Crash Narratives

Goetzmann, Kim & Shiller (2024)

Discussion by Aditya Chaudhry

The Ohio State University

2nd January 2025

Overview

How do media “narratives” that recall historical crashes impact investor beliefs?

- **Narrativity:** “The manner in which facts are presented, affecting how they are interpreted and contextualized”

Overview

How do media “narratives” that recall historical crashes impact investor beliefs?

- **Narrativity:** “The manner in which facts are presented, affecting how they are interpreted and contextualized”

Use machine learning to measure daily narrativity in Wall Street Journal

- **Main measure:** Use Doc2Vec to measure similarity of news text to text used to describe October 19, 1987 crash

Overview

How do media “narratives” that recall historical crashes impact investor beliefs?

- **Narrativity:** “The manner in which facts are presented, affecting how they are interpreted and contextualized”

Use machine learning to measure daily narrativity in Wall Street Journal

- **Main measure:** Use Doc2Vec to measure similarity of news text to text used to describe October 19, 1987 crash

Empirically document that crash narrativity is associated with volatility and surveyed beliefs

- **Volatility:** Crash narrativity today predicts VIX tomorrow
- **Beliefs:** Crash narrativity today predicts surveyed subjective crash probabilities tomorrow

Overview

How do media "narratives" that recall historical crashes impact investor beliefs?

- **Narrativity:** "The manner in which facts are presented, affecting how they are interpreted and contextualized"

Use machine learning to measure daily narrativity in Wall Street Journal

- **Main measure:** Use Doc2Vec to measure similarity of news text to text used to describe October 19, 1987 crash

Empirically document that crash narrativity is associated with volatility and surveyed beliefs

- **Volatility:** Crash narrativity today predicts VIX tomorrow
- **Beliefs:** Crash narrativity today predicts surveyed subjective crash probabilities tomorrow

Contribution: Provide evidence for a particular mechanism of belief formation about rare disasters

Overall: Very Interesting Paper

Paper uses innovative methodology to tackle difficult question

1. Measuring narrativity is difficult
2. Isolating effect of narrativity vs. contemporaneous market conditions is difficult

Overall: Very Interesting Paper

Paper uses innovative methodology to tackle difficult question

1. Measuring narrativity is difficult
2. Isolating effect of narrativity vs. contemporaneous market conditions is difficult

Discussion: Suggestions on how paper can do even more to address these difficulties

- Provide more detail on how well Doc2Vec captures narrativity & potentially consider alternative methods
- Potentially use more structure to isolate narrativity vs. market conditions

Difficulty 1: Measuring Narrativity

What is Narrativity?

My interpretation: Holding facts fixed, are facts expressed as an analogy to past events?

- “When large shocks to the financial markets occur, historical references often play a role in news stories... Comparisons to past catastrophes make salient the gravity of current events and focus public attention on a singular narrative about what the future may bring.”

What is Narrativity?

My interpretation: Holding facts fixed, are facts expressed as an analogy to past events?

- "When large shocks to the financial markets occur, historical references often play a role in news stories... Comparisons to past catastrophes make salient the gravity of current events and focus public attention on a singular narrative about what the future may bring."

Contrast:

- Event X has occurred and may have outcome Z

What is Narrativity?

My interpretation: Holding facts fixed, are facts expressed as an analogy to past events?

- "When large shocks to the financial markets occur, historical references often play a role in news stories... Comparisons to past catastrophes make salient the gravity of current events and focus public attention on a singular narrative about what the future may bring."

Contrast:

- Event X has occurred and may have outcome Z

- Event X has occurred and is similar to historical event Y. Y had outcome Z. Thus, X may now also have outcome Z.

What is Narrativity?

My interpretation: Holding facts fixed, are facts expressed as an analogy to past events?

- “When large shocks to the financial markets occur, historical references often play a role in news stories... Comparisons to past catastrophes make salient the gravity of current events and focus public attention on a singular narrative about what the future may bring.”

Contrast:

- Event X has occurred and may have outcome Z

- Event X has occurred and is similar to historical event Y. Y had outcome Z. Thus, X may now also have outcome Z.

Higher-order linguistic relationship

- Must understand order of words to understand implicit causal chain being invoked

What is Doc2Vec?

Machine learning method to map text documents to numerical representations

- Neural network learns vector representations of whole document & individual words that best predict word sequences
- Example sentence in corpus: “the cat sat on the mat”
  - Use vectors for “the”, “cat”, “sat” & whole document to predict “on”

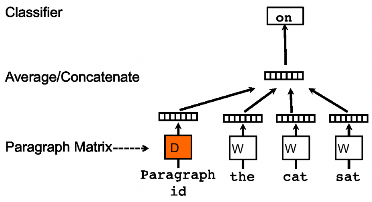

Figure: Le & Mikolov (2014)

What is Doc2Vec?

Machine learning method to map text documents to numerical representations

- Neural network learns vector representations of whole document & individual words that best predict word sequences
- Example sentence in corpus: “the cat sat on the mat”
  - Use vectors for “the”, “cat”, “sat” & whole document to predict “on”

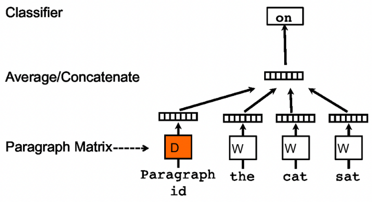

Figure: Le & Mikolov (2014)

Document vectors encode general meaning/topic of document

What is Doc2Vec?

Machine learning method to map text documents to numerical representations

- Neural network learns vector representations of whole document & individual words that best predict word sequences
- Example sentence in corpus: “the cat sat on the mat”
  - Use vectors for “the”, “cat”, “sat” & whole document to predict “on”

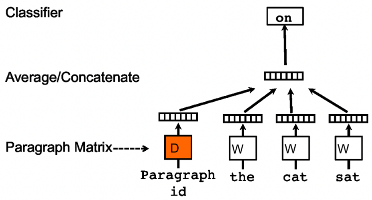

Figure: Le & Mikolov (2014)

Document vectors encode general meaning/topic of document

- Day t crash narrativity: How close are document vectors from Wall Street Journal news on t to those from Wall Street Journal news around 1987 crash

Can Doc2Vec Capture Narrativity?

Potentially to some extent

- Document vectors are trained to predict word sequences

Can Doc2Vec Capture Narrativity?

Potentially to some extent

- Document vectors are trained to predict word sequences
- Document vectors are trained to predict **short** word sequences (paper uses max distance of five words)

Can Doc2Vec Capture Narrativity?

Potentially to some extent

- Document vectors are trained to predict word sequences
- Document vectors are trained to predict **short** word sequences (paper uses max distance of five words)
- Implicit cause and effect relationships in narratives may play out over several sentences

Can Doc2Vec Capture Narrativity?

Potentially to some extent

- Document vectors are trained to predict word sequences
- Document vectors are trained to predict **short** word sequences (paper uses max distance of five words)
- Implicit cause and effect relationships in narratives may play out over several sentences

Suggestion: Validate Doc2Vec performance with tests of narrative recognition

- Randomly sample, for example, 500 documents
- Manually classify as “Contains 1987 crash narrative” or “Does not contain 1987 crash narrative”
- Report Doc2Vec performance in identifying which documents contain crash narratives

Can Doc2Vec Capture Narrativity?

Potentially to some extent

- Document vectors are trained to predict word sequences
- Document vectors are trained to predict **short** word sequences (paper uses max distance of five words)
- Implicit cause and effect relationships in narratives may play out over several sentences

Suggestion: Validate Doc2Vec performance with tests of narrative recognition

- Randomly sample, for example, 500 documents
- Manually classify as "Contains 1987 crash narrative" or "Does not contain 1987 crash narrative"
- Report Doc2Vec performance in identifying which documents contain crash narratives

Suggestion: Potentially consider more sophisticated models

- Transformer models (e.g. BERT) capture long-range meaning of text much better
  - But are more costly to train
  - Authors train own Doc2Vec models using only text available up to time *t* to avoid lookahead bias

Can Doc2Vec Capture Narrativity?

Potentially to some extent

- Document vectors are trained to predict word sequences
- Document vectors are trained to predict **short** word sequences (paper uses max distance of five words)
- Implicit cause and effect relationships in narratives may play out over several sentences

Suggestion: Validate Doc2Vec performance with tests of narrative recognition

- Randomly sample, for example, 500 documents
- Manually classify as "Contains 1987 crash narrative" or "Does not contain 1987 crash narrative"
- Report Doc2Vec performance in identifying which documents contain crash narratives

Suggestion: Potentially consider more sophisticated models

- Transformer models (e.g. BERT) capture long-range meaning of text much better
  - But are more costly to train
  - Authors train own Doc2Vec models using only text available up to time *t* to avoid lookahead bias
- Sarkar (2024) has released pre-trained, time-stamped BERT models to avoid lookahead bias

Difficulty 2: Narrativity vs. Market Conditions

Goal: Measure Impact of Crash Narratives on Investor Beliefs

Simplified empirical framework

$$\text{Subjective Crash Probability} = \beta \text{Crash Narrativity} + f\,(\text{Market Conditions}) + \epsilon$$

- Subjective crash probabilities obtained from surveys of individual and institutional investors
  - Likely depend on general market conditions
  - Want to test if they depend on crash narrativity ($\beta > 0$)

Goal: Measure Impact of Crash Narratives on Investor Beliefs

Simplified empirical framework

$$\text{Subjective Crash Probability} = \beta \text{Crash Narrativity} + f\left(\text{Market Conditions}\right) + \epsilon$$

- Subjective crash probabilities obtained from surveys of individual and institutional investors
  - Likely depend on general market conditions
  - Want to test if they depend on crash narrativity ($\beta > 0$)

Potential endogeneity concern: Selection

- Journalists choose to invoke crash narratives in states of the world with certain market conditions

$$\text{Crash Narrativity} = h\left(\text{Market Conditions}\right) + \nu$$

Goal: Measure Impact of Crash Narratives on Investor Beliefs

Simplified empirical framework

$$\text{Subjective Crash Probability} = \beta \text{Crash Narrativity} + f\left(\text{Market Conditions}\right) + \epsilon$$

- Subjective crash probabilities obtained from surveys of individual and institutional investors
  - Likely depend on general market conditions
  - Want to test if they depend on crash narrativity ($\beta > 0$)

Potential endogeneity concern: Selection

- Journalists choose to invoke crash narratives in states of the world with certain market conditions

$$\text{Crash Narrativity} = h\left(\text{Market Conditions}\right) + \nu$$

- Failure to control for $f\left(\cdot\right)$ or $h\left(\cdot\right)$ will bias estimates of β

Goal: Measure Impact of Crash Narratives on Investor Beliefs

Simplified empirical framework

$$\text{Subjective Crash Probability} = \beta \text{Crash Narrativity} + f\left(\text{Market Conditions}\right) + \epsilon$$

- Subjective crash probabilities obtained from surveys of individual and institutional investors
  - Likely depend on general market conditions
  - Want to test if they depend on crash narrativity ($\beta > 0$)

Potential endogeneity concern: Selection

- Journalists choose to invoke crash narratives in states of the world with certain market conditions

$$\text{Crash Narrativity} = h\left(\text{Market Conditions}\right) + \nu$$

- Failure to control for $f\left(\cdot\right)$ or $h\left(\cdot\right)$ will bias estimates of β
- $f\left(\cdot\right)$ and $h\left(\cdot\right)$ are potentially complex
  - Beliefs and text are complex (that is why we use highly non-linear machine learning methods)

Paper Tries to Address this Concern in Multiple Ways

$$\text{Subjective Crash Probability} = \beta\text{Crash Narrativity} + f\left(\text{Market Conditions}\right) + \epsilon$$
$$\text{Crash Narrativity} = h\left(\text{Market Conditions}\right) + \nu$$

Paper Tries to Address this Concern in Multiple Ways

$$\text{Subjective Crash Probability} = \beta \text{Crash Narrativity} + f\,(\text{Market Conditions}) + \epsilon$$
$$\text{Crash Narrativity} = h\,(\text{Market Conditions}) + \nu$$

Compare institutional vs. individual investor beliefs to difference out $f\,(\cdot)$(imposes structure on $f\,(\cdot)$)

- Both investor types respond homogeneously to market conditions (have same $f\,(\cdot)$)

Paper Tries to Address this Concern in Multiple Ways

$$\text{Subjective Crash Probability} = \beta\text{Crash Narrativity} + f\left(\text{Market Conditions}\right) + \epsilon$$
$$\text{Crash Narrativity} = h\left(\text{Market Conditions}\right) + \nu$$

Compare institutional vs. individual investor beliefs to difference out $f\left(\cdot\right)$(imposes structure on $f\left(\cdot\right)$)

- Both investor types respond homogeneously to market conditions (have same $f\left(\cdot\right)$)

Linear controls to absorb $h\left(\cdot\right)$ (imposes structure on $h\left(\cdot\right)$)

- Standard market condition proxies (e.g. lagged returns, lagged VIX, etc.)
- Similarity of word frequencies (i.e. news content) to news around 1987 crash
- Crash narrativity driven by non-lede paragraphs

Paper Tries to Address this Concern in Multiple Ways

$$\text{Subjective Crash Probability} = \beta\text{Crash Narrativity} + f\left(\text{Market Conditions}\right) + \epsilon$$
$$\text{Crash Narrativity} = h\left(\text{Market Conditions}\right) + \nu$$

Compare institutional vs. individual investor beliefs to difference out $f(\cdot)$(imposes structure on $f(\cdot)$)

- Both investor types respond homogeneously to market conditions (have same $f(\cdot)$)

Linear controls to absorb $h(\cdot)$ (imposes structure on $h(\cdot)$)

- Standard market condition proxies (e.g. lagged returns, lagged VIX, etc.)
- Similarity of word frequencies (i.e. news content) to news around 1987 crash
- Crash narrativity driven by non-lede paragraphs

Suggestion: Impose more structure on $h(\cdot)$

- Model selection problem of when journalists invoke crash narratives (Heckman (1979); Kelly, Manela & Moreira (2021))

Paper Tries to Address this Concern in Multiple Ways

$$\text{Subjective Crash Probability} = \beta \text{Crash Narrativity} + f\left(\text{Market Conditions}\right) + \epsilon$$
$$\text{Crash Narrativity} = h\left(\text{Market Conditions}\right) + \nu$$

Compare institutional vs. individual investor beliefs to difference out $f(\cdot)$(imposes structure on $f(\cdot)$)

- Both investor types respond homogeneously to market conditions (have same $f(\cdot)$)

Linear controls to absorb $h(\cdot)$ (imposes structure on $h(\cdot)$)

- Standard market condition proxies (e.g. lagged returns, lagged VIX, etc.)
- Similarity of word frequencies (i.e. news content) to news around 1987 crash
- Crash narrativity driven by non-lede paragraphs

Suggestion: Impose more structure on $h(\cdot)$

- Model selection problem of when journalists invoke crash narratives (Heckman (1979); Kelly, Manela & Moreira (2021))
- Three benefits:
  - Clarify implicit assumptions to separate impact of narrativity vs. market conditions
  - Yield even deeper understanding of when crash narratives arise & how do they propagate
  - Reveal potential feedback looks from narratives to market conditions and vice versa

Minor Comments

More details on Doc2Vec implementation would be useful

- E.g. Clarify if this is the Distributed Memory or Distributed Bag of Words version?

More discussion of why certain functional forms are chosen would be useful

- Why define adjusted similarity as the difference between natural log of one plus average cosine similarities?
- Why use quadratic projections (e.g. when projecting FolkMotif onto '87 Narrative)?

More discussion of magnitudes would be useful

- E.g. How large is one unit increase in '87 Narrative $\rightarrow$ 1.529 higher crash probability (Table 7, Column 1)

Results on crash narrativity & market-implied probabilities would be interesting

- Under suitable assumptions, the market-implied crash probabilities can be interpreted as marginal investor's beliefs

Typos

- Link in footnote 19 is broken
- Pg. 27: I think you mean to say Dot-Com bubble burst twelve years after 1987 crash, not two years

Conclusion

Very interesting paper

Creative methodology to address difficult and important question

- How do media "narratives" that recall historical crashes impact investor beliefs?
- Use machine learning to measure crash narrativity in Wall Street Journal text

Main comments

- Suggestions on how authors can provide even more empirical support for their main findings