Discussion:

Large Language Models and Return Prediction in China

Lin Tan, Huihang Wu, and Xiaoyan Zhang

Yinan Su
Johns Hopkins University Carey Business School

ABFER Webinar Series – Capital Market Development: China and Asia
November 20, 2024

Big picture

- Information is central to finance
  - efficient market hypothesis (EMH) is about information efficiency

Big picture

- Information is central to finance
  - efficient market hypothesis (EMH) is about information efficiency
  - asset prices are forward-looking valuations conditional on information

$$price_t = \mathbb{E}(\cdots|\text{time } t \text{ information}) \tag{1}$$

- Language is a prominent way to convey information in the securities market
  - EMH idea works with "natural" intelligence
- AI (LLMs) is profoundly changing this process
  - how information is processed and incorporated into asset prices
  - how economists learn capital market dynamics

Big picture

- Information is central to finance
  - efficient market hypothesis (EMH) is about information efficiency
  - asset prices are forward-looking valuations conditional on information

$$price_t = \mathbb{E}(\cdots|\text{time } t \text{ information}) \tag{1}$$

- Language is a prominent way to convey information in the securities market
  - EMH idea works with "natural" intelligence
- AI (LLMs) is profoundly changing this process
  - how information is processed and incorporated into asset prices
  - how economists learn capital market dynamics
- This paper provides important new insights on LLM applications to the Chinese stock market
  - important topic
  - extensive analysis
  - high-quality and interesting research!

This paper in the big picture

Explain price movements with (textual) information
To what extent?

- In ideal settings, presumably:
    - concurrent: $\approx$100%

This paper in the big picture

Explain price movements with (textual) information
To what extent?

- In ideal settings, presumably:
  - concurrent: $\approx$100%
    1988 AFA Presidential Address (Roll, "R^2")
    "... one could presumably "explain" stock price movements with an astronomical degree of accuracy. The R2 should be close to 1.0."
  - predictive: 0%
    if efficient market hypothesis holds
- In practice

This paper in the big picture

Explain price movements with (textual) information
To what extent?

- In ideal settings, presumably:
  - concurrent: ≈100%
    1988 AFA Presidential Address (Roll, "R^2")
    "... one could presumably "explain" stock price movements with an astronomical degree of accuracy. The R2 should be close to 1.0."
  - predictive: 0%
    if efficient market hypothesis holds
- In practice
  - *recent textual analysis in finance and economics is the* **tip of the iceberg***...an exciting research agenda [in the future], in which economists gradually expand sourced text corpora and* **increasingly refine their ability** *to elicit information from that text*
    (Chen, Kelly, Xiu, 2023)

This paper in the big picture

Explain price movements with (textual) information
To what extent?

- In ideal settings, presumably:
  - concurrent: ≈100%
    1988 AFA Presidential Address (Roll, "R^2")
    "... one could presumably "explain" stock price movements with an astronomical degree of accuracy. The R2 should be close to 1.0."
  - predictive: 0%
    if efficient market hypothesis holds
- In practice
  - *recent textual analysis in finance and economics is the* **tip of the iceberg***...an exciting research agenda [in the future], in which economists gradually expand sourced text corpora and* **increasingly refine their ability** *to elicit information from that text*
    (Chen, Kelly, Xiu, 2023)
  - predictive:
    This paper: > 0**!**, for daily stock-level returns
    highly profitable trading strategies
  - implication: inefficient prices

Summary

Two main points:

- News can predict stock returns by applying LLM, and yield a highly profitable trading strategy
- LLMs can be helpful in processing public news, and thus contribute to overall market efficiency

How is the forecasting done, in a nutshell

- Step 1, text representation with LLM
  news articles → embeddings (a long vector of numbers representing the semantics of the text)

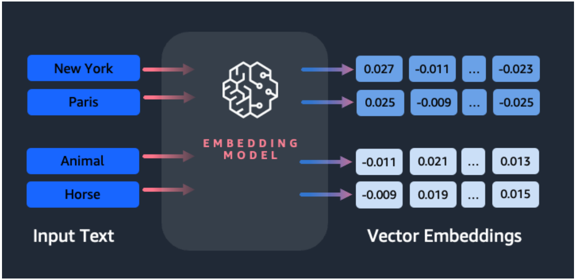

illustration source: AWS Machine Learning Blog. https://aws.amazon.com/blogs/machine-learning/getting-started-with-amazon-titan-text-embeddings/

How is the forecasting done, in a nutshell

- Step 1, text representation with LLM
  news articles $\rightarrow$ embeddings (a long vector of numbers representing the semantics of the text)

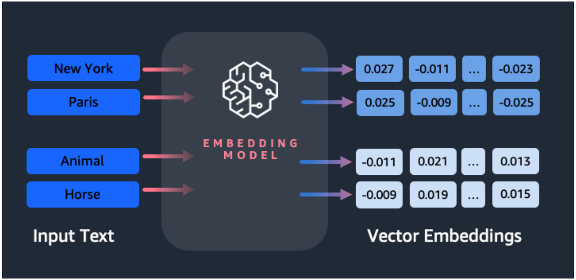

illustration source: AWS Machine Learning Blog. https://aws.amazon.com/blogs/machine-learning/getting-started-with-amazon-titan-text-embeddings/

  (link news article to stock-day)

- Step 2, return prediction with ML
  (high-dimensional regression with ridge regularization)

  $$\text{return}_{i,t+1} \text{ onto embeddings}_{i,t}$$

Implications on market efficiency

Four "steps":

- *"LLM signals can predict future earnings surprises, which suggest that the LLM signals contain information regarding firm fundamentals"*
- *Predictive power of LLM signals is higher when firms' information environments are less transparent, when retail investor holdings are higher and when news are more complex*
- *the predictive power of LLM signals remains positive over two days after news releases, but diminishes after the third day.*
- *different investors load their trades differently on the LLM signals, and some of them can benefit from LLM signals*

"These findings suggest LLMs can be helpful in processing public news, and thus contribute to overall market efficiency."

Discussion points

- implications on market efficiency
- uses of LLM embeddings and potential applications in finance

Implications on market efficiency

Finding:

- LLM signals can predict future returns

Implications:

- market is not perfectly efficient
- investors can make money by applying the LLM signals
- **LLMs can contribute to overall market efficiency**

  how? and how I should interpret this claim?

Implications on market efficiency

Finding:

- LLM signals can predict future returns

Implications:

- market is not perfectly efficient
- investors can make money by applying the LLM signals
- **LLMs can contribute to overall market efficiency**

  how? and how I should interpret this claim?

Combining the large quantity of retail traders and their low financial literacy, it is reasonable to expect that these retail investors have difficulty in processing public information and trading on public news. ...by rapidly analyzing and disseminating news, LLMs can play a significant role in accelerating this process and fastening price discovery.

Implications on market efficiency

- Lack of market inefficiency due to the lack of LLM? or many other market frictions?
  - again, human intelligence is sufficient for EMH logic, and practically the markets have been approximately efficient

Implications on market efficiency

- Lack of market inefficiency due to the lack of LLM? or many other market frictions?
  - again, human intelligence is sufficient for EMH logic, and practically the markets have been approximately efficient
- How will the use of LLM improve market efficiency?
  - LLM-quant trading is likely not for retail investors
  - how will LLM adoption by sophisticated instituition investors necessarily improve market efficiency?

    LLM essentially makes information processing faster

Implications on market efficiency

- Lack of market inefficiency due to the lack of LLM? or many other market frictions?
  - again, human intelligence is sufficient for EMH logic, and practically the markets have been approximately efficient
- How will the use of LLM improve market efficiency?
  - LLM-quant trading is likely not for retail investors
  - how will LLM adoption by sophisticated instituition investors necessarily improve market efficiency?

  LLM essentially makes information processing faster
    - "it takes two days for LLM signals to be incorporated into stock prices."
    - LLM signals are still trained by supervising on future return,
      - statistical learning relies on human learning (reading and inference) albeit slower
      - still requires trading and market clearing for price formation
      - limited arbitrage capital, risk-bearing capacity, and other limits to arbitrage are still there
    - "Quantity, Risk, and Return" (WP, 2024)

      sophisticated and noise traders interact

      incorporate quantity information from market trading activities to factor models and return prediction

Implications on market efficiency

At a high level:

- This webinar: “Capital Market Development: China and Asia”
- What are AI’s roles?

Implications on market efficiency

At a high level:

- This webinar: "Capital Market Development: China and Asia"
- What are AI's roles?

  My thoughts inspired by the paper:

  - **small** pockets of **highly profitable** "hedge funds"
  - economists can better learn in what ways the market is inefficient, by using LLM as a proxy for human information processing in the market, or as a tool to incorporate textual data into economic research
    - great potential for research and policy impact
  - "Narrative Asset Pricing: Interpretable Systematic Risk Factors from News Text" (RFS 2023)

Text Generation LLMs and Embedding LLMs

	Text Generation LLMs	**Embedding LLMs**
Output	Text (e.g., sentences, paragraphs, dialogue).	Numerical vectors (fixed-length representations).
Purpose	Generate coherent, contextually relevant text.	Convert text into high-dimensional vector embeddings for tasks like similarity, search, and clustering.
Applications	Creative writing, chatbots, summarization, question-answering, coding.	Semantic search, recommendation systems, document clustering, and machine learning features.
Examples	GPT models, Llama 2, Claude.	OpenAI's text-embedding-ada-002, Sentence-BERT.

Text Generation LLMs and Embedding LLMs

	Text Generation LLMs	**Embedding LLMs**
Output	Text (e.g., sentences, paragraphs, dialogue).	Numerical vectors (fixed-length representations).
Purpose	Generate coherent, contextually relevant text.	Convert text into high-dimensional vector embeddings for tasks like similarity, search, and clustering.
Applications	Creative writing, chatbots, summarization, question-answering, coding.	Semantic search, recommendation systems, document clustering, and machine learning features.
Examples	GPT models, Llama 2, Claude.	OpenAI's text-embedding-ada-002, Sentence-BERT.

- Still, one can also extract vector embeddings from text generation models

  Text generation LLMs (e.g. GPT) process text through many layers attention neural networks

  Neuron activations produce high-dimensional representations of the text

Zero-shot: "learning" without sample and estimation

Learning the predictive models in this paper:

- Still high-dimensional estimation:

  regress return$_{i,t+1}$ onto embeddings$_{i,t}$, in the historical sample

- quite similar to the traditional (quantitative) prediction
- quite different from human inference of news

Zero-shot: "learning" without sample and estimation

Learning the predictive models in this paper:

- Still high-dimensional estimation:

  regress $\text{return}_{i,t+1}$ onto $\text{embeddings}_{i,t}$, in the historical sample

  - quite similar to the traditional (quantitative) prediction
  - quite different from human inference of news

New advance with LLMs:
zero-shot learning with generalizable representations of knowledge

- prompt:

  "read this article, do you think stock price will go up or down tomorrow?"

  (potential for retail investors adoption)

- embeddings + cosine similarity

  embed queries with directional semantics, such as "bullish", "optimistic", "growth", "positive"

  calculate article-level forecasts as semantic similarity to the queries

Semantic similarity

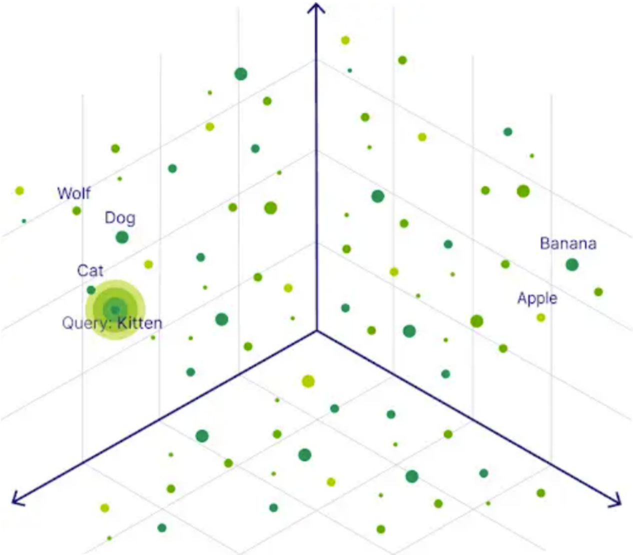

illustration source: From prototype to production: Vector databases in generative AI applications.
https://stackoverflow.blog/2023/10/09/from-prototype-to-production-vector-databases-in-generative-ai-applications/

Tracking narratives with LLM embeddings (in progress)

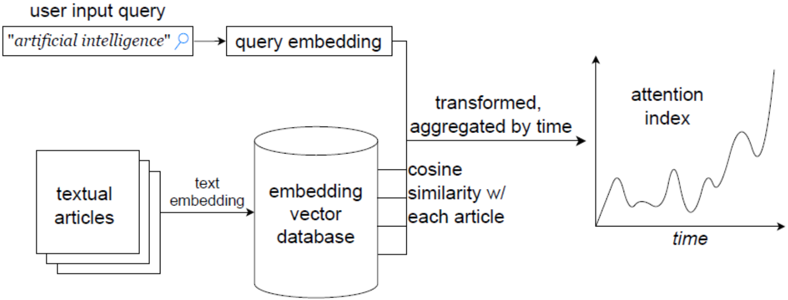

- any textual query
- web-based service open to all

Example: recession

Figure 3: Replicating Economic Recession Index

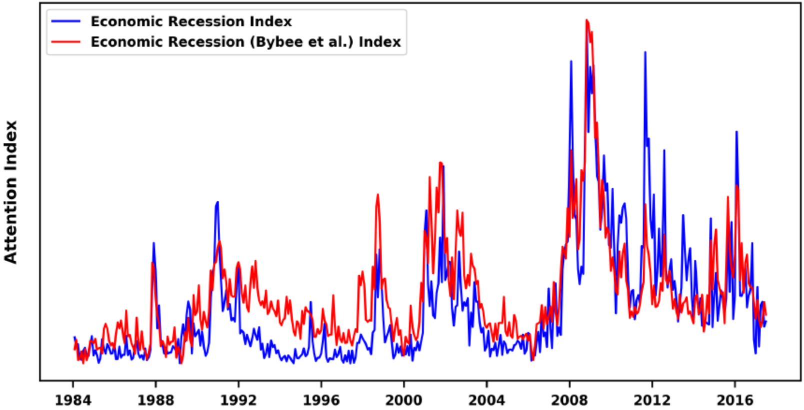

Example: ESG

Figure 5: ESG Index

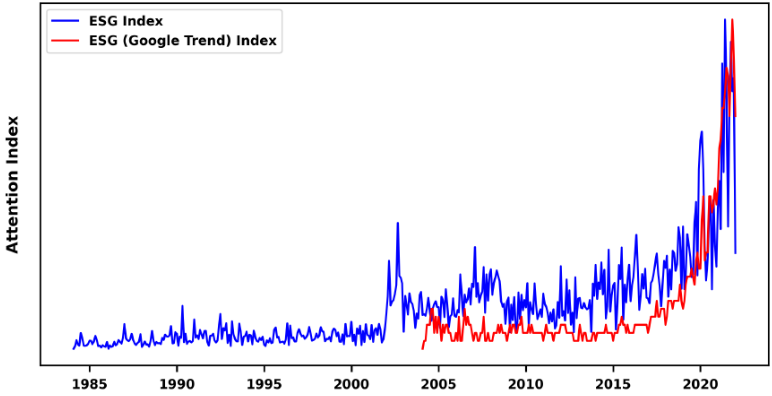

Example: financial stress

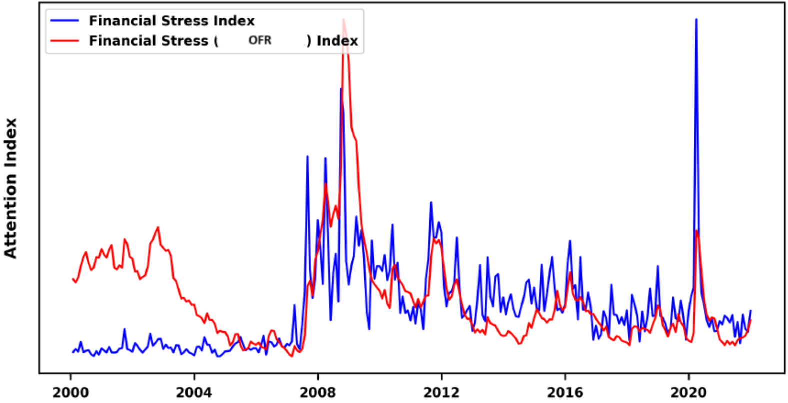

Discussion:

Large Language Models and Return Prediction in China

Lin Tan, Huihang Wu, and Xiaoyan Zhang

Yinan Su
Johns Hopkins University Carey Business School

ABFER Webinar Series – Capital Market Development: China and Asia
November 20, 2024