

# EconNLI: Evaluating Large Language Models on Economics Reasoning

Yue Guo Yi Yang

The Hong Kong University of Science and Technology  
yguoar@connect.ust.hk imyiyang@ust.hk

## Abstract

Large Language Models (LLMs) are widely used for writing economic analysis reports or providing financial advice, but their ability to understand economic knowledge and reason about potential results of specific economic events lacks systematic evaluation. To address this gap, we propose a new dataset, natural language inference on economic events (EconNLI), to evaluate LLMs' knowledge and reasoning abilities in the economic domain. We evaluate LLMs on (1) their ability to correctly classify whether a premise event will cause a hypothesis event and (2) their ability to generate reasonable events resulting from a given premise. Our experiments reveal that LLMs are not sophisticated in economic reasoning and may generate wrong or hallucinated answers. Our study raises awareness of the limitations of using LLMs for critical decision-making involving economic reasoning and analysis. The dataset and codes are available at <https://github.com/Irenehere/EconNLI>.

## 1 Introduction

Economics is a social science that studies the behavior and interactions of economic agents such as individuals, firms, and nations (Krugman and Wells, 2013). It has significant social impact as it covers nearly all vital aspects of human life, including social production, distribution, and consumption (Marshall, 2009; Parkes and Wellman, 2015; Mankiw and Scarth, 2010; Mankiw et al., 2020). Studying economics enables people to understand human society better and forecast future trends.

Recently, with the development of Large Language Models (LLMs), many financial institutions and investment firms are using these models to assist in writing economic analysis reports or providing financial advice (Yang et al., 2023b; Zhao et al., 2024; Lee et al., 2024). For example, an LLM may be used to analyze the impact of changes in interest rates on the stock market or to predict the

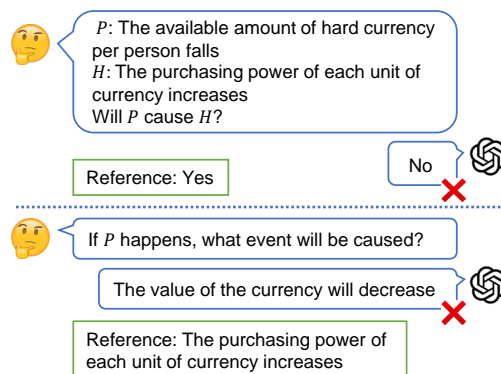


Figure 1: One example from EconNLI. Top: We evaluate whether the LLMs can correctly classify the causal relationship between the premise and the hypothesis. Bottom: We evaluate whether the LLMs can generate a correct consequent event given a premise.

potential effects of a new government policy on a particular industry. These tasks require LLMs to understand economic knowledge and possess the ability to reason and infer the potential results that may arise from certain economic events or actions. However, there lacks a systematic evaluation of such tasks, and how well the LLMs can understand the economic knowledge and reason for the economic events is unknown.

To fill this research gap, we propose a new task, natural language inference on economic events (EconNLI), to evaluate the LLMs' knowledge and reasoning ability on economic events. To facilitate this task, we construct an EconNLI dataset that contains sentence pairs of premise and hypothesis, both of which are economic events. Figure 1 provides an example ChatGPT cannot correctly answer. In this example, the premise is the event that the available amount of hard currency per person falls, and the hypothesis is the purchasing power of each unit of currency increases. Different from the traditional NLI datasets (Dagan et al., 2005; Bowman et al., 2015; Williams et al., 2018), the premise does not entail the hypothesis in semantics; instead,

it requires understanding the economic theory to conduct the inference. Inference on the example in Figure 1 is based on the quantity theory of money. In the case of a decrease in the available amount of hard currency per person, the money supply in the economy will decrease, leading to a decrease in the price level. As a result, the purchasing power of each unit of currency will increase.

We assess the performance of LLMs on EconNLI from two aspects: classification and generation. As shown in Figure 1, in terms of classification, we provide the premise and the hypothesis to the LLMs and evaluate whether the LLMs can correctly classify the causal relationship between the premise and the hypothesis. For generation evaluation, we ask the LLMs to generate the possible resulting events of the given premise and compare the generated events with the hypothesis of the positive example, which serves as a reference.

We conduct extensive experiments on various language models, from open-source to commercial LLMs, and from the general to the finance-specific domain models. For the classification task, we find that without supervised fine-tuning on the training set, the open-source LLMs perform closely to the random guess, and the advanced commercial LLMs such as ChatGPT and GPT-4 also have unsatisfactory performance. As for the generation task, we find that the LLMs can generate the wrong or hallucinated answers. Our findings suggest that LLMs are not sophisticated in economic reasoning, and there remains room for improving the reliability of LLMs in financial and economic analysis.

In conclusion, this paper presents a novel dataset designed to evaluate the reasoning ability of LLMs in the economic domain. The results of our experiments suggest that these models are far from perfect in their ability to economic reasoning. Given that LLMs are now widely applied in writing financial reports or providing financial advice without carefully evaluating their correctness, our study shows LLMs can sometimes lead to errors or unreliable results in economic analysis. Therefore, it is essential to be aware of the limitations of these models when using them for decision-making.

## 2 Related Works

**NLP for Finance and Economics.** The development of LLMs has revolutionized the solutions for many natural language processing (NLP) tasks (OpenAI, 2023; Touvron et al., 2023a,b),

with no exception in the financial domain (Guo et al., 2023a; Lee et al., 2024; Liu et al., 2024). Many domain-specific models have been developed to provide solutions to financial tasks. BloombergGPT (Wu et al., 2023) is the first LLM trained on vast financial data. Later, FinMA (Xie et al., 2023) adapts LLAMA (Touvron et al., 2023a) with instruction-tuning on the various financial NLP tasks. FinGPT (Yang et al., 2023a) is a series of models developed to solve some specific financial tasks, such as sentiment analysis or financial relation extraction. InvestLM (Yang et al., 2023b) is instruction-tuned to provide investment suggestions specifically.

