NBER WORKING PAPER SERIES

CRASH NARRATIVES

William N. Goetzmann
Dasol Kim
Robert J. Shiller

Working Paper 30195
http://www.nber.org/papers/w30195

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2022, Revised November 2024

We thank Francesco Fabozzi, Milad Nozari, Karen Spitzer and James Tyler for valuable support. We thank Nico Forteza, Hongqi Liu, Greg Phelan, Wenting Song, Yinan Su, and Jessica Wachter for the valuable comments. We also thank participants in seminars at Erasmus, Amsterdam, Lehigh, Yeshiva, Villanova, Seoul National University as well as the participants in the CIBC and IVRSiF seminars, participants in the Miami Behavioral Conference, OFR-Federal Reserve Bank of Cleveland Financial Stability Conference, NBER Summer Institute, ECONDAT Conference, Society for Financial Studies Cavalcade Conference, Financial Intermediation Research Society Conference, and European Finance Association Meetings. We thank the International Center for Finance at the Yale School of Management for support. The opinions expressed are those of the authors and do not necessarily reflect those of the Office of Financial Research, the U.S. Department of the Treasury, or the National Bureau of Economic Research.

Crash Narratives
William N. Goetzmann, Dasol Kim, and Robert J. Shiller
NBER Working Paper No. 30195
July 2022, Revised November 2024
JEL No. E03, G00, G02, G11, G23

ABSTRACT

Narratives can serve as linguistic mechanisms for transmitting knowledge about important historical events. Shiller (2017) argues that narratives are potent mechanisms for the propagation of beliefs about the economy. We use a controlled context approach to test the effect of media narratives about historic stock market crashes on current investor beliefs and choices. Our methodology avoids the problem of large language model leakage that can lead to look-ahead biases. We find that crash narratives propagate broadly once they appear in news articles, and that they predict market volatility. We exploit investor heterogeneity using survey data to distinguish the effects of narrativity from response to fundamental factors. Finally, we develop a measure of pure narrativity to examine when the financial press is more likely to employ narratives.

William N. Goetzmann
School of Management
Yale University
Box 208200
New Haven, CT 06520-8200
and NBER
william.goetzmann@yale.edu

Dasol Kim
Department of the Treasury
717 14th Street NW
Washington D.C., 20005
dasol.kim@ofr.treasury.gov

Robert J. Shiller
Yale University, Cowles Foundation
Box 208281
30 Hillhouse Avenue
New Haven, CT 06520-8281
and NBER
robert.shiller@yale.edu

1. Introduction

Precisely because stock market crashes are infrequent, the history of past events, whether personally experienced or not, plays a crucial role in the assessment of extreme crash probabilities. Rare disaster concerns and the probabilistic assessment of them have been posed as an explanation for various findings in asset pricing.[4] Yet most people experience only a handful of extreme market shocks in their lifetimes, but of course they learn about others in various ways – the financial press being one important source. In this paper we explore the effect of media recall of historic shocks on investor belief formation about rare disasters. Shiller (2017) highlights the important role of narratives in the transmission of economic beliefs. Narratives are a social medium through which an idea can spread rapidly through conversation and potentially affect aggregate beliefs about asset prices.[5]

In non-literate cultures, "collective memory" – i.e. social knowledge of past events and mechanisms for recall -- is carried by stories, myths and folk-takes. In interviews with the Ongee people of Little Andaman Island following the catastrophic tsunami of 2004, Pandya (2005) learned that native Andamans largely avoided the disaster because of a traditional story about the conflict between the living and spirits of the dead. The withdrawal of the ocean immediately before the tsunami was a sign for people to flee to high ground because "angry spirits would come down to hunt us." The danger signal was embedded in a narrative that was compelling enough to be passed down through generations – in Shiller (1984) terms, it was infectious. More recent settlers to these remote islands retained no such tradition and suffered from the disaster in far greater proportion. The account of the Andaman islanders highlights not only the role of collective memory as a potentially life-preserving mechanism, but also the relevance of narrative devices such as plot, characters and disequilibrium in the maintenance and stimulus of collective memory.

Historically, newspapers have been an important vehicle for the spread of ideas. Newspaper articles are stories that collate facts related to current events. Their narrativity relates to the manner in which those facts are presented, affecting how they are interpreted and

[4] Cf. Reitz (1988), Barro (2006), Weitzmann (2007), Santa-Clara & Yan (2010), Berkman et. al. (2011), Bollerslev & Todorov (2011), Wachter (2013), Anderson et. al. (2020), Welch (2015), and Goetzmann, Kim & Shiller (2024).

[5] Shiller (1984) argues that the existence of smart-money investors in the economy will not entirely mitigate the effect of average investors subject to fads and fashions.

contextualized by readership. When large shocks to the financial markets occur, historical references often play a role in news stories. The role of the financial press as a mechanism for collective memory is potentially important. 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.

For example, the stock market crash on October 19, 1987 is an event that represents the largest one-day decline in the Dow Jones Industrial Index [DJIA] since the stock market crash on October 24, 1929. Figure 1 displays monthly counts of references in *The Wall Street Journal* to the crashes of October 24th, 1929 and October 19th, 1987.[6] Notice that the biggest spike in references to the crash of '29 occurred in October, 1987. Journalists the day after the '87 crash immediately invoked the specter of the Great Depression.[7] In the years that followed, however, the crash of October '87 – larger in magnitude that that of 1929 – came to dominate historical crash references. The pattern in Figure 1 is unmistakable: references to past crashes clearly varied with market conditions and were frequently associated with big drops in stock prices.

For readers who did not experience the crash of 1987 personally, the financial press renewed awareness of this specific event more than 2,000 times in the years that followed. A natural question is whether and how this collective memory mechanism affected investor attitudes and beliefs. Press references to the 1987 crash and the 1929 crash were embedded in rich, written narratives designed to interest and engage readers. Simple counts tell us very little about this narrative content.

We posit that investors use narratives to inform their beliefs and, in turn, their choices. Narratives often elicit causal relationships between a sequence of events connecting one state to another with things that happen in between. They thus suggest an analogy for how and why things are happening and what the future outcome might be. Investors may be especially influenced by

---

[6] We search the Eastern Edition of the Wall Street Journal in ProQuest from January, 1984 to April, 2022 on the terms "1929 crash, "crash of 1929", "Black Thursday", "1987 crash", "crash of 1987" and "Black Monday". Note that the terms Black Monday and Black Thursday have been occasionally used to describe other events, so there is some noise in these counts.

[7] "The Dow Jones Industrial Average plummeted an astonishing 508 points, or 22.6%, to 1738.74. The drop far exceeded the 12.8% decline on the notorious day of Oct. 28, 1929, which is generally considered the start of the Great Depression." The Crash of '87: Stocks Plummet 508.32 Amid Panicky Selling --- Percentage Decline Is Far Steeper Than '29; Bond Prices Surge. Tim Metz, Alan Murray, Thomas E. Ricks and Beatrice E. Garcia. Wall Street Journal, Eastern edition; New York, N.Y. [New York, N.Y]. 20 Oct 1987: 1

narratives related to rare events, such as stock market crashes, compared to common ones. The psychology literature provides evidence that individuals are susceptible to a variety of behavioral biases in evaluating low probability events.[8] Investors may seek a framework to contextualize rare events, and narratives conveyed by media may offer one. This suggests not only a role for the financial press in the propagation of narratives, but also, as we show below, a mechanism for aggregate market reaction.

To test this hypothesis, we develop a higher-order measure of news narrativity specifically related to major crashes. Detecting narrativity requires estimating structural relationships within highly unstructured data. Recent advances in computational linguistics allow us to quantify distributed representations of these relationships using high-dimensional semantic space models. We extract narratives that appear in the news in the days following the 1987 and 1929 crashes and use the model to estimate their prevalence in news articles each day for over 30 years. We argue that this measure captures "crash narratives" in the media. We show that the measure corresponds with major market events throughout the sample period, including the LTCM collapse, 9/11, the Great Financial Crisis, the European Debt Crisis, and the COVID-19 pandemic. This suggests that journalists broadly employ narratives used following the 1987 crash in spite of differences in contexts. The measure has a significant relationship with lagged, market-based crash indicators in the post-1987 period. Manual inspection suggests that these patterns are unlikely to be driven by the prevalence of specific words alone, but rather by the sequencing and distribution of words representing narrative structure.

Leakage is an important methodological challenge to the use of pre-trained large language models. Current popular generative AI models are trained on vast, temporally undifferentiated corpuses. As such they can exploit linguistic relationships that extend across time. However recent research has shown how this can bias hypothesis tests that condition only upon past information.[9] In finance, this can bias tests of predictive models that should not rely on knowledge of future events.

Due to limitations of machine power and data access, it is not currently feasible to construct a time-indexed version of a comprehensive large language model like ChatGPT that uses restricts

[8] Cf. Barberis (2013).
[9] CF. Bybee (2023), Glasserman & Lin (2023), Lopez-Lira & Tang (2023), Halawi et. al. (2024) & Sarkar & Vada (2024).

itself to prior information as of a given date. Previous researchers (eg. Bybee, 2023, Lopez-Lira & Tang, 2023) have addressed the leakage problem by using earlier versions of ChatGPT to estimate models and conduct out of sample tests. This typically leaves a short out of sample window. Our methodology addresses this. We construct a Doc2Vec model that focuses on comparisons with articles published prior to the sample period used for the analysis. Doing so limits the influence of forward-looking information in the analysis. The model is based upon the Wall Street Journal [WSJ], the dominant US financial newspaper in the 20th century, under the assumption that it pre-screens, curates and constructs narratives about the U.S. stock market that investors find relevant to decision-making. As a robustness check, we using a rolling estimation window to construct a series of Doc2Vec models that control for leakage by using nested, annual time-indexed contextual embeddings that are restricted to backward-looking WSJ data.[10] We confirm that our approach is effective in mitigating the influence of look-ahead bias.

Our first set of tests provide suggestive evidence that the prevalence of crash narratives in the media explains a significant amount of predictive variation in investor attention to stock market crashes and market volatility. Crash attention is defined as the daily internet search volumes of broad-based crash-related terms. We do not focus on terms related to individual stocks as the crash narrative measure is likely capturing aggregate market crash concerns. We find that crash narratives that appear the previous day have a positive and significant association with crash attention. These tests isolate variation in daily innovations in crash attention and account for potential biases due to serial correlation in the market-based controls. The effects almost double in magnitude when conditioning on periods of high general stock market attention, squaring with received wisdom. These results suggest broad propagation of crash narratives in the financial press. They also suggest that crash narratives prime investor crash concerns, and, in turn, market reaction. We find that crash narratives have a positive and significant association with future VIX levels however unlike Manela & Moreira (2017) we find no significant relationship to future market returns.

A potential issue with the crash attention tests is that they may not necessarily be attributable to investors since we cannot identify internet search volumes by individual users.

[10] This approach is closest in spirit to Jha et al (2023) who use time-indexed Google Books 5-grams to measure the evolution of sentiment about finance. It also allows us evaluate a number of things of interest to research about leakage bias: we quantify the stability of model weights and isolate the biasing effects of the leakage component.

Additionally, in spite of a large battery of controls, it is difficult to rule out the potential influence of unobservable fundamental heterogeneity unaccounted for by the models. We address these concerns by using investor survey data to directly examine the effects of crash narratives on subjective crash assessments, and exploit investor heterogeneity to directly account for potential omitted factors. We argue that fundamental factors should affect individual and institutional investor beliefs similarly. However, we also expect susceptibility to the influence of crash narratives on beliefs should decrease in investor sophistication.[11] Accordingly, we test for differential effects of crash narratives on individual investor beliefs relative to those of institutional investors.

We find evidence that crash narratives significantly affect investor crash beliefs. Investor crash probabilities are elevated following the appearance of crash narratives the previous day. There is no effect for institutional investors, but a positive and significant effect for individual investors. We go further by decomposing the crash narrative measure according to the position that the narratives appear in the articles. Specifically, we differentiate narratives that appear in the lede from non-lede paragraphs. Given that readers are more likely to read the lede paragraph compared to other parts of the article, we are able to conduct tests based on intra-article variation in salience. We find consistent effects for lede paragraph crash narratives, even when controlling for crash narratives that appear in other parts of the article. Finally, we show that investors' own narratives, provided in response to open survey questions, echo crash narratives in the financial press.

In the final set of tests, we examine periods when journalists are more likely to employ general narrativity. For these tests, we construct a pure measure of narrativity using information that is plausibly exogenous to financial market conditions. Specifically, we construct narrative measures based on classical folk tales published almost two centuries prior to our sample period. The pure narrativity measure is based on the similarity in semantic distributions between those folk tales and each article in the sample. We find a strong correspondence between the crash narrative measure and the pure narrativity measure, consistent with our interpretation of the crash narrative measure. We also show that journalists are more likely to employ general narratives

[11] Experimental evidence from the social psychology literature suggest that the prevalence of various biases is pronounced when individuals perceive that they have low expertise (Ottati and Isbell (1996), Sedikides (1995)), or low processing capacity (Greifeneder & Bless (2007), Siemer & Reisenzein (1998)).

during periods of high general stock market attention. This is consistent with the view that narratives are more likely to be used during periods of more diverse readership, and that narratives may be used to make the news more accessible. Finally, we show that the pure narrativity measure has a similar effect on investor crash beliefs, but only the component that is associated with crash narratives.

1.1 Related Literature

1.1.1 Narrative

Despite their obvious centrality to human society, narratives are difficult to precisely define. Bruner (1991) argues stories are the primary way that humans make sense of the world. Narrative scholars often define narratives as a constellation of specific characteristics. For example, Herman (2009) lists four key features of a narrative: it involves the act of telling, it represents a sequence of events that involve disequilibrium, it involves context – a world within which the sequence takes place, and it evokes a sense of experiencing or living through the events. Different narrative scholars emphasize some of these features more than others. Lively (2019) highlights the narrative act. Kermode (1967) highlights plot. The experiments in Morag & Loewenstein (2021) suggest that narrative framing can significantly affect valuation decisions.

Recently, some researchers have applied language processing tools to explore narrative, topical content and sentiment in financial news. Manela and Moreira (2017) use textual analysis to construct a news-implied proxy for the VIX fear index from *The Wall Street Journal* [WSJ] headlines and test their correlation to collective uncertainty and future market return. Bybee et al. (2021) extract topic models from WSJ news and find that news attention to topics explain a significant proportion of stock returns. Similarly, Bertesch et. al. (2021) test the Shiller (2017) hypothesis about fluctuating narratives in the business cycle. Flynn & Shastry (2002) test for the macroeconomic effects of narratives proposed in Shiller (2017). Larsen et. al. (2021) document the effect of news on inflation expectations. Dierckx et al. (2021) identify narratives in the financial press that forecast implied market volatility. Bertsch et. al. (2021) find interesting asymmetries in the narrative structures for boom vs. bust periods. Chahrour et. al. (2021) show that financial press narratives play an important motivating role in business cycle dynamics. Our work differs from these papers by focusing on two factors: the degree of narrativity in the text and the narrative mechanism of historical recall.

One of the challenges in our analysis is to distinguish between topical and "pure" narrativity i.e. narrativity that derives from structural association. We address this in two ways. The first approach is based on the presumption that a story is not a random collection of words – that the relationship among words defines a narrative. To that end we test a null that a non-contextual model, i.e. topic modeling based on bag-of-words analysis can identify relationships that the contextual Doc2Vec model has identified. If the content rather than the structure of the text explains the crash narrative measure, then a methodology that ignores word sequence and structure is likely to work as well as a contextual model.

Second, we use a high-narrativity text corpus completely unrelated in topic to financial and economic news: folk-motifs derived from an archetypal corpus of fairly-tales. Propp (1968), *Morphology of the Folktale* pioneered a structural approach to identifying recurring motifs: plot structures in traditional Russian folktales. Folklore scholars have refined this approach to study the structure and function of folktales.[12] Tales of the Brothers Grimm are widely used as archetypal narrative corpora.[13] We construct a folk-motif narrativity metric from a classification of the Tales of the Brothers Grimm and use it as an instrument for narrativity.

1.1.2 Rare Disasters

There is a rich literature on the effect of rare disasters expectations on financial markets. We refer readers to Tsai and Wachter (2015) for a review the re-emergence of attention to rare disaster risk in asset pricing, particularly focusing on dynamic models that have been proposed and tested.[14] Reitz (1988) and Barro (2006) show how rational rare disaster concerns could explain the equity premium puzzle. Recently, researchers have increased focus on mechanisms by which beliefs about rare disasters can form. Bordalo et. al. (2021) show how selective recall can induce overestimation of the probability of unlikely events. Enke et. al. (2020) provide experimental evidence in support of the memory channel. Goetzmann, Kim & Shiller (2024) find that the affect

[12] Cf. Thompson (1989), Wurtzback (1997) for example.
[13] Cf. Garry (2017).
[14] Cf. Reitz (1988), Barro (2006), Gabaix (2008), Wachter (2013), Tsai & Wachter (2015), Manela & Moreira (2016).

heuristic documented by Johnson & Tversky (1983) plays a role in crash probability assessment, and that it is mediated by financial press sentiment. van Binsbergen et. al. (2022) model the belief formation process in a Bayesian framework and show how biased beliefs can form when crashes fail to occur in small sample.

1.1.3 Memory

Our work also relates to research on memory as a mechanism for directing attention. Kahana (2012) summarizes the considerable experimental evidence that human memory works according to a few basic rules that together are referred to as *retrieved context*. Retrieved context relies on a few operators: recency, semantic similarity, and temporal contiguity. *Recency* refers to humans relying on more recent events when accessing memory. *Semantic similarity* refers to humans accessing events that are like those they are currently experiencing or asked to recall. *Temporal contiguity* refers to humans accessing a memory through broader contextual similarities.

Memory models have had considerable recent influence on research in behavioral economics. For example, Mullainathan (2002) proposes an associative-memory mechanism by which current news triggers selective memories of past events. The model generates biased forecasts based upon the degree of selectivity. Gennaioli and Shleifer (2010) demonstrate that a bounded-memory heuristic leads to "local thinking" that reduces the state space from which probabilistic beliefs are formed. Bordalo et. al. (2017) show how an associative-memory framework can explain a variety of behavioral biases. In an experimental setting, Enke et. al. (2019) test key implications of associative memory for imperfect forecasting and verify the key hypothesis that associative memory can induce overreaction to news.

Vector space representations of text similar to those used in this paper, are based on ensembles of contextual associations of words extracted from large text corpuses. They have been studied as a computational analog to human judgement processes, which rely on contextual associations of words and ideas. Bhatia (2017), for example, demonstrates that vector space models, can mimic documented prediction fallacies used in judgement under uncertainty. Wachter and Kahana (2019) show how a model based on contextual association and retrieval can shape expectations and influence asset prices. Goetzmann, Watanabe & Watanabe (2022) use a vector space representation to test the Wachter and Kahana model on a panel of macroeconomic forecasters and find that the contextual similarity of macroeconomic conditions to past states significantly predicted forecasts of the stock market.

One of the challenges to testing the asset pricing implications of memory models is that they describe individual behavior, but asset prices are driven by aggregate supply and demand. In memory models, current news triggers access to personal memory, which in turn affects beliefs. This indexing process need not be limited to personal recall. One need not have lived through the 1929 stock market crash to "recall" it via its contextual and semantic similarity to contemporary economic events, particularly if the news points out the similarities. Reference to a common narrative corpus like the news can address this challenge. While personally-experienced past events may be more salient, events "experienced" by hearing about them or reading about them can be subject to the same process of accessing past events via contextual association. In this case, reading about or hearing a story about an event substitute for personally experiencing it and thus can still potentially affect prediction via the same proposed recall mechanisms.

The balance of the paper is organized as follows. Section 2 describes the data and provides details on the methodology. Section 3 presents the results for the crash attention and market outcome tests. Section 4 presents the tests based on the survey data. Section 5 describes the methodology for the pure narrativity measure and presents the results of the narrativity tests. Section 6 concludes.

2. Data and Measure Construction

2.1 Data Sources

We use ProQuest to search the Eastern Edition of the *Wall Street Journal* [WSJ] from January 1, 1987 through December 31, 2020. This is the only edition available on ProQuest for that period. We assume that it corresponds reasonably well to the national edition.[15] We searched articles containing words and phrases associated with the stock market, yielding a total of 189,921 articles with a word count of at least 200 words.[16] We download the full text of these articles, their headlines, lede paragraphs and, for a sub-set, topic tags.

We also use ProQuest to download images of articles that appear in the WSJ from October 15, 1929 through October 31, 1929. The data available are digitized images of the main section. All articles that appear in the main section are manually transcribed. This process yields the full text and headlines for a total of 663 articles published during this period.

For stock market data, we use daily data on the value-weighted index of the NYSE-AMEX-Nasdaq-Arca universe from the Center for Research in Security Prices (CRSP) from January 2, 1987 through December 31, 2020. The daily returns of each index are used to empirically measure market volatility and the occurrence of extreme events. Market volatility implied by the OEX and VIX is obtained from the Chicago Board Options Exchange from January 2, 1987 through December 31, 2020. Daily, closing values for OEX are used through December 29, 1989, after which the VIX is used for the remainder of the sample period. Option-implied crash probabilities based on the S&P 500 are constructed using the methodology of Martin (2017) from January 4, 1996 through December 31, 2020. Specifically, options price data are collected from OptionMetrics for the SPDR S&P 500 ETF Trust.

In order to capture investor attention, we collect daily data on internet search volumes related to stock market crashes from January, 2004 to December, 2020. Da, Engleberg and Gao (2011) argue that search volume measures the attention of individual investors. In our application of this measure, search volumes may be indicative of individual investor attention to crashes,

[15] Data for the 1987 through 2015 period were collected in the weeks of May 17 through June 11 of 2016. Data for the 2016 through 2020 period were collected in the weeks of July 11 through December 5 of 2021.
[16] The Proquest search term used to identify the articles is: "(stock NEAR/5 market) OR SU(stock) OR SU(securities)". We did not use broader search terms, such "SU(markets)", because they yielded articles on other asset markets, such as for bonds and commodities.

which, in turn, may influence beliefs in the likelihood of a stock market crash.[17] Individuals may search for general terms related to a stock market crash if they are concerned about broader, market-wide trends as opposed to idiosyncratic price movements of individual stocks. We construct an analogous measure of general stock market attention which is also used in the analysis.

Robert Shiller's Stock Market Confidence Indices are based on survey data collected continuously since 1989; semi-annually for a decade and then monthly by the International Center for Finance at the Yale School of Management since July, 2001. Shiller (2000a) describes the indices constructed from these surveys and compares them to other sentiment indicators and studies their dynamics in the aggregate. In this paper, we use the disaggregated survey responses used to construct those indices. About 300 questionnaires each month are mailed to individuals identified by a market survey firm as high-net-worth investors and institutional investors. They may fill it in when they wish, but they are asked to mark the date on which they complete the survey. It is not a longitudinal survey. Each month comprises a different sample of respondents with the sampling goal of 20 to 50 responses by each of the two types – individual and institutional. The combined sample used in this paper contains 13,555 responses.

The current study focuses on responses to the survey question:

"What do you think is the probability of a catastrophic stock market crash in the U. S., like that of October 28, 1929 or October 19, 1987, in the next six months, including the case that a crash occurred in the other countries and spreads to the U.S.? (An answer of 0% means that it cannot happen, an answer of 100% means it is sure to happen.)

Probability in U. S.: ______________%"

The phrasing of this question has not been significantly altered during the sample period we examine.[18] Thus it has the advantage of consistency, during which time the stock market, the

[17] We download daily search volume indices using the Google Trends API. We restrict searches to the finance category. Because we cannot request daily data for the full sample period as calls to the API using daily frequency are limited, we download overlapping, 9-month subsamples that span the full period. The values of the search index are specific to each subsample, and so need to be adjusted in order to make them comparable to other subsamples. We address this issue by adjusting the daily values using the monthly frequency data, which is not affected by this issue. We adjust the daily values based on monthly averages to match with the monthly values. The overlapping periods are then averaged, resulting in the time-consistent series.

[18] This wording has remained the same since 1994. Prior to 1994, the question is phrased as: "What do you think is the probability of a catastrophic stock market crash, like that of October 28, 1929 or October 19, 1987, in the next six months?" Only approximately 10% of the observations used in the analysis are associated with the earlier wording. The results are not sensitive to the exclusion of these observations.

macro-economy and the financial system has experienced considerable variation. Additional discussion and analysis related to the survey data can be found in Goetzmann, Kim & Shiller (2024).

The Stock Market Confidence Index questionnaires have room for additional commentary. A sub-set of respondents filled these in and we hand-keyed their written responses. This provides us with narrative data from individual and institutional investors themselves. Notice that the wording of the question is ideal for motivating recall or attention to a rare disaster because it specifically refers to the two largest one-day crashes in U.S. capital market history. Thus, even without any influence from the financial press, we might expect contextual retrieval of knowledge about those two events.

For the construction of folk motifs, we use a classification of *Grimm's Fairy Tales* according to the Thompson Motif Index (TMI), a widely used classification system to attribute motifs to folktales.[19] The advantage of this system is that it allows for multiple motifs to be assigned to the same story. Moreover, the categories are quite granular though we only use the major categories for simplicity. The complete set of TMI attributions for the corpus comes from *TalesUnlimited.com*. We implemented a web scraper to collect the data.[20]

2.2 Crash Narrative Measure Construction

In this section we explain how we construct a crash narrative score. Because narratives involve sequences, settings and characters, they require a modeling tool that captures sequencing and pattern. The most commonly used natural language processing tools called "bag-of-words" methods classify documents by word frequencies alone.[21] They do not use sequencing or locational information within a document, except in-so-far as these leave traces in the types of words that appear.

[19] https://sites.ualberta.ca/~urban/Projects/English/Motif_Index.htm

[20] One potential concern is that Grimm's tales in translation may be noisy. Dehghani et. al. (2017) shows that neural responses to fairy tales is robust to translation. They also show that Doc2Vec representations of the stories allowed subjects to differentiate them based on neurological (fMRI) responses.

[21] Cf. Gentzkow et. al. (2019) for an overview of using text as data in economic research, including bag-of-words methods.

Recent advances in contextual modeling and computing power offer ways to exploit higher-order linguistic relationships within paragraphs, documents and corpora. Le and Mikolov (2014) introduce an approach called Doc2Vec that contextualizes words, paragraphs and documents via a machine-learning algorithm to do such things as predict missing words and classify documents according to higher-order semantics. More specifically, Doc2Vec is an algorithm that captures high-dimensional relationships between the textual units (e.g., words, sentences, paragraphs) by a text vectorization procedure that uses numeric representations of documents within a corpus. In particular, Doc2Vec does not process individual words in isolation but rather constructs context vectors for words and paragraphs in the corpus and uses machine learning to train a predictive algorithm on the multiple contexts in which textual units appear. Thus, it has the potential to do a better job at capturing overall meaning than non-contextual bag-of-words approaches. The class of algorithms that estimate semantic similarity through context vectors have been shown to be effective at higher level classification of news.[22]

Based on the Doc2Vec approach, we identify articles that have high semantic association with articles that were published shortly after the two largest one day crashes in U.S. stock market history: October 24, 1929 and October 19th, 1987. The idea is that articles with high similarity based on language used by the press around the 1929 and 1987 stock market crashes, regardless of whether the event was explicitly mentioned in the article, should prime the retrieval of collective and perhaps personal memories associated with the event. We focus on these events given that neither was anticipated, particularly the 1987 stock market crash. This ensures a relatively clean structural break in the narratives used by journalists before and after these events to ensure that our approach is capturing crash-related stock market narratives.

A key advantage of the approach is that it allows for the identification of narratives that mitigates look-ahead bias. By focusing on the similarity of articles published prior to sample period for which the survey data is available, it limits the influence of recent data that could have been used to inform the document embeddings and induce look-ahead bias. Moreover, though a decade old, Doc2Vec remains a robust tool for testing hypotheses about document similarity: its primary application. Contrast this approach to the use of querying a generative AI model, such as Chat GPT, to score the narrativity of an article. Responses provide no natural statistic for hypothesis

[22] Cf. Kim et. al. (2017) and Kabir et. al. (2019) for tests of doc2vec methods on news text.

testing. More importantly, most generative AI models have been trained using data that extends across unspecified text and time periods. Thus, properly segmenting the training data and training a new model is difficult to do given the computational resources required. In contrast, Doc2Vec models can be feasibly trained on computing systems available to most researchers. We further discuss these methodological issues later in the paper.

We estimate a Doc2Vec model using the WSJ corpus from 1987 and constructed a daily series of average similarity metrics.[23]

The following is a summary of the implementation:

1. Train the Doc2Vec model over all WSJ articles from 1987 to 2020. The vectorization of these articles provides the basis for the estimation of cosine similarity across articles or a daily corpus of articles.
2. For each article over the sample period, we calculate the average cosine similarity for each article to articles published in *The Wall Street Journal* between October 20 to October 23, 1987.
3. Vernacular and other structural aspects of writing styles are likely to have evolved over the sample period. That is, two articles may differ substantially due to factors unrelated to narrativity. We address this issue by adjusting the similarity calculations in (2) based on the average cosine similarity for each article to articles appearing in the WSJ from October 5, 1987 through October 9, 1987. The adjusted similarity is defined as the difference between the natural log of one plus the average cosine similarities.
4. For each day, we calculate the average adjusted similarity across articles, and refer to this quantity as *'87 Narrative*.

It is useful to consider an example of texts with high and low cosine similarity scores. Below we take an article from September 16th, 2008 and find the best and worst matches among articles published on October 20th, 1987 – the day following the 21% single-day crash. We note that none of the articles include the term, "crash."

[23] The models were estimated in Python using the *gensim* library. For the analysis presented, the model is trained over 100 iterations based on a feature vector size of 250, maximum distance of 5, minimum word frequency of 10 across all articles, sub-sampling threshold of 10^{-5}, and a negative sample of 5.

The following is the article from September 16, 2008.

> *If you learn nothing else from the last few harrowing days, you should learn the difference between what is obvious and what is inevitable.*
>
> *In the heat of the moment, the two perceptions seem identical. It was obvious that investors would panic as they absorbed the news about Bloody Sunday on Wall Street, so it was inevitable that the market would take a slashing. It was obvious that Lehman Brothers had to go bust, so a bankruptcy filing was inevitable. It was obvious that Merrill Lynch could no longer make it on its own, so it was inevitable that a bigger institution like Bank of America would take it over.*
>
> *But investors -- at least individual investors -- don't actually panic in times like these. Instead, they freeze. In July (the latest month for which final numbers are available), mutual-fund investors pulled out just $2.62 of every $100 they had invested in stock funds. That was less than they took out of bond funds, even though the stock market had just gone through a nauseating summer swoon.*

Of all articles published on October 20, 1987, the following is the article with the *highest* cosine similarity (29%).

> *The little investor doesn't know where to turn, although bigger ones are putting up a brave front.*
>
> *"I'm scared," says Julie Ianotti, an executive secretary in Houston. "My stock is my nest egg, for a house or something. Should I sell? Tell me, should I sell?"*
>
> *When Ms. Ianotti's mother suggested the price drop could portend a new depression, the young Ms. Ianotti at first scoffed at the idea. "I said, 'Yeah, sure, Mom. It'll go back up.'" But as the Dow Jones Industrial Average was taking a record plunge yesterday, Ms. Ianotti says she began giving more thought to her mother's stories about subsisting on water and sugar in the Depression.*

Of all articles published on October 20, 1987, the following is the article with the *lowest* cosine similarity (1%).

> *IC Industries Inc., expanding a restructuring of the company, said it will buy as much as $1 billion of its common shares and is considering selling its aerospace unit, valued at more than $1.5 billion.*
>
> *Karl D. Bays, chairman and chief executive officer, said the moves are being undertaken to focus the company around its most profitable businesses, which include food products, soft-drink bottling and auto repair. He said the steps aren't related to recent market reports that Minneapolis investors Irwin Jacobs and Carl Pohlad may seek control of the company.*

In September, IC announced plans to divest its Illinois Central Gulf Railroad unit by distributing shares in the division to IC holders. It estimated the railroad's market value at $250 million.

This example highlights the benefit of using a methodology of identifying higher-order document similarity. While all three articles have characters engaged in a sequence of actions in a setting, The best matched article feels more like a story; it has a greater sense of urgency.

2.3. Summary Statistics

Table 1 displays the variable descriptions used in the analysis and summary statistics.

2.4. Time-series Properties of the Crash Narratives Measure

Figure 2 displays the monthly averages of *'87 Narrative* (thick black) and other market-based crash indicators from January, 1987 through December, 2020. *'87 Narrative* is the cosine similarity of articles published on day *t* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987. The survey-based crash probabilities (red) are from the Shiller survey. The market-based crash indicators include the minimum daily market return for a given month (gray), the volatility of the daily market returns (light blue), the average daily option-based crash probabilities based on Martin (2017) (dark blue), and the average daily VIX (yellow). All measures are standard normalized.

'87 Narrative corresponds strongly with the other crash indicators. Notably the measure exhibits large spikes during exceptional periods of market stress, identified in the text overlayed on the figure, despite contextual differences from the 1987 crash. The seeming incongruity suggests that journalists may employ crash narratives irrespective of specific causes. The correlations based on monthly frequency vary between 37.5% for VIX to 39.9% for market returns volatility. The correlation with the survey-based crash probabilities is 46.8%. Interestingly, there are periods where *'87 Narrative* diverges from the market-based indicators, particularly during the 2003-2006 and 2013-2019 periods. These periods are associated with abnormally low volatility conditions. They may reflect concerns expressed through other channels that are in turn conveyed by financial media, of elevated market risks in spite of low realized volatility.

We next examine the correlation structure between the crash narrative measure and lagged market-based crash indicators based on the daily data. This analysis only uses data for trading days as some of the variables are unavailable for non-trading days. The results are displayed in Panel A of Table 2. The calculations exclude data prior to 1988. The correlations with the lagged volatility measures range between 26.2% to 28.9%, while the correlation with lagged market returns is negative: -8.7%. There is also a significant, positive correlation with lagged market turnover, or 33.8%. Correlations based on 30-day changes in the variables are somewhat weaker though consistent with the results in Panel A. Panel B of Table 2 displays the results. The correlations for changes in the volatility measures range between 11.5% to 25.2%. The correlation for the 30-day cumulative market returns is -2.3%, while the correlation for changes in market turnover is 9.8%. These results indicate that the crash narrative measure has a strong correspondence with market-based risk, return, fear and disagreement indicators based on the unconditional correlations.

Table 3 displays the regression analysis results. Again, the calculations exclude data prior to 1988. The regressions include controls for two-way fixed effects to account for day-of-week and month-of-year factors. Columns (1) through (5) display the regression coefficients when including each market-based indicator individually, and are consistent with the unconditional correlation results. Column (6) includes all the measures in the same model. All of the measures are statistically significant at least at the 1% level. The high degree of collinearity between the measures likely affects some of the estimates, most noticeably the option-implied crash probability coefficient.

Given that the sample spans more than 30 years following the 1987 crash, we perform subperiod analysis to assess how the associations between the crash narrative measure and the market-based indicators change over time. In untabulated results, we find that the associations remain significant throughout the full sample period. Moreover, R^2 is relatively similar across the subperiods. This suggests that the semantic structure, possibly the types of narratives that prevailed during the 1987 crash continued to be used by the financial press, even into the most recent period.

3. Crash Attention and Market Dynamics

In this section, we examine whether crash narratives are salient to investors. If so, do they affect aggregate market outcomes? We define *CrashAttention* as the daily search volume index for the search term "stock market crash". We chose this term given that it is highly correlated with alternative, related search terms.[24]

Figure 3 displays the monthly values of *CrashAttention* (green) and *'87 Narrative* (black). After the large spike in *CrashAttention* in 2008, there were subsequent spikes, particularly in the 2015 through 2020 sample period. This suggests that crash concerns can arise even during periods of market advance. *CrashAttention* is positively correlated with *'87 Narratives*: 14.2% at the daily frequency. It is also significantly correlated with market-based crash indicators, ranging from 9.5% for option-implied crash probabilities to 13.8% for market returns squared, at the daily frequency. Notice the *CrashAttention* measure is episodic. Attention to crashes spiked dramatically only about six or seven times during the fifteen-year period, but was "out-of-sight, out-of-mind" for long periods. This suggests that attention is a potentially important conditioning variable for belief formation.

To assess the predictive power of *'87 Narrative* on *CrashAttention*, we estimate the following regression model:

$$\begin{aligned} CrashAttention_t = {} & \alpha_1 \times \text{'87 } Narrative_{t-1} + \Sigma_{j=1}^{5} \boldsymbol{\alpha_{2,j}} \times \boldsymbol{X_{t-j}} \\ & + \Sigma_{j=1}^{5} \alpha_{3,j} \times CrashAttention_{t-j} + \boldsymbol{\phi^{DOW}_t} + \boldsymbol{\phi^{Month}_t} + \varepsilon_t. \end{aligned} \tag{1}$$

We control for the first five lags of *CrashAttention* to mitigate any spurious associations due to serial correlation in *CrashAttention*. Because *'87 Narrative* is also correlated with market-based crash indicators, we control for the following in (**X**): VIX, market returns, market returns squared, and market turnover. Most recent values for the controls are used for non-trading days. The first five lags of these variables are included in the model. We control for day-of-the week and month effects by including two-way fixed effects in the model. Finally, we use Newey-West standard

[24] These search terms include "market crash," "stock crash," "1987 crash," "1929 crash," and others.

errors to account for any remaining serial correlation in the residuals that may bias the test statistics.

The results support the hypothesis that attention to the stock market may be a conditioning factor on the effects of crash narratives. If investors' attention to the market is not activated, they are unlikely to seek out news about the market and, in turn, react to it.[25] As such, we expect the effects of crash narratives to be stronger during periods of high attention to the stock market. We construct a measure of general stock market attention, *StockMarketAttention*, using the daily search volume index for a set of search terms related to the stock market.[26] We augment Equation (1) with *StockMarketAttention* and an interaction term with *'87 Narrative* in the following manner:

$$\begin{aligned} CrashAttention_t = {} & \beta_1 \times \text{'87 } Narrative_{t-1} + \beta_2 \times StockMarketAttention_{t-1} \\ & + \beta_3 \times \text{'87 } Narrative_{t-1} \times StockMarketAttention_{t-1} \\ & + \Sigma_{j=1}^{5} \boldsymbol{\beta_{4,j}} \times \boldsymbol{X_{t-j}} + \Sigma_{j=1}^{5} \beta_{5,j} \times CrashAttention_{t-j} + \boldsymbol{\kappa}^{\mathbf{DOW}}_t + \\ & \boldsymbol{\kappa}^{\mathbf{Month}}_{\mathbf{t}} + \eta_t. \end{aligned} \tag{2}$$

These tests allow us to assess whether the effects of crash narratives are larger during periods of high investor attention via the coefficient on interaction term is positive ($\beta_3 > 0$).

3.1 Results

Table 4 displays the results. Column (1) presents the results from univariate tests. The coefficient for *'87 Narrative* is positive and statistically significant at the 1% level. The specification for Column (2) includes the interaction terms with *StockMarketAttention*. The coefficients on both the un-interacted and interacted. *'87 Narrative* terms are positive and statistically significant at the 1% level. The coefficient for the un-interacted *StockMarketAttention* term is also statistically significant. The results suggest that the media crash narratives are predictive of investor attention to stock market crashes, and the effects are stronger during periods

[25] Cf. Bybee et. al. (2021).

[26] Specifically, we restrict searches to the finance category, and is based on the union of following search terms: "djia", "dow jones", "nasdaq", "new york stock exchange", "nyse", "stocks", "s&p", and "sp500".

where investors are already paying attention to stock market news. Column (3) repeats the tests in Column (2) but also includes controls for fundamental conditions when crash concerns are likely to be higher. The results are qualitatively similar and remain significant, though the point estimates attenuate. To control for the strong right-skew of *CrashAttention* we examine an alternative specification that maps *CrashAttention* to a dummy, where non-zero values are coded as one and zero otherwise, in order to determine whether distributional asymmetry is driving significance. Column (4) displays the results; they are consistent with those of Column (3).

The types of narratives used by financial media have likely evolved throughout time. Older narratives may be unfamiliar to modern audiences, and so may lack salience. They may also be less salient because they are simply removed in time. We directly evaluate this conjecture by repeating the analysis in Table 4 using the crash narrative measure based on the October, 1929 stock market crash, or *'29 Narrative*.

The 1987 and 1929 crash narrative measures exhibit a strong positive association. Figure 4 displays the results for monthly averages. The correlation between the two measures is 49.2% at a monthly frequency, and 48.2% at a daily frequency. This suggests a common, long-lived component in narratives used in articles published in 1929 and 1987. There are also some interesting differences. *'29 Narrative* displays a more episodic pattern compared to *'87 Narrative.* It spiked after the Dot Com bubble and the 2008 financial crisis as with the *'87 Narrative* measure but is lower during the period in between the two and in the period since 2016.

Table A1 displays the results for a similar specification to Table 4. The coefficients for *'29 Narrative* are positive and statistically significant in the specifications without the control variables, though two to three times smaller compared to those for *'87 Narrative* in Table 4. Additionally, the interaction term is statistically insignificant at the 10% level. With the inclusion of the control variables, the coefficients for the *'29 Narrative* terms become insignificant. In other words, the explanatory power of *'29 Narrative* appears to be highly correlated with fundamental factors.

The different effects of the *'29 Narrative* and the *'87 Narrative* on crash attention are interesting in that the former preceded the Great Depression and the latter had little knock-on effects. The market mostly recovered by the end of the year and the Dot-Com bubble burst more than two years later. One might expect that 1929 would attract the attention of those concerned about the economic consequences of a crash. The *'87 Narrative* in the years that followed may be a story of disaster and recovery. Differences in narrative effects may reflect differences in outcomes. Additionally, the 1929 crash may have been anticipated by some investors. There were warning signs in the period leading up to the event, some of which coincided with signs of market instability as late as September 1929.[27] Given that our strategy assumes that the period prior to the crash is relatively calm, it would suggest that the *'29 Narrative* measure is not fully capturing crash-related narratives used to describe the 1929 crash.

3.2 Market Outcomes

We next examine whether crash narratives affect market volatility and returns. The results from the crash attention tests suggest that financial media narratives are salient. However, it is unclear whether the effects are sufficiently powerful to affect market dynamics.

To evaluate, we directly examine associations for market volatility and returns with lagged values of *'87 Narratives*. As with the earlier tests, we account for serial correlations and correlations with other market-based indicators that may generate spurious correlations with *'87 Narratives*. Given that the outcome variables are only available on trading days, non-trading days are excluded for these tests. Additionally, we examine how the associations differ when accounting for conditional distributions. Given that stock market crashes are generally associated with large, unexpected spikes in market volatility, we use quantile regressions to evaluate the effects at different points of the distribution.

We find that crash narratives have a strong positive association with volatility the following day, and the effect diminishes over the subsequent days. Table 5 presents the results. Column (1) displays the results for the OLS regressions. *'87 Narrative* has a positive effect the day prior, though the effects reverse in subsequent days. We next consider the effects of crash narratives on the conditional volatility distribution. Again, we conjecture that the effects should be most pronounced at higher levels. Columns (2), (3) and (4) display the results at the 50th, 75th and 90th percentiles, respectively. Across these columns, the initial effect of *'87 Narratives* become larger,

---

[27] Most notably, Roger Babson's famous forecast of a crash on September 8, 1929, citing speculation and the economic cycle c.f. "Babson Predicts Crash in Stocks Sooner or Later". *The Owensboro Messenger* (Owensboro, Kentucky). September 8, 1929. p. 2.

and the effects persist much longer. These results suggest that crash narratives may amplify the effects of volatile periods.

We next examine the effects on market returns and find no significant association. Table 6 presents the results and is formatted similarly to Table 5. Columns (2), (3), and (4) present the quantile regression results and are associated with specifications at the 50th, 25th, and 10th percentiles, respectively. Across the specifications, the coefficients for *'87 Narrative* is statistically insignificant at the 10% for the day prior, and mostly insignificant for earlier days. The results indicate that there are no significant effects on market returns.[28]

4. Investor Beliefs

The results of the previous section suggest that crash narratives impact investor crash concerns and market dynamics. Those tests rely on a battery of controls in order to mitigate omitted variable issues that could bias the tests. However, these controls may not be sufficient to capture other fundamental factors impounded in market prices. In addition, the outcomes examined do not necessarily correspond with investor beliefs. Time-variation in internet search volumes could reflect non-investors as well.

We address these issues by performing tests on investor survey data. The data allow us to directly examine investor beliefs, and utilize cross-sectional variation based on investor types. The survey responses include crash probability assessments that are directly comparable to estimates based on stock options data. We are able to identify two types of investors: individual and institutional. Because both sets of investors observe the same market conditions, any narrative effects associated with changes in fundamentals should be reflected in both. However, there is a large literature in behavioral finance suggesting that investor sophistication has an important role in mediating susceptibility of psychological biases. In particular, individuals may rely more on judgement heuristics when domain expertise is low (Ottati & Isbell, 1996; Sedikides, 1995). