.03589</i> .		818
765			819
766	Matthew Gentzkow, Bryan Kelly, and Matt Taddy. 2019. Text as data. <i>Journal of Economic Literature</i> , 57(3):535–74.	Kelvin JL Koa, Yunshan Ma, Ritchie Ng, and Tat-Seng Chua. 2024. Learning to generate explainable stock predictions using self-reflective large language models. In <i>Proceedings of the ACM on Web Conference 2024</i> , pages 4304–4315.	820
767			821
768			822
769	Vahid Gholampour and Eric van Wincoop. 2019. Exchange rate disconnect and private information: What can we learn from euro-dollar tweets? <i>Journal of international Economics</i> , 119:111–132.		823
770			824
771		Fiacc Larkin and Conor Ryan. 2008. Good news: Using news feeds with genetic programming to predict stock prices. In <i>Genetic Programming: 11th European Conference, EuroGP 2008, Naples, Italy, March 26-28, 2008. Proceedings 11</i> , pages 49–60. Springer.	825
772			826
773	Vincent Gurgul, Stefan Lessmann, and Wolfgang Karl Härdle. 2023. Forecasting cryptocurrency prices using deep learning: Integrating financial, blockchain, and text data. <i>arXiv preprint arXiv:2311.14759</i> .		827
774			828
775			829
776		Chien-Cheng Lee, Chun-Li Tsai, and Yu-Chun Chiang. 2009. Federal funds rate prediction: A comparison between the robust rbf neural network and economic models. <i>Journal of Information Science &amp; Engineering</i> , 25(3).	830
777	James D Hamilton and Oscar Jorda. 2002. A model of the federal funds rate target. <i>Journal of Political Economy</i> , 110(5):1135–1167.		831
778			832
779			833
780	Amy Handlan. 2020. Text shocks and monetary surprises: Text analysis of fomc statements with machine learning. <i>Published Manuscript</i> .	Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. <i>arXiv preprint arXiv:1907.11692</i> .	834
781			835
782			836
783	Shwai He and Shi Gu. 2021. Multi-modal attention network for stock movements prediction. <i>arXiv preprint arXiv:2112.13593</i> .		837
784			838
785			839
786	Clayton Hutto and Eric Gilbert. 2014. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In <i>Proceedings of the international AAAI conference on web and social media</i> , volume 8, pages 216–225.	Alejandro Lopez-Lira and Yuehua Tang. 2023. Can chatgpt forecast stock price movements? return predictability and large language models. <i>arXiv preprint arXiv:2304.07619</i> .	840
787			841
788			842
789			843
790		Robert E Lucas and Thomas J Sargent. 1981. <i>Rational expectations and econometric practice</i> , volume 2. U of Minnesota Press.	844
791	Eleni Kalamara, Arthur Turrell, Chris Redl, George Kapetanios, and Sujit Kapadia. 2022. Making text count: economic forecasting using newspaper text. <i>Journal of Applied Econometrics</i> , 37(5):896–919.		845
792			846
793		Alistair Macaulay and Wenting Song. 2023. Narrative-driven fluctuations in sentiment: Evidence linking traditional and social media. Technical report, Bank of Canada.	847
794			848
			849
			850

851	Donato Masciandaro, Davide Romelli, Gaia Rubera, et al. 2021. <i>Monetary policy, Twitter and financial markets: evidence from social media traffic</i> . BAFFI CAREFIN, Centre for Applied Research on International Markets Banking . . . .	904
852		905
853		906
854		907
855		908
856	Tajebe Tsega Mengistie and Deepak Kumar. 2021. Deep learning based sentiment analysis on covid-19 public reviews. In <i>2021 International Conference on Artificial Intelligence in Information and Communication (ICAIIIC)</i> , pages 444–449. IEEE.	909
857		910
858		911
859		912
860		913
861	Helena Le Mezo and Massimo Minesso Ferrari. 2021. Text-based recession probabilities.	914
862		915
863	Rickard Nyman, Sujit Kapadia, and David Tuckett. 2021. News and narratives in financial systems: exploiting big data for systemic risk assessment. <i>Journal of Economic Dynamics and Control</i> , 127:104119.	916
864		917
865		918
866		
867	OpenAI. 2021. ChatGPT-3.5 API. <a href="https://platform.openai.com/docs/models/gpt-3-5">https://platform.openai.com/docs/models/gpt-3-5</a> . Sep 2021 version.	919
868		920
869		921
870	OpenAI. 2022. ChatGPT. <a href="https://chat.openai.com/chat">https://chat.openai.com/chat</a> . Jan 2022 version.	922
871		923
872	Yao Qin, Dongjin Song, Haifeng Chen, Wei Cheng, Guofei Jiang, and Garrison Cottrell. 2017. A dual-stage attention-based recurrent neural network for time series prediction. <i>arXiv preprint arXiv:1704.02971</i> .	924
873		925
874		926
875		
876		
877	Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. <i>The Journal of Machine Learning Research</i> , 21(1):5485–5551.	927
878		928
879		929
880		930
881		931
882		
883	Eghbal Rahimikia, Stefan Zohren, and Ser-Huang Poon. 2021. Realised volatility forecasting: Machine learning via financial word embedding. <i>arXiv preprint arXiv:2108.00480</i> .	932
884		933
885		934
886		935
887	Manoj Sethi, Sarthak Pandey, Prashant Trar, and Prateek Soni. 2020. Sentiment identification in covid-19 specific tweets. In <i>2020 International Conference on Electronics and Sustainable Communication Systems (ICESC)</i> , pages 509–516. IEEE.	936
888		937
889		
890		
891		
892	Robert J Shiller. 2017. Narrative economics. <i>American economic review</i> , 107(4):967–1004.	938
893		939
894	S Soumya and KV Pramod. 2020. Sentiment analysis of malayalam tweets using machine learning techniques. <i>ICT Express</i> , 6(4):300–305.	940
895		941
896		
897	Manaswini Swamy, Arunima Shukla, and James Purtilo. 2023. Llm-based stock market trend prediction.	942
898		943
899	Twitter. Twitter API. <a href="https://developer.twitter.com/en/docs/twitter-api">https://developer.twitter.com/en/docs/twitter-api</a> .	944
900		945
901	Domonkos Vamossy and Rolf Skog. 2021. Emtract: Investor emotions and market behavior. <i>arXiv preprint arXiv:2112.03868</i> .	946
902		947
903		948
		949
		950
		951
		952
		953
		954
		955
	Qiping Wang, Tingxuan Su, Raymond Yiu Keung Lau, and Haoran Xie. 2023. Deepemotionnet: Emotion mining for corporate performance analysis and prediction. <i>Information Processing &amp; Management</i> , 60(3):103151.	
	Yaser Maher Wazery, Hager Saleh Mohammed, and Essam Halim Houssein. 2018. Twitter sentiment analysis using deep neural network. In <i>2018 14th international computer engineering conference (ICENCO)</i> , pages 177–182. IEEE.	
	Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. 2022. Chain-of-thought prompting elicits reasoning in large language models. <i>Advances in neural information processing systems</i> , 35:24824–24837.	
	Justin Wolfers and Eric Zitzewitz. 2004. <b>Prediction markets</b> . <i>Journal of Economic Perspectives</i> , 18(2):107–126.	
	Huizhe Wu, Wei Zhang, Weiwei Shen, and Jun Wang. 2018. Hybrid deep sequential modeling for social text-driven stock prediction. In <i>Proceedings of the 27th ACM international conference on information and knowledge management</i> , pages 1627–1630.	
	Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambarur, David Rosenberg, and Gideon Mann. 2023. Bloomberggpt: A large language model for finance. <i>arXiv preprint arXiv:2303.17564</i> .	
	X. X (Twitter) Platform. <a href="https://twitter.com/">https://twitter.com/</a> .	
	Qianqian Xie, Weiguang Han, Xiao Zhang, Yanzhao Lai, Min Peng, Alejandro Lopez-Lira, and Jimin Huang. 2023. Pixiu: A large language model, instruction data and evaluation benchmark for finance. <i>arXiv preprint arXiv:2306.05443</i> .	
	Yahoo Finance. <a href="https://finance.yahoo.com/">https://finance.yahoo.com/</a> .	
	Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. 2023a. Fingpt: Open-source financial large language models. <i>arXiv preprint arXiv:2306.06031</i> .	
	Xin Yang, Metoh Adler Loua, Meijun Wu, Li Huang, and Qiang Gao. 2023b. Multi-granularity stock prediction with sequential three-way decisions. <i>Information Sciences</i> , 621:524–544.	
	Zhihao Ye, Gongyao Jiang, Ye Liu, Zhiyong Li, and Jin Yuan. 2020. Document and word representations generated by graph convolutional network and bert for short text classification. In <i>ECAI 2020</i> , pages 2275–2281. IOS Press.	
	Mu-Hui Yu and Jheng-Long Wu. 2019. <b>Ceam: A novel approach using cycle embeddings with attention mechanism for stock price prediction</b> . In <i>2019 IEEE International Conference on Big Data and Smart Computing (BigComp)</i> , pages 1–4.	

956 Xinli Yu, Zheng Chen, Yuan Ling, Shujing Dong,  
957 Zongyi Liu, and Yanbin Lu. 2023. Temporal data  
958 meets llm-explainable financial time series forecast-  
959 ing. *arXiv preprint arXiv:2306.11025*.

## A Twitter Datasets Details

960

**Twitter topics' keywords:** Table 4 represents the keywords used to collect the Twitter datasets. Each topic has a corresponding list of keywords related to the topic. See Section 2 for further explanation.

961

962

Topic	Key Words
Politics	politics, enough, occupy, coup, demonstration, protest, protesters, corruption, active, embassy, government
Business	retail, business, business owner, business relationship, ecommerce, entrepreneurs, entrepreneurship, CEO, management, invest, founder, innovation, patent, economic, finance, economy, financial, funding, stock
Economy	macroeconomics, yield, inflation, mortgage, recession, debt, interest rate, loan
Current Affairs	human interest, news, all news, breaking news, whats happening, notify, news desk, world news
Natural Disasters	global disaster, natural disaster, world, evacuation, hit, shaking, aftershock, safe, disaster, evacuate, earthquake, tsunami, flooding, volcano, eruption, massive, damage, destroyed
Human Disasters	explosion, terror, attack, horror, crash, shot, shots, shooting, terrible, shocking, police, killed, national military, national security, national terrorism

Table 4: Collected topics and their corresponding keywords.

**Additional Keywords for the Post-2021 Twitter Dataset:** occupy, invest, business, economic, finance, economy, financial, stock, macroeconomics, yield, inflation, mortgage, recession, interest rate, loan, GDP, unemployment, federal reserve, fed, exchange rate, monetary policy, FOMC, central bank.

963

964

965

**Data Pre-Processing:** The tweets underwent minimal pre-processing including emoji-to-text conversion, duplicates removal, links removal, spelling corrections, lemmatization and tokenization.

966

967

## B Financial Indicators Details

968

**Federal Funds Rate (FFR):** The FFR is the interest rate at which depository institutions, such as banks, lend reserve balances to each other overnight to meet reserve requirements. The FFR serves as a critical monetary policy tool employed by the Federal Reserve to influence economic activity by adjusting the target range of the FFR in quarter-point increments following Federal Open Market Committee (FOMC) meetings. Changes in the FFR can have significant impacts on various sectors of the economy, including consumer spending, investments, housing, and financial markets. Changes in the FFR can have significant impacts on various sectors of the economy, including consumer spending, investments, housing, and financial markets, making it a widely monitored indicator of overall economic health. The dataset was downloaded from the publicly available economic data at FRED® website ([Fred](#)). To address historical distribution changes, mainly due to the Great Recession and the COVID-19 pandemic, we employ 'blocking split for time-series data', separating periods with consistent distributions. The longest block (December 2008-January 2015) yields 2207 samples and was used in experiments. Figure 5 shows the blocking splits of the FFR. The second block (December 2008-January 2015) is the longest and was used for experiments.

969

970

971

972

973

974

975

976

977

978

979

980

981

982

**S&P 500:** This stock market index measures the performance of the 500 largest publicly traded companies in the US. It is widely regarded as one of the best indicators of the overall health and performance of the US stock market and, by extension, the broader economy. Changes in the index value can indicate shifts in investors' narratives, confidence in economic growth prospects, risk appetite, and expectations for future economic conditions. This dataset was downloaded from Yahoo Finance® website ([Yahoo Finance](#)). We utilize the closing price of the asset. Additionally, to mitigate the impact of the COVID-19 pandemic, we only utilize post-pandemic data.

983

984

985

986

987

988

989

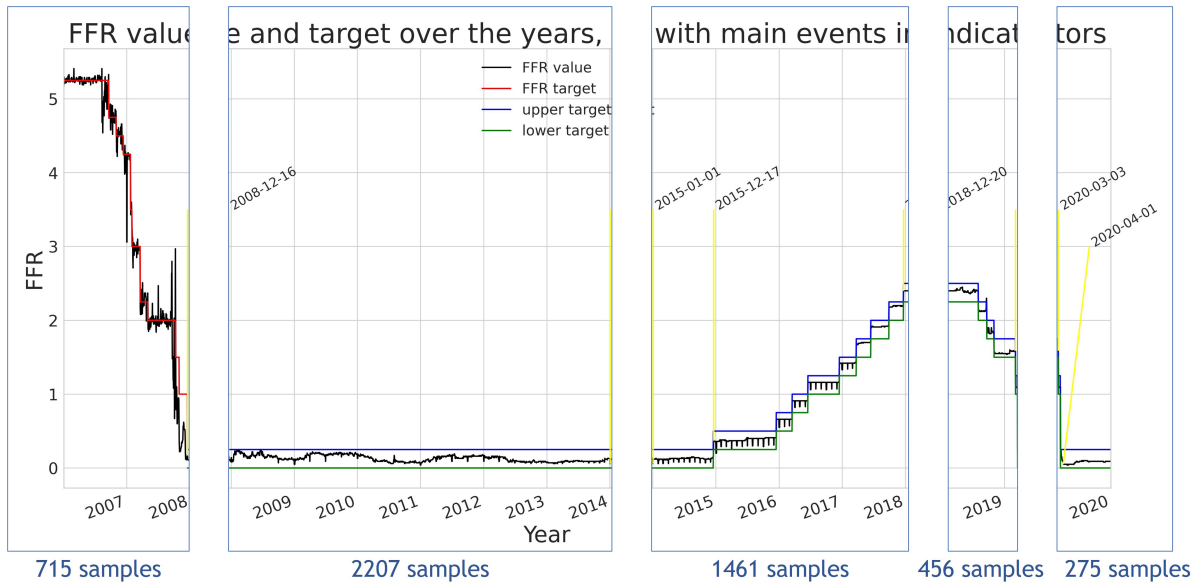


Figure 5: Blocking split of FFR data, isolating time-periods with consistent distributions.

**CBOE (Chicago Board Options Exchange) Volatility Index (VIX):** The VIX measures market expectations of future volatility based on S&P 500 options prices. It reflects investors’ perceptions of the risk and uncertainty in the stock market over the next 30 days. The VIX is known as the ‘fear gauge’ as it tends to rise during times of market stress and uncertainty and fall during periods of market stability. Analyzing changes in the VIX provides insights into investors’ sentiments, risk appetite, and economic confidence. This dataset was downloaded from Yahoo Finance® website (Yahoo Finance). Similar to the S&P 500, we utilize the close price of the index, either pre- or post-pandemic data, avoiding their integration within the same model.

### C Datasets Characteristics

Table 5 summarizes all datasets and their characteristics. Each dataset has a different role in the experiments. The full descriptions of the datasets are presented in Section 2.

Dataset Type	Dataset Name	Time Span	Characteristics
Twitter	Pre-Pandemic Twitter Dataset	2007-2020	Collected daily using 6 topics, minimal followers threshold
Twitter	Post-2021 Twitter Dataset	2021-2023	Collected monthly using a single topic, minimal followers threshold, tweets poster after training cutoff of the used LLM
Financial	Federal Funds Rate	2007-2020 (2008-2015)	Daily macroeconomic indicator, impacts various sectors, widely monitored
Financial	S&P 500	2021-2023	Daily macroeconomic indicator, its fluctuations indicate shifts in investors’ narratives
Financial	Volatility Index	2007-2020 2021-2023	Daily macroeconomic indicator, reflects investors’ uncertainty of the market

Table 5: Summary of datasets with their characteristics.

## D LLM Details

1001

Here we describe implementation details for our LLM-generated analyses pipeline. This Appendix extends the information in Section 4.2. We elaborate on the prompt and parameters, and present a prompt example.

1002

1003

**Prompt:** The prompt of the LLM contains instructions, a set of tweets from the analyzed month, and a set of the financial indicator values (VIX or S&P 500) from the same time period. We feed the LLM with both tweets and indicator values to allow it to potentially understand connections or correlations between them while generating the output. We limit to up to 10 tweets per day to avoid a long prompt. The sets of tweets and financial values are given in a dictionary format keyed by their date: 'yyyy-mm-dd: tweet or indicator value'. Each part of the prompt begins and ends with < > and </ > tags to enable easy post-processing. We instruct the model to generate a summarised analysis of the given input. We specify that the output will be used for predicting future indicator values. We explain about the input data and its structure. Finally, we break the task into two steps and instruct first to analyze the attributes that appear in the tweets and then analyze their potential effects on the close-future indicator values. A full example of a prompt is presented in Figure 6.

1004

1005

1006

1007

1008

1009

1010

1011

1012

1013

1014

**LLM Parameters:** Initial experiments suggest that a temperature of 0.5 balances well between generating diverse and creative outputs while maintaining a factual analysis. All other parameters are the model's default.

1015

1016

1017

### LLM prompt example

1018

```
You are a financial and NLP expert, assisting on creating a summarised
analysis on textual and financial data.
Your task is to create an analysis on given tweets from Twitter and on S&P 500
values from the same time period.
Your output will be used for producing S&P 500 predictions in the close-future
.
Your input 'Financial values' is a dictionary of S&P 500 values with their
corresponding date (yyyy-mm-dd), from a time period of a month.
Your input 'Tweets' is a dictionary of tweets from Twitter with their
publication date (yyyy-mm-dd), from the same time period. The tweets were
posted by opinion leaders and discuss about the news, current affairs,
economy, finance, and politics.
To produce this analysis, first analyse the fear, stability and stress
expressed in the tweets, and then analyse their possible effects on the
close-future S&P 500 value.
Produce your output in the format: <Analysis of Tweets>PLACE HOLDER</Analysis
of Tweets>
<Potential Effects on S&P 500>PLACE HOLDER</Potential Effects on S&P 500>

Input:
<Financial values>
2021-09-01: 4524.09
2021-09-02: 4536.95
...
2021-10-01: 4357.04
</Financial values>

<Tweets>
2021-09-01: It is time for a total economic boycott of Texas and Texas-based
businesses.
...
2021-09-29: Pfizer is the 6th most owned stock by Congress. Surprised?
</Tweets>

Output:
```

Figure 6: LLM prompt for generating the summarized analysis of tweets and financial data of a monthly window.

1019  
1020  
1021  
1022  
1023  
1024

### E T5 Integration Prompt

Figure 7 shows a prompt example for the T5 model classifying 'direction change' of the S&P 500 for 1-day prediction horizon. This model is used for the experiments in Section 7.2. The prompt contains both LLM-generated analysis and historical financial values, separated using leading titles and special tokens. The prompt ends with a title of the task itself. More details about the  $T$  and  $F$  integration can be seen in Section 4.3.

```
[CLS] Summary of recent tweets: The tweets from the given time period cover a wide range of topics including politics, economy, finance, and social issues. There are tweets expressing concerns about the economy, such as discussions on inflation, high energy prices, and the impact of government policies on businesses. There are also tweets discussing the need for financial reforms, including calls for canceling student loan debt and increasing taxes on the wealthy and corporations. Additionally, there are tweets highlighting the consequences of defaulting on debt and the potential impact on the economy. Overall, the tweets reflect a mix of opinions and concerns regarding the current economic and financial landscape. [SEP] Recent S&P 500 directions of change: decrease=0, increase=1, increase=1, increase=1, decrease=0, decrease=0, decrease=0. [EOS] Predict S&P 500 direction of change tomorrow:</s>
```

Figure 7: Prompt example of the T5 prediction model receiving both LLM-generated analysis and historical financial values, in order to classify 'direction change' of the S&P 500.

1025  
1026

### F Additional Narratives Analysis from Text

This section presents complementary analyses to the ones presented in Section 5.

1027

#### F.1 Twitter Analysis

1028  
1029  
1030  
1031  
1032  
1033

To inspect the information contained in our Pre-Pandemic Twitter dataset, we conduct an analysis of the 100 most frequent words associated with each topic of the dataset. This analysis reveals temporal shifts in the distribution of most frequent words, coinciding with significant world events. For instance, Figure 8 illustrates the monthly number of occurrences of the term 'debt' in the 'economics' dataset, presenting peaks in July 2011 and October 2013, corresponding to major debt ceiling crises experienced by the US during these times.

1034  
1035

Figure 9 shows (a) sentiment scores created by VADER (Hutto and Gilbert, 2014) sentiment model (ranged [-1,1]) for the 'business' related tweets; and (b) the sentiment scores divided to bins of size 0.5.

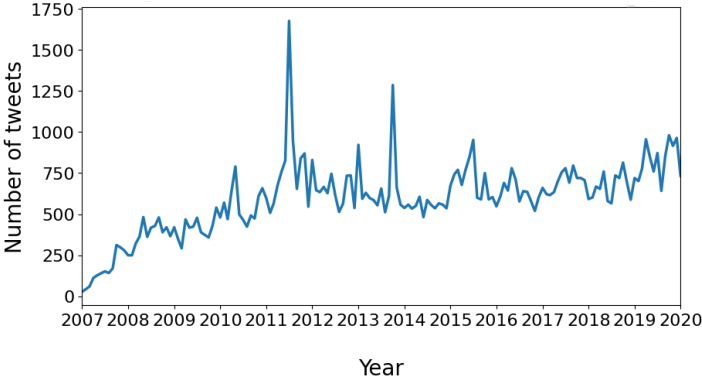
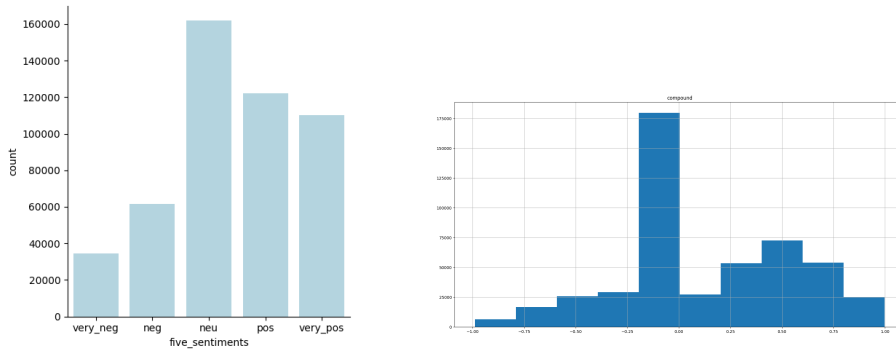


Figure 8: Frequency of the word 'debt' in the 'economics' dataset over the years. Observed peaks in July 2011 and October 2013 are aligned with the major debt ceiling crises the US experienced in these times.





(a) Sentiment scores divided to negative and positive bins.

(b) Raw sentiment scores.

Figure 9: VADER (Hutto and Gilbert, 2014) sentiment scores of the 'business' dataset.

## F.2 LLM-Based Narrative Analyses Examples

1036

Figures 10 and 11 present LLM's outputs examples. Each analysis is constructed of several sentences describing the tweets and financial values given in the prompt. In both examples we can see the analyses recognize the described world events and analyze the impact of them on the financial activity, such as raising concerns and frustration, stress about the state of the economy market, or support in the government actions, which can impact the financial indicator.

1037  
1038  
1039  
1040  
1041

#### Analysis of Tweets:

The tweets from the given time period cover a wide range of topics including Brexit, inflation, debt, unemployment, tax cuts, climate change, and business practices. There is a mix of opinions expressed, with some tweets criticizing government policies and economic decisions, while others highlight the need for investment in working families and fair economic opportunities.

Many tweets express concerns about the impact of Brexit on the UK, citing issues such as gas price hikes, food shortages, export crisis, and chaos in various industries. These tweets indicate a sense of fear and stress regarding the economic consequences of Brexit.

There are also tweets discussing the US economy, particularly the national debt and the potential consequences of defaulting on it. Some tweets criticize Republicans for their stance on the debt ceiling and the potential negative effects on jobs and unemployment. Other tweets highlight the need for tax reform and wealth redistribution to address income inequality.

Climate change and environmental issues are mentioned in several tweets, with some praising China's decision to stop building coal plants overseas and others criticizing the lack of action on climate change by certain politicians and corporations.

In terms of business practices, there are tweets discussing stock buybacks, corporate bailouts, and the influence of pharmaceutical companies in Congress. Some tweets call for increased regulation and transparency in the financial sector.

Overall, the tweets reflect a mix of concerns about the state of the economy, the role of government, and the need for equitable economic policies.

#### Potential Effect on S&P 500:

Based on the analysis of the tweets, it is evident that there are several factors that could potentially impact the S&P 500 in the close future. These factors include:

1. Brexit-related concerns: The tweets expressing fear and stress about the economic consequences of Brexit could lead to increased market volatility and uncertainty, which may negatively impact the S&P 500.
2. Debt ceiling and default risks: The tweets discussing the potential consequences of defaulting on the national debt highlight the risks to the economy and job market. If the debt ceiling issue is not resolved in a timely manner, it could lead to market instability and potentially impact the S&P 500.
3. Climate change and environmental policies: The tweets mentioning climate change and the need for clean energy transition may have implications for certain sectors of the economy, such as fossil fuels and renewable energy. Changes in government policies and regulations related to climate change could impact the performance of companies in these sectors, which in turn could affect the S&P 500.
4. Business practices and regulations: The tweets criticizing stock buybacks, corporate bailouts, and the influence of pharmaceutical companies may contribute to calls for increased regulation and scrutiny of these practices. Any regulatory changes or increased scrutiny on these issues could impact the performance of companies involved, potentially affecting the S&P 500.

It is important to note that the impact of these factors on the S&P 500 will depend on various other economic and market conditions. However, considering the sentiments expressed in the tweets, it is advisable to closely monitor these factors and their potential effects on the S&P 500 in the close future.

Figure 10: LLM-generated analysis for dates 21/09/2021 to 21/10/2021. In this time period the US Labor Department reported that the unemployment rate fell to the lowest level since the start of the COVID-19 pandemic, the US Federal Reserve announced that it would begin tapering its asset purchases, and the US House of Representatives passed the Build Back Better Act (a bill for investments in climate change, healthcare, education and child care) and a week after the US senate failed to pass it.

#### Analysis of Tweets:

The tweets cover a wide range of topics including politics, finance, economy, and current affairs.

There are tweets discussing the actions and statements of political figures such as Marco Rubio, Donald Trump, and Boris Johnson.

Some tweets express concerns about inflation, rising interest rates, and the impact on the economy and personal finances.

There are mentions of financial fraud, insider trading, and the role of big corporations in driving inflation.

Several tweets highlight the impact of government policies on various sectors, such as healthcare, student loans, and housing.

Some tweets express skepticism towards central banks and their role in the economy.

There are mentions of potential economic crises, recessions, and the need for fiscal stimulus.

A few tweets discuss the impact of global events, such as the Russian mobilization and the European Central Bank's digital euro.

Some tweets express frustration with the current state of the economy and the financial hardships faced by individuals.

#### Potential Effect on S&P 500:

- The tweets reflect a mix of positive and negative sentiments towards the economy and financial policies. This sentiment can influence investor confidence and market sentiment, potentially impacting the S&P 500.
- Concerns about inflation, rising interest rates, and economic instability expressed in the tweets may lead to increased market volatility and potential declines in the S&P 500.
- Tweets discussing financial fraud, insider trading, and negative perceptions of big corporations may erode investor trust and confidence, potentially impacting the S&P 500.
- The mention of potential economic crises, recessions, and the need for fiscal stimulus may create uncertainty and caution among investors, potentially affecting the S&P 500.
- Global events mentioned in the tweets, such as the Russian mobilization and the European Central Bank's digital euro, may have indirect effects on the S&P 500 through their impact on global markets and investor sentiment.
- Overall, the sentiment and concerns expressed in the tweets suggest a cautious outlook towards the economy and financial markets, which may influence S&P 500 predictions in the close future.

Figure 11: LLM-generated analysis for dates 29/08/2022 to 28/09/2022. In this time period the Federal Reserve raised the interest rates in an effort to combat inflation, the US Supreme Court ruled that the Biden administration could not extend the pause on student loan payments, the US unemployment rate fell to the lowest level since 1969, and the Biden administration announced an aid package for Ukraine.

## G Full Results for 'Sentiment-Based Next-Day Prediction' Experiment

1042

This appendix shows the full results for the experiment described in Section 6.1. Tables 6 and 7 presents the performance of all compared models for both 'direction change' and 'next value' tasks. For the summary of results and takeaways, please refer to the main experiment section.

1043

1044

1045

## H Full Description for 'LLMs for End-to-End Prediction' Experiment

1046

This appendix is complementary to Section 7.1.

1047

**Method.** Motivated by the limitations observed in prior experiments, we explore the utility of LLMs in predicting the VIX index by utilizing their ability to process long textual sequences and perform multi-step reasoning, obviating the need for separate components for text analysis and financial prediction. We utilize the web chat version of GPT (OpenAI, 2022) for continuous reasoning.

1048

1049

1050

1051

We predict the average VIX value for the following week based on a monthly window of tweets paired with corresponding VIX values. Formally, given data covering dates  $T_1, \dots, T_{30}$ , the LLM is instructed to predict the average value of the week between  $T_{30+7} - T_{30+14}$  (horizon=7). This target mitigates random

1052

1053

1054

Type	Model	Accuracy
F-baselines	As-previous	0.812
	Train-majority	0.885
F	Logistic Regression Binary	0.884
F	Logistic Regression Numeric	0.896
F	SVM Binary	0.883
F	SVM Numeric	0.899
F	Random Forest Binary	0.878
F	Random Forest Numeric	<b>0.936</b>
TF	Logistic Regression Binary	0.884
TF	Logistic Regression Numeric	0.895
TF	SVM Binary	0.883
TF	SVM Numeric	0.899
TF	Random Forest Binary	0.880
TF	Random Forest Numeric	<b>0.939</b>
T	Logistic Regression	0.885
T	SVM	0.885
T	Random Forest	0.879

Table 6: ‘Direction change’ classification of FFR. Financial features are encoded either as ‘Binary’ (increase/decrease) or ‘Numeric’. The TF and *F* models perform nearly identically, outperforming the *T* models and the *F* baselines. Best performing models are *F* & *TF* Random Forest with numeric financial features.

Type	Model	MSE
F-baselines	As-previous	28.712
	Train-mean	<b>15.661</b>
F	Linear Regression	18.842
F	Lasso	18.135
F	Ridge	18.842
F	SVM	<b>15.416</b>
F	Random Forest	37.081
TF	Linear Regression	19.152
TF	Lasso	18.135
TF	Ridge	19.121
TF	SVM	<b>15.416</b>
TF	Random Forest	37.576
T	Linear Regression	15.932
T	Lasso	15.661
T	Ridge	15.898
T	SVM	<b>15.36</b>
T	Random Forest	24.629

Table 7: ‘Next value’ prediction of FFR. Best *T*, *F*, *TF* models are comparable to each other and to ‘train-mean’ *F* baseline.

1055 daily fluctuations common in economic indicators, enabling the model to capture broader trends.

1056 Notably, similar behaviors were seen for S&P 500 prediction, and are omitted due to space constraints.

1057 **1. Few-Shot Learning:** Initially, we condition the LLM with examples of tweet-VIX pairs, instructing  
1058 it to generate similar predictions for unseen inputs.

1059 **Results.** While the model provides descriptive analyses of inputs, although not being instructed to  
1060 do so, it faces challenges in financial prediction. It sometimes refuses to provide predictions, citing its  
1061 limitations as a language model and its inability to provide financial advice or prediction. In other cases  
1062 it predicts a range of values mirroring the input range, neglecting potential effects of current tweets’  
1063 narratives which it successfully analyzes.

1064 **2. Multi-Step CoT Reasoning:** To address the above challenges, we adopt a multi-step approach using  
1065 chain-of-thought (CoT) instructions. The LLM is instructed to first analyzing tweets, then explaining  
1066 their potential influence on the VIX, and finally predicting the average VIX for the following week, while  
1067 providing a rationale for the prediction. This unified approach aims to encourage the model to generate  
1068 analyses directly relevant to the final prediction task.

1069 **Results.** The LLM produces multi-step outputs, demonstrating meaningful narrative extraction and

comprehension of their financial implications. However, prediction performance is inconsistent, with occasionally producing reasonable predictions and explanations, alongside instances where previous mentioned challenges still arise.

*Analyzing Economic Periods:* Since the financial indicator can be volatile during a monthly period, we assess the LLM’s performance during stable (declining VIX), stressed (increasing VIX), trend reversal (declining present and increasing future-surprising peak), and volatile (fluctuating VIX) periods, assessing its prediction abilities in definite periods. We compare predictions with and without tweet data (*F* vs. *TF* models). Figure 12 presents the VIX values during such periods, together with the LLM’s predicted range.

*Stable Period:* Both *F* and *TF* models correctly predict a continued decline, with the *TF* model providing richer analysis.

*Stressed Period:* The LLM fails to capture the upward trend, despite identifying stress in tweets and a current moderate-high VIX values. Revises predictions upon prompting favor the higher end of the current value range.

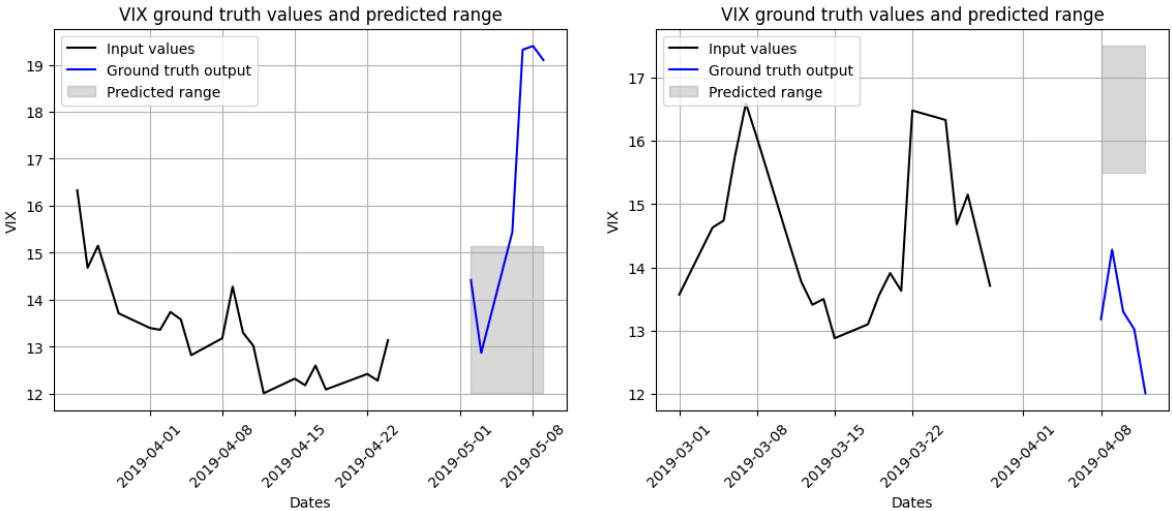
*Trend Reversal:* Despite recognizing the declining trend and increased stress, the LLM predicts a continuation of the declining trend, indicating limitations in incorporating dynamic changes.

*Volatile Period:* Both *F* and *TF* models predict wide ranges, mirroring the input, with the *TF* model offering richer analysis, capturing emotional cues, economic discussions, and political events, all potentially contributing to volatility. While it suggests a slight increase, the wide range limits its precision.

*Multi-Step Few-Shot:* This periods analyses reveal that the textual data seems to enrich the LLM’s understanding, leading to more nuanced explanations but only occasionally offering tailored numerical predictions. To address this, we revisited the few-shot setting, providing the LLM with both illustrative input-output pairs and multi-step instructions. This aims to quantify the magnitude and direction of tweets’ narratives influence on future financial values. Yet, the LLM’s predicted ranges remain incongruent with its tweets’ analysis.

*Synthetic Scenarios:* We manipulate input data to investigate the LLM’s grasp of causal relationships, feeding it with texts coupled with corresponding financial changes (see the ‘synthetic narratives’ baseline in § 3). While the LLM recognizes the direction of impact of these scenarios, it struggles to quantify the actual magnitude of the change.

**Takeaway.** The LLM shows potential for financial forecasting, with its ability to analyze textual data and identify potential economic impacts, offering promising avenues for market insights. Yet, the model faces challenges in prediction reluctance and inconsistency in incorporating textual insights. Addressing these issues can unlock the full potential of LLMs for robust and narratives-insightful financial forecasting.



(a) Trend Reversal Period. (b) Volatile period.  
 Figure 12: VIX economic definite periods, together with the LLM’s predicted ranges.