Meanwhile, some works evaluate these LLMs' capability on the financial tasks. Callanan et al. (2023) find the ChatGPT and GPT-4 struggle in answering the mock exam questions of the Chartered Financial Analyst (CFA) Program. Cheng et al. (2023) show GPT-4 performs comparably to human experts as a data analyst. Guo et al. (2023b) shows while some LLMs demonstrate notable performance on some financial NLP tasks, they generally lag behind the fine-tuned expert models. Besides, several datasets are developed to evaluate LLMs in various financial NLP tasks, such as sentiment analysis (Malo et al., 2014; FiQA), text classification (Sinha and Khandait, 2020; Shah et al., 2023) or question answering (Chen et al., 2021, 2022). However, none of them involve assessing the reasoning ability of LLMs in the financial and economic domain. To our knowledge, we are the first to assess LLMs from the aspect of reasoning on the relations between the economic events.

**Natural Language Inference.** Natural language inference (NLI), which aims to reason the entailment relationship between text fragments, has long been a key task in natural language processing (Yu et al., 2023). Various natural language inference tasks are proposed to address the inference ability from different aspects, such as RTE (Dagan et al., 2005), SNLI (Bowman et al., 2015), MultiNLI (Williams et al., 2018), SCITail (Khot et al., 2018), Abductive NLI (Bhagavatula et al., 2019), defeasible NLI (Rudinger et al., 2020), SciNLI (Sadat and Caragea, 2022). However, none of them address the inference in economics and finance. Unlike traditional NLI tasks that rely on linguistic or common sense knowledge for reasoning, EconNLI demands an understanding of economic theories to perform event reasoning.

Wiki Page	Source Sentence	Premise	Hypothesis
Unemployment	For example, minimum wage laws raise the cost of some low-skill laborers above market equilibrium, resulting in increased unemployment as people who wish to work at the going rate cannot	Minimum wage laws raise the cost of some low-skill laborers above market equilibrium	Increased unemployment as people who wish to work at the going rate cannot
Valuation of options	Higher volatility increases the option premium <u>because of</u> greater risk it brings to the seller	Higher volatility brings greater risk to the seller	Higher volatility increases the option premium
Crowding out (economics)	One channel of crowding out is a reduction in private investment that occurs <u>because of</u> an increase in government borrowing	An increase in government borrowing	One channel of crowding out is a reduction in private investment
Hedge (finance)	Therefore, the farmer has reduced his risks to fluctuations in the market of wheat <u>because</u> he has already guaranteed a certain number of bushels for a certain price	The farmer has guaranteed a certain number of bushels for a certain price	The farmer has reduced his risks to fluctuations in the market of wheat
Productivity	Increasing national productivity can raise living standards <u>because</u> more real income improves people’s ability to purchase goods and services, enjoy leisure, improve housing, and education and contribute to social and environmental programs	More real income improves people’s ability to purchase goods and services	Increasing national productivity can raise living standards

Table 1: Five random positive examples from EconNLI. In these examples, the premise can cause the hypothesis. The Wiki Page and Source Sentence are the sources where the (premise, hypothesis) pair are extracted.

### 3 Natural Language Inference on Economic Events

We study the problem of causal relation inference on the economic events in the natural language form. In our study, we define an event as somebody or something having some actions or states. Economic events are the events that are related to the economic subject or phenomenon, for example, events describing production, distribution, and consumption (Krugman and Wells, 2013).

We construct a dataset, EconNLI, to evaluate LLMs’ causal relation inference ability on economic events. It consists of sentence pairs of (*premise*, *hypothesis*), both of which represent economic events. In positive pairs, the occurrence of the premise event leads to the occurrence of the hypothesis event based on economic theories. For negative pairs, no such causal relationship is supported by any theory. Table 1 provides five random positive examples from the EconNLI test set, illustrating when the premise can cause the hypothesis.

Note that in the real world, multiple premises may cause the hypothesis. For example,  $premise_1$  and  $premise_2$  may be the joint causes of the *hypothesis*. However, in our benchmarking, we only consider whether  $premise_1$  causes the *hypothesis* and  $premise_2$  causes the *hypothesis*

separately. If  $premise_1$  and  $premise_2$  are joint causes of the *hypothesis*, we exclude this example in our dataset.

## 4 EconNLI Dataset Construction

In this section, we describe the construction of EconNLI. The outline of the construction framework is shown in Figure 2. We first prepare a corpus about the economics from Wikipedia, and extract the events from the corpus using a supervised fine-tuned LLAMA2 model. Next, we generate candidates for premises and hypotheses based on a given reference sentence and assign a label. To ensure the dataset’s scale and quality, we use ChatGPT and GPT-4 to assign labels to the training set and annotate the testing set with manual inspection to align with human domain knowledge.

### 4.1 Corpus Preparation

To build a domain-specific corpus, we collected all pages related to economics and finance from Wikipedia, resulting in a corpus of 862 articles. To identify sentences containing causal relationships, we follow Sadat and Caragea (2022) and extract sentences with causal linking phrases. The linking phrases are provided in Appendix A. In total, we identify 5,077 sentences to construct the dataset.

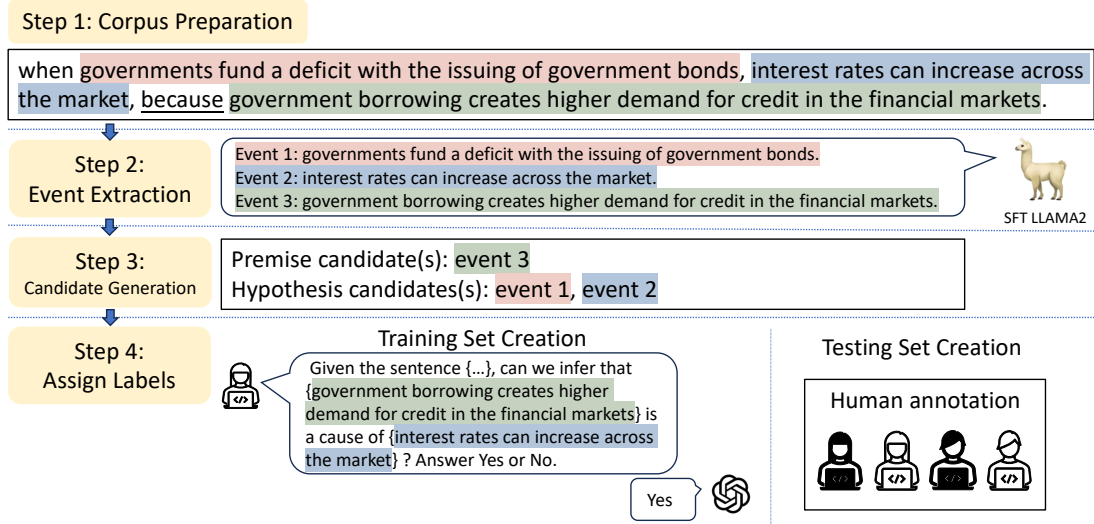


Figure 2: Our framework of generating the EconNLI dataset. We first prepare a corpus about the economics from Wikipedia, and extract the events using a supervised fine-tuned LLAMA2 model. Then we generate the candidates for the premises and hypotheses, and assign labels to them with both LLMs and human annotations.

## 4.2 Event Extraction

As our goal is to conduct inference among economic events, a crucial step involves extracting meaningful events from the source sentence. Directly splitting the sentence into two segments by the linking phrase can cause several issues: first, a sentence segment may not accurately express an event or contain multiple events; second, it can lead to potential grammar mistakes or unclear pronoun references. These problems can create confusion in the subsequent reasoning task.

To address this issue, we propose summarizing events with LLMs from the source sentence rather than generating events through syntax patterns. To ensure the LLM can summarize events as required, we manually write down all possible events for 125 sentences and perform supervised fine-tuning (SFT) on an LLAMA2 model for event generation. We use the following prompt for SFT: *### Question: An event is defined as somebody/something having some actions/states. List all events that can be inferred from the given sentence. If no event exists, simply respond "No event." The sentence: {source sentence} \n ### Answer: {human label}*. After the SFT, we apply the fine-tuned LLAMA2 model to the corpus prepared in 4.1 and extract all possible events from every sentence in the corpus.

## 4.3 Positive Pairs Construction

After summarizing the events from the sentences, we construct positive pairs in EconNLI.

### 4.3.1 Candidates Generation

To create positive pairs in EconNLI, we first generate candidates for premises and hypotheses based on syntax patterns. This step makes creating prompts for asking ChatGPT and GPT-4 to generate training set labels easier in the next stage.

Given a sentence containing a causal linking phrase (e.g.,  $s_1$  because  $s_2$ ), we split the sentence into two segments by the linking phrase (because) and denote the first sentence segment as  $s_1$ , and the second segment as  $s_2$ . We assume the premise event is contained in  $s_2$  and the hypothesis event is in  $s_1$ . Under this assumption, for an event  $e$  summarized by LLAMA2, we compute the cosine similarity between the bag-of-words representation of event  $\vec{v}_e$  and the two segments  $\vec{v}_{s_1}$ ,  $\vec{v}_{s_2}$  respectively. If  $\text{similarity}(\vec{v}_e, \vec{v}_{s_1}) < \text{similarity}(\vec{v}_e, \vec{v}_{s_2})$ , we consider  $s_2$  to include the event  $e$ , and regard  $e$  as a premise candidate; otherwise, we regard  $e$  as a hypothesis candidate.

Likewise, for a sentence containing a linking phrase of effects (e.g.  $s_1$  resulting in  $s_2$ ), If  $\text{similarity}(\vec{v}_e, \vec{v}_{s_1}) < \text{similarity}(\vec{v}_e, \vec{v}_{s_2})$ , we assign the event  $e$  as a premise candidate; otherwise, we assign it as a hypothesis candidates.

Using this approach, after processing all events in a source sentence, we generate a (source sentence, premise candidates, hypothesis candidates) tuple for each source sentence, where the premise and hypothesis candidates are sets of events, respectively. Next, we assign the labels.



### 4.3.2 Assign Labels

To create the training set, for each possible (source sentence, premise, hypothesis) tuple derived from the (source sentence, premise candidates, hypothesis candidates), we ask ChatGPT and GPT-4 to assign a label. Using the prompt *Given the sentence {source sentence}, can we infer that {premise} is a cause of {hypothesis}? Answer Yes or No*, we ask ChatGPT and GPT-4 to classify whether the event in the premise can cause the event in the hypothesis based on the source sentence. By providing the reference sentence, we enhance the reliability of LLMs’ answers and reduce the chance of generating a hallucinated answer.

Creating the testing set requires human expertise in economics and finance. Therefore, instead of using crowd-sourced workers, the authors with relevant backgrounds manually annotate the testing set. We select sentences from the training set that ChatGPT identifies as positive and assess if the premise leads to the hypothesis based on relevant economic theories. We manually annotate 513 positive examples and then exclude them from the training set.

### 4.4 Negative Pairs Construction

We also include negative pairs in EconNLI, where the premise is not a valid cause of the hypothesis. We create negative hypotheses by sampling from all possible hypotheses. For each premise from the examples that ChatGPT identified as positive, we first retrieve 20 hypotheses with the highest cosine similarity to the target premise as candidates, where the similarity is based on the BERT embeddings. We then exclude candidates from the same Wikipedia page as the premise to avoid false negatives. Finally, we randomly sample one hypothesis from the candidates as the negative hypothesis.

To improve the dataset quality, we use ChatGPT and GPT-4 to classify the relation between the sampled premise and hypothesis using the prompt *Can we infer that {premise} is a cause of {hypothesis}? Answer Yes or No* on the training set. We only preserve the examples where either ChatGPT or GPT-4 answers no. Authors manually check all negative pairs for the testing set to ensure validity.

### 4.5 Dataset Statistics

Table 2 presents the descriptive statistics of the EconNLI dataset. The dataset is divided into two sets: the training set and the test set. The training

	# Examples	# Topics	# Avg. Words
Train	10,810	554	17.84
Test	1,026	211	19.54

Table 2: Descriptive statistics of EconNLI dataset.

set includes 10,810 examples, covering 554 different economic topics (Wikipedia pages), with an average of 17.84 words per (premise, hypothesis) example. We provide two labels for each example in the training set, one from ChatGPT and the other from GPT-4. The Cohen’s kappa coefficient between ChatGPT and GPT-4 labels is 0.731.

The test set contains 1,026 examples covering 211 economic topics, with an average of 19.54 words per example. We also compare EconNLI with other popular financial NLP datasets in Appendix B.

## 5 EconNLI Evaluation

In this section, we use EconNLI to benchmark the LLMs’ performance on economic event reasoning. We examine the language models’ capabilities in inferring the correct relations between the event pairs and the correctness in generating the possible resulting events of a given premise.

### 5.1 Evaluated Language Models

Our experiment includes a wide range of language models, from the encoder-only to the decoder-only architecture. Specifically, we examine the following LLMs:

**Encoder-only Language Models.** For general domain LLMs, we consider (1) BERT (Devlin et al., 2019) and (2) RoBERTa (Liu et al., 2019): two popular language models for sequence classifications. Besides, we also evaluate the LLMs of the financial domain, including (1) FinBERT (Yang et al., 2020), a domain-specific model trained on financial text data. (4) FLANG-BERT and FLANG-ELECTRA (Shah et al., 2022) which use finance domain-specific pre-training with preferential masking to build more robust representations for the domain.

**Decoder-only Language Models.** We also evaluate the performance of various popular decoder-only language models: (1) LLAMA2 (Touvron et al., 2023b) is a popular open-source LLM pre-trained on extensive online data, and we use the LLAMA2-Chat version, which is optimized for dialogue use cases. (2) Alpaca (Taori et al., 2023) is fine-tuned from a LLAMA-7B model (Touvron

	Negative Pairs			Positive Pairs			Average		
	Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
Random	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500
BERT	0.769 $\pm$ 0.006	0.832 $\pm$ 0.013	0.800 $\pm$ 0.004	0.818 $\pm$ 0.009	0.750 $\pm$ 0.012	0.783 $\pm$ 0.004	0.794 $\pm$ 0.003	0.791 $\pm$ 0.003	0.791 $\pm$ 0.003
RoBERTa	0.822 $\pm$ 0.01	0.702 $\pm$ 0.005	0.758 $\pm$ 0.006	0.740 $\pm$ 0.004	0.848 $\pm$ 0.01	0.790 $\pm$ 0.007	0.781 $\pm$ 0.007	0.775 $\pm$ 0.006	0.774 $\pm$ 0.006
FinBERT	0.740 $\pm$ 0.015	0.814 $\pm$ 0.04	0.775 $\pm$ 0.011	0.794 $\pm$ 0.027	0.713 $\pm$ 0.035	0.751 $\pm$ 0.007	0.767 $\pm$ 0.007	0.764 $\pm$ 0.003	0.763 $\pm$ 0.002
FLANG-BERT	0.778 $\pm$ 0.011	0.810 $\pm$ 0.022	0.794 $\pm$ 0.007	0.802 $\pm$ 0.015	0.769 $\pm$ 0.02	0.785 $\pm$ 0.006	0.790 $\pm$ 0.006	0.789 $\pm$ 0.005	0.789 $\pm$ 0.005
FLANG-ELECTRA	0.827 $\pm$ 0.012	0.692 $\pm$ 0.025	0.753 $\pm$ 0.01	0.736 $\pm$ 0.013	0.855 $\pm$ 0.017	0.791 $\pm$ 0.002	0.781 $\pm$ 0.003	0.774 $\pm$ 0.005	0.772 $\pm$ 0.006
LLAMA2-chat(7B)	0.889 $\pm$ 0.003	<b>0.867<math>\pm</math>0.008</b>	<b>0.878<math>\pm</math>0.004</b>	<b>0.870<math>\pm</math>0.007</b>	0.891 $\pm$ 0.004	<b>0.881<math>\pm</math>0.003</b>	<b>0.879<math>\pm</math>0.003</b>	<b>0.879<math>\pm</math>0.003</b>	<b>0.879<math>\pm</math>0.003</b>
LLAMA2-chat(13B)	<b>0.919<math>\pm</math>0.012</b>	0.818 $\pm$ 0.06	0.865 $\pm$ 0.031	0.838 $\pm$ 0.045	<b>0.927<math>\pm</math>0.014</b>	0.880 $\pm$ 0.021	0.878 $\pm$ 0.021	0.873 $\pm$ 0.025	0.872 $\pm$ 0.026

Table 3: Classification results of the language models after supervised fine-tuning.

et al., 2023a) on instruction-following data generated by the technique of Self-Instruct (Wang et al., 2023). (3) FINMA (Xie et al., 2023) is a financial LLM based on fine-tuning LLAMA (Touvron et al., 2023a) with instruction data. (4) ChatGPT (Ouyang et al., 2022) and GPT-4 (OpenAI, 2023) are two advanced LLMs pre-trained on a wide array of textual data and reinforced by human feedback.

## 5.2 EconNLI for Classification Evaluation

In this subsection, we use EconNLI as a classification benchmark to evaluate the LLM’s capability to reason about the cause-effect relations between economic events. Specifically, given a premise  $p$  and a hypothesis  $h$ , the LLM performs as a classifier  $f : (p, h) \rightarrow \{0, 1\}$ , where the label 1 represents the happening of  $p$  can result in the happening of  $h$ , and label 0 otherwise.

### 5.2.1 Experiment Setup

There are two main techniques for utilizing LLMs for sequence classification: one is supervised fine-tuning, typically for encoder-only models or some relatively small decoder-only language models, and the other is zero-shot or few-shot prompting, typically for large decoder-only models. We experiment with the classification of EconNLI with these two techniques in the following setup:

**Supervised fine-tuning.** For fine-tuning the encoder-only LLMs, we use the BERT (base, uncased), RoBERTa (base), FinBERT (pre-train), FLANG-BERT and FLANG-ELECTRA from Huggingface<sup>1</sup>, and the model fine-tuning is implemented via Trainer<sup>2</sup>. Besides, we supervised fine-tuning (SFT) the LLAMA2-chat (7B, 13B) with the Huggingface SFT Trainer<sup>3</sup> using LoRa (Hu et al., 2022). The hyperparameters setting is

<sup>1</sup><https://huggingface.co/>

<sup>2</sup>[https://huggingface.co/docs/transformers/main\\_classes/trainer](https://huggingface.co/docs/transformers/main_classes/trainer)

<sup>3</sup>[https://huggingface.co/docs/trl/sft\\_trainer](https://huggingface.co/docs/trl/sft_trainer)

given in Appendix C. The prompt we used for SFT is provided in Appendix D.1. For all models, We randomly select 10% examples from the training set as the validation set and fine-tune the model for three epochs, and the model with the best validation performance is selected for evaluation. We repeat the fine-tuning experiments three times for each LLM, and report the average performance along with the standard deviation on the test set.

**Prompting the decoder-only LLMs.** We experiment with three types of prompting: zero-shot prompting, in-context learning (ICL) prompting, and chain-of-thought (CoT) prompting (Wei et al., 2022). The prompts for zero-shot setting, ICL, and CoT are provided in Appendix D.2, D.3 and D.4 respectively. The demonstrations in ICL, containing one positive example and one negative example, are randomly sampled from the training set. For ChatGPT and GPT-4, We use the "gpt-3.5-turbo" and "gpt-4" model API from OpenAI, respectively, retrieved in January 2024. Other open-source models are experimented on huggingface, and the answer generation is based on the greedy search.

### 5.2.2 Results

Table 3 presents the classification results of the language models after supervised fine-tuning. Despite being fine-tuned on 10,810 training data, encoder-only language models - from general to domain-specific - failed to perform well on EconNLI classification. The best encoder-only model, BERT, scored less than 0.8 F1 on this binary classification task. However, supervised fine-tuning on LLAMA2-chat models, which have a larger number of parameters, significantly improved the results from encoder-only models, as the 7B and 13B LLAMA2 models achieved similar results with around 0.87 F1 score. Nonetheless, the language models’ economic reasoning ability still has room for improvement.

Table 4 displays the classification results of the

		Negative Pairs			Positive Pairs			Average		
		Precision	Recall	F1	Precision	Recall	F1	Precision	Recall	F1
Random		0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500	0.500
LLAMA2 (7B)	Zero	0.702	0.092	0.162	0.514	0.961	0.670	0.608	0.526	0.416
	ICL	0.000	0.000	0.000	0.500	1.000	0.667	0.250	0.500	0.333
	CoT	0.765	0.172	0.280	0.535	0.948	0.684	0.650	0.560	0.482
LLAMA2 (13B)	Zero	0.648	0.308	0.417	0.546	0.832	0.660	0.597	0.570	0.538
	ICL	0.000	0.000	0.000	0.500	1.000	0.667	0.250	0.500	0.333
	CoT	0.683	0.328	0.443	0.564	0.851	0.678	0.624	0.590	0.561
Alpaca	Zero	0.547	0.263	0.355	0.515	0.782	0.621	0.531	0.522	0.488
	ICL	0.000	0.000	0.000	0.500	1.000	0.667	0.250	0.500	0.333
	CoT	0.541	0.618	0.577	0.555	0.476	0.512	0.548	0.547	0.545
FINMA	Zero	0.582	0.807	0.676	0.685	0.419	0.520	0.633	0.613	0.598
	ICL	0.000	0.000	0.000	0.500	1.000	0.667	0.250	0.500	0.333
	CoT	0.580	0.774	0.663	0.660	0.439	0.527	0.620	0.606	0.595
ChatGPT	Zero	0.874	0.613	0.720	0.703	0.912	0.794	0.788	0.763	0.757
	ICL	0.848	0.587	0.694	0.684	0.895	0.775	0.766	0.741	0.734
	CoT	0.807	0.662	0.728	0.716	0.844	0.775	0.762	0.753	0.751
GPT-4	Zero	0.895	0.762	0.823	0.793	0.910	0.848	0.844	0.836	0.835
	ICL	0.899	0.784	0.838	0.808	0.912	0.857	<b>0.854</b>	<b>0.848</b>	<b>0.847</b>
	CoT	0.908	0.616	0.734	0.710	0.937	0.808	0.809	0.777	0.771

Table 4: Classification results of the language models with various prompting methods.

language models with various prompting methods. The open-source models, with parameters ranging from 7B to 13B, all demonstrate performance similar to random guesses. The open-source model with the best performance is FINMA, indicating that tuning on financial instructions improves the model’s capability in economic event reasoning. ChatGPT and GPT-4 significantly surpass the open-source models, possibly due to their larger model size, richer pre-training corpora, and the knowledge gained from human feedback. Nevertheless, the results of both ChatGPT and GPT-4 still notably lag behind the LLAMA2 model after supervised fine-tuning, suggesting potential improvement from acquiring more related knowledge.

Upon comparing various prompting strategies, we have observed that different strategies have varying effects on open-source and GPT-series models. With ICL prompting, the open-source models consistently predict the positive label, which is undesired as it does not provide valid information. On the other hand, CoT prompting improves the results compared to zero-shot prompting, indicating that prompting the model to reason step by step can enhance reasoning accuracy. For the GPT-series model, ICL prompting hinders the performance of ChatGPT but improves the performance in GPT-4. In contrast, CoT prompting hindered the performance of both models, likely because these LLMs with reinforcement learning feedback can under-

stand the intention in zero-shot prompts, and providing extra reasoning steps is unnecessary.

### 5.3 EconNLI for Generation Evaluation

EconNLI is also an ideal dataset for evaluating the generation quality of the LLMs, as it provides sentence pairs describing the economic events and their consequential effects. Therefore, it can be used to assess whether the LLMs can generate the possible resulting event provided a premise event. As the LLMs are usually used to provide advice for economic analysis, correctly inferring the consequences of specific economic phenomena is crucial for providing reliable advice and solutions in the economic domain.

#### 5.3.1 Experiment Setup

To evaluate the LLM’s generation quality, we provide an event from the EconNLI premises of the positive examples and ask LLM to generate the possible resulting event with the following prompt: *If the event {premise} happens, what event will be caused? Answer the possible resulting event very briefly in one sentence. answer:.*

As the real-world economic system is complicated, one economic event can have multiple consequences. As a result, the LLM’s answer may be an alternative event that is also correct but not included in our reference. Therefore, we avoid using the BLUE (Papineni et al., 2002) or ROUGE (Lin,

Premise	Hypothesis(Reference)	LLM's Answer	Explanation
The subordinate bonds are allocated any losses from the collateral before losses are allocated to the senior bonds	The senior bonds are given a credit enhancement	ChatGPT: The resulting event would be a higher probability of default for the senior bonds.	If losses occur, they are first allocated to the subordinate bonds, resulting in a lower probability of default for the senior bonds
The available amount of hard currency per person falls	The purchasing power of each unit of currency increases	ChatGPT: The value of the currency will decrease, leading to inflation.	The premise event makes money more scarce, and consequently, the purchasing power of each unit of currency increases
demand curve shifts to the left	price decreases	LLAMA2: ... Increase in the price of the event tickets.	The demand curve shifts to the left typically leads to a decrease in the price of the product or service.
Their exports pay for fewer imports	The importing country's current account deficits widen	LLAMA2: If the event Their exports pay for fewer imports happens, the possible resulting event could be: They have a trade surplus.	If a country's exports pay for fewer imports, it means that the country is experiencing a trade deficit, not a trade surplus
The 2020 Russia-Saudi Arabia oil price war	the substantial decrease in the price of oil	FINMA: The 2020 Russia-Saudi Arabia oil price war will have no real-world consequences.	the 2020 Russia-Saudi Arabia oil price war caused a significant drop in oil prices

Table 5: Five cases that LLMs generate wrong or hallucinated resulting events given a premise from EconNLI. Each row includes the premise, the correct hypothesis (reference), the LLM's answer, and an explanation of why the LLM's answer is incorrect. The errors demonstrate misunderstandings or logical fallacies that LLMs can exhibit in the context of economic reasoning.

	Ent.	Con.	Irr.	Con./Ent.(%)
LLAMA2-chat(7B)	321	28	164	8.72%
LLAMA2-chat(13B)	324	26	163	8.02%
Alpaca	304	28	181	9.21%
FINMA	41	29	443	70.73%
ChatGPT	322	20	171	6.21%
GPT-4	<b>382</b>	<b>18</b>	113	<b>4.71%</b>

Table 6: Evaluation results for generating consequent economic events. In this evaluation, we provide the generative language models with a premise and ask it to generate a possible resulting event. We use GPT-4 to classify the generated event as entailment (Ent.), contradiction (Con.), or irrelevant (Irr.) compared to the reference (hypothesis).

2004) metrics that measure the similarity between the generated sequence and the reference sentence. Instead, we use GPT-4 to evaluate the relations between the LLM's answer and our hypothesis (reference). We ask GPT-4 to classify the relations into entailment, contradiction, and irrelevant. The entailment class stands that the LLM's answer and our reference describe similar economic phenomena, while the contradiction class means they describe the opposite economic phenomena. The irrelevant class means that the LLM provides an event that is different or unrelated to our reference, implying the correctness of the LLM answer is unknown. This setup considers the complications of the economic system, where the irrelevant answer may also be a

possible resulting event. The prompt we used on GPT-4 to verify the LLM's answer is provided in Appendix D.5.

### 5.3.2 Results

Table 6 presents the quality of generating the consequent economic events given a premise. The results show that GPT-4 had the best generation quality, with the lowest contradiction/entailment rate of 4.71%. LLAMA2-chat (7B) and LLAMA2-chat (13B) had similar rates of 8.72% and 8.02%, respectively, while Alpaca had a slightly higher rate of 9.21%. The model FINMA had the highest rate of contradiction/entailment at 70.73%, probably because it is instruction-tuned on the classification datasets, which impedes its ability to generate fluent and relevant results. The experiment results suggest that all generative language models have the probability of generating incorrect answers when analyzing the potential consequences of economic events. While some models performed better than others, none of them achieved perfect performance. Therefore, it is important for people to carefully check the correctness of the generated content when using LLMs to provide financial advice for economic analysis. This will help avoid misunderstandings and errors in decision-making and ensure the generated content is reliable and accurate.



### 5.3.3 Case Study

Table 5 presents a case study evaluating the performance of different LLMs in generating possible resulting events based on given premises. Table 5 identifies five cases where the LLMs generate wrong or hallucinated answers, providing an intuitive understanding of how LLMs fail the economics reasoning. The case study demonstrates the potential limitations of using LLMs for generating content related to economic events. While LLMs can help generate hypotheses and explore potential outcomes, it is important to recognize their limitations and potential for generating inaccurate or misleading results. As such, it is recommended that LLMs be used in conjunction with human expertise and judgment and that the generated content be carefully evaluated and verified before being used for decision-making.

## 6 Conclusions

In conclusion, we propose a new task to evaluate (financial) LLMs' ability to reason about economics. To do so, we construct a new dataset, EconNLI, and apply it to benchmark LLM's ability to perform classification and generation tasks. Through experiments, we find that LLMs have limitations in understanding economic knowledge and reasoning about the potential results of specific economic events. This work provides a new benchmark and insights for constructing financial LLMs. By emphasizing the need for continued research in this area, we hope to inspire future studies that will further enhance the reasoning capabilities of LLMs in the financial domain.

## 7 Limitations

There are two main limitations to this work that should be considered. Firstly, the focus of the study is on the ability of LLMs to reason about economics. The study does not address the ability of LLMs to reason about other domain-specific tasks, such as legal or clinical reasoning. Therefore, the findings of this study should only be generalized to other domains with further evaluation.

Secondly, the dataset was constructed based on Wikipedia, which may not fully represent the complexity and variability of real financial reports. The reasoning scenarios in writing a financial report may have unique and important problems that are not captured in the dataset. Future work should explore using more diverse and representative cor-

pora to provide a more comprehensive evaluation of LLMs' performance in the financial domain.

## 8 Acknowledgement

We acknowledge Jiaxin Liu and Hanyu Duan for providing helpful feedback on labeling the EconNLI dataset.

## References

- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Scott Wen-tau Yih, and Yejin Choi. 2019. Abductive commonsense reasoning. *arXiv preprint arXiv:1908.05739*.
- Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. 2015. [A large annotated corpus for learning natural language inference](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 632–642. The Association for Computational Linguistics.
- Ethan Callanan, Amarachi Mbakwe, Antony Papadimitriou, Yulong Pei, Mathieu Sibue, Xiaodan Zhu, Zhiqiang Ma, Xiaomo Liu, and Sameena Shah. 2023. [Can gpt models be financial analysts? an evaluation of chatgpt and gpt-4 on mock cfa exams](#). *arXiv preprint arXiv:2310.08678*.
- Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan R. Routledge, and William Yang Wang. 2021. [Finqa: A dataset of numerical reasoning over financial data](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 3697–3711. Association for Computational Linguistics.
- Zhiyu Chen, Shiyang Li, Charese Smiley, Zhiqiang Ma, Sameena Shah, and William Yang Wang. 2022. [Convinqa: Exploring the chain of numerical reasoning in conversational finance question answering](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 6279–6292. Association for Computational Linguistics.
- Liyang Cheng, Xingxuan Li, and Lidong Bing. 2023. [Is gpt-4 a good data analyst?](#) *arXiv preprint arXiv:2305.15038*.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. [The PASCAL recognising textual entailment challenge](#). In *Machine Learning Challenges, Evaluating Predictive Uncertainty, Visual Object Classification and Recognizing Textual Entailment, First*

- PASCAL Machine Learning Challenges Workshop, MLCW 2005, Southampton, UK, April 11-13, 2005, Revised Selected Papers*, volume 3944 of *Lecture Notes in Computer Science*, pages 177–190. Springer.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. **BERT: pre-training of deep bidirectional transformers for language understanding**. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- FiQA. Financial question answering. <https://sites.google.com/view/fiqa>.
- Yue Guo, Chenxi Hu, and Yi Yang. 2023a. **Predict the future from the past? on the temporal data distribution shift in financial sentiment classifications**. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing, EMNLP 2023, Singapore, December 6-10, 2023*, pages 1029–1038. Association for Computational Linguistics.
- Yue Guo, Zian Xu, and Yi Yang. 2023b. **Is chatgpt a financial expert? evaluating language models on financial natural language processing**. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 815–821.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022. **Lora: Low-rank adaptation of large language models**. In *The Tenth International Conference on Learning Representations, ICLR 2022, Virtual Event, April 25-29, 2022*. OpenReview.net.
- Tushar Khot, Ashish Sabharwal, and Peter Clark. 2018. **Scitail: A textual entailment dataset from science question answering**. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32.
- P.R. Krugman and R. Wells. 2013. *Economics*. Macmillan higher education. Worth Publishers.
- Jean Lee, Nicholas Stevens, Soyeon Caren Han, and Minseok Song. 2024. **A survey of large language models in finance (finllms)**.
- Chin-Yew Lin. 2004. **ROUGE: A package for automatic evaluation of summaries**. In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Jiaxin Liu, Yi Yang, and Kar Yan Tam. 2024. **Beyond surface similarity: Detecting subtle semantic shifts in financial narratives**. *CoRR*, abs/2403.14341.
- 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**. *CoRR*, abs/1907.11692.
- Pekka Malo, Ankur Sinha, Pekka J. Korhonen, Jyrki Wallenius, and Pyry Takala. 2014. **Good debt or bad debt: Detecting semantic orientations in economic texts**. *J. Assoc. Inf. Sci. Technol.*, 65(4):782–796.
- N Gregory Mankiw, Ronald D Kneebone, and Kenneth J McKenzie. 2020. *Principles of microeconomics*. Nelson Education Ltd.
- N Gregory Mankiw and William M Scarth. 2010. *Macroeconomics (Canadian Edition)*. Macmillan.
- Alfred Marshall. 2009. *Principles of economics: unabridged eighth edition*. Cosimo, Inc.
- Rajdeep Mukherjee, Abhinav Bohra, Akash Banerjee, Soumya Sharma, Manjunath Hegde, Afreen Shaikh, Shivani Shrivastava, Koustuv Dasgupta, Niloy Ganguly, Saptarshi Ghosh, and Pawan Goyal. 2022. **Ectsum: A new benchmark dataset for bullet point summarization of long earnings call transcripts**. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 10893–10906. Association for Computational Linguistics.
- OpenAI. 2023. **GPT-4 technical report**. *CoRR*, abs/2303.08774.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. **Training language models to follow instructions with human feedback**. In *NeurIPS*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics, July 6-12, 2002, Philadelphia, PA, USA*, pages 311–318. ACL.
- David C Parkes and Michael P Wellman. 2015. **Economic reasoning and artificial intelligence**. *Science*, 349(6245):267–272.
- Rachel Rudinger, Vered Shwartz, Jena D. Hwang, Chandra Bhagavatula, Maxwell Forbes, Ronan Le Bras, Noah A. Smith, and Yejin Choi. 2020. **Thinking like a skeptic: Defeasible inference in natural language**. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4661–4675, Online. Association for Computational Linguistics.
- Mobashir Sadat and Cornelia Caragea. 2022. **Scinli: A corpus for natural language inference on scientific text**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7399–7409.

- Agam Shah, Suvan Paturi, and Sudheer Chava. 2023. [Trillion dollar words: A new financial dataset, task & market analysis](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2023, Toronto, Canada, July 9-14, 2023, pages 6664–6679. Association for Computational Linguistics.
- Raj Sanjay Shah, Kunal Chawla, Dheeraj Eidnani, Agam Shah, Wendi Du, Sudheer Chava, Natraj Ramman, Charese Smiley, Jiaao Chen, and Diyi Yang. 2022. [When FLUE meets FLANG: benchmarks and large pretrained language model for financial domain](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022, Abu Dhabi, United Arab Emirates, December 7-11, 2022*, pages 2322–2335. Association for Computational Linguistics.
- Ankur Sinha and Tanmay Khandait. 2020. [Impact of news on the commodity market: Dataset and results](#). *CoRR*, abs/2009.04202.
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. 2023. Stanford alpaca: An instruction-following llama model. [https://github.com/tatsu-lab/stanford\\_alpaca](https://github.com/tatsu-lab/stanford_alpaca).
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, Aurélien Rodriguez, Armand Joulin, Edouard Grave, and Guillaume Lample. 2023a. [Llama: Open and efficient foundation language models](#). *CoRR*, abs/2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *CoRR*, abs/2307.09288.
- Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A. Smith, Daniel Khashabi, and Hannaneh Hajishirzi. 2023. [Self-instruct: Aligning language models with self-generated instructions](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2023, Toronto, Canada, July 9-14, 2023, pages 13484–13508. Association for Computational Linguistics.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. [Chain-of-thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems 35: Annual Conference on Neural Information Processing Systems 2022, NeurIPS 2022, New Orleans, LA, USA, November 28 - December 9, 2022*.
- Adina Williams, Nikita Nangia, and Samuel R Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of NAACL-HLT*, pages 1112–1122.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhanjan Kambadur, David S. Rosenberg, and Gideon Mann. 2023. [Bloomberggpt: A large language model for finance](#). *CoRR*, abs/2303.17564.
- 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](#). *CoRR*, abs/2306.05443.
- Hongyang Yang, Xiao-Yang Liu, and Christina Dan Wang. 2023a. [Fingpt: Open-source financial large language models](#). *FinLLM Symposium at IJCAI 2023*.
- Yi Yang, Yixuan Tang, and Kar Yan Tam. 2023b. [Investlm: A large language model for investment using financial domain instruction tuning](#). *CoRR*, abs/2309.13064.
- Yi Yang, Mark Christopher Siy Uy, and Allen Huang. 2020. [Finbert: A pretrained language model for financial communications](#). *CoRR*, abs/2006.08097.
- Fei Yu, Hongbo Zhang, Prayag Tiwari, and Benyou Wang. 2023. [Natural language reasoning, a survey](#).
- Huaqin Zhao, Zhengliang Liu, Zihao Wu, Yiwei Li, Tianze Yang, Peng Shu, Shaochen Xu, Haixing Dai, Lin Zhao, Gengchen Mai, Ninghao Liu, and Tianming Liu. 2024. [Revolutionizing finance with llms: An overview of applications and insights](#).

Dataset	# total	Task
FPB (Malo et al., 2014)	4,840	Sentiment Analysis
FiQA-SA (FiQA)	1,173	Sentiment Analysis
Headlines (Sinha and Khandait, 2020)	10,570	Text Classification
FOMC (Shah et al., 2023)	2,480	Text Classification
FinQA (Chen et al., 2021)	8,281	Question Answering
ConvFinQA (Chen et al., 2022)	14,115	Question Answering
ECTSum (Mukherjee et al., 2022)	2,425	Text Summarization
EconNLI	11,836	Economics Reasoning

Table 7: Information about some popular datasets for financial NLP.

## A Linking Phrase

**Linking phrase of causes:** because, because of, owing to, due to, caused by

**Linking phrase of effects:** therefore, hence, thus, as a result, as a consequence, consequently

## B Comparisons with Other Financial NLP Datasets

Table 7 compares EconNLI with other well-known datasets for financial NLP. As a domain-specific dataset, EconNLI stands out due to its relatively large scale compared to other financial NLP datasets. Additionally, EconNLI is unique in that it addresses the problem of reasoning in economics, which has not been previously explored in other datasets. This makes EconNLI a valuable resource for researchers and practitioners interested in developing and evaluating large language models for economic reasoning.

## C Hyper-parameters for SFT

SFT on BERT, RoBERTa, FinBERT, FLANG-BERT: We fix the learning rate as  $2 \times 10^{-5}$ , weight decay as 0.01, and the batch size as 48. Other hyperparameters remain the default in Trainer.

SFT on LLAMA2: We set the hyperparameters  $r = 16$ ,  $\alpha = 32$  in LoRa, and train the model with  $5 \times 10^{-5}$  learning rate and 24 batch size for three epochs.

All experiments are run on a machine with four Nvidia 3090 GPUs.

## D Prompts in Experiments

### D.1 SFT prompting

### Question: Conduct inference on economic events. We provide a premise and a hypothesis, both of them are economical events. Infer whether the premise can cause the happening of the hypothesis. Only answer 'Yes' or 'No'. premise: {premise}, hypothesis: {hypothesis}.

### Answer: {answer} + tokenizer.eos\_token

Note that the loss is only calculated on the completion after the term "### Answer:".

### D.2 Zero-shot Prompting

Conduct inference on economic events. We provide a premise and a hypothesis, both of them are economical events. Infer whether the premise can cause the happening of the hypothesis. Only answer 'Yes' or 'No'. Premise: {provided premise}, hypothesis: {provided hypothesis}, answer:

### D.3 ICL Prompting

Conduct inference on economic events. We provide a premise and a hypothesis, both of them are economical events. Infer whether the premise can cause the happening of the hypothesis. Only answer 'Yes' or 'No'. Here are some examples: premise:{premise of the positive example}, hypothesis: {hypothesis of the positive example}, answer:Yes

premise: {premise of the negative example}, hypothesis: {hypothesis of the negative example}, answer:No

Conduct inference on the following premise and hypothesis: premise: {premise for prediction}, hypothesis: {hypothesis for prediction}, answer:

### D.4 COT Prompting

### Question: Conduct inference on economic events. We provide a premise and a hypothesis, both of them are economic events. Infer whether the premise can cause the hypothesis to happen. Write the reasoning chain on the first line, and summarize the answer as 'Yes' or 'No' in the second line. premise: demand increases, hypothesis: price increases.

### Answer: When demand for a product or service increases, more people want to buy it. This creates a situation where there are more buyers than available supply, which leads to an increase in competition among buyers. As a result, sellers can raise their prices because they know that buyers are willing to pay more to get the product or service they want. \n Yes. \n

### Question: Conduct inference on economic events. We provide a premise and a hypothesis, both of them are economic events. Infer whether the premise can cause the hypothesis to happen. Write the reasoning chain on the first line, and summarize the answer as 'Yes' or 'No' in the second line. premise: government borrowing cre-



ates higher demand for credit in the financial markets, hypothesis: interest rates decreases across the market.

### Answer: When the government borrows money, it creates higher demand for credit in the financial markets. This is because the government is competing with other borrowers for available funds, which can drive up interest rates. Therefore, it is unlikely that government borrowing would cause interest rates to decrease across the market.

\n No. \n

### Question: Conduct inference on economic events. We provide a premise and a hypothesis, both of them are economic events. Infer whether the premise can cause the hypothesis to happen. Write the reasoning chain on the first line, and summarize the answer as 'Yes' or 'No' in the second line. premise: {premise}, hypothesis: {hypothesis}

### Answer:

#### **D.5 Prompt for Generation Quality Evaluation**

The prompt we used on GPT-4 to verify the LLM's answer is: I will give you two economical events. Determine the connection between these two events, choosing from 'entailment', 'contradiction', and 'irrelevant'. 'Entailment' means the two events describe similar economic phenomena, 'contradiction' means the two events describe the opposite economic phenomena, 'irrelevant' means the two events are unrelated or do not belong to one of the above two classes. Only provide the label. Event 1: {LLM answer}, event 2: {hypothesis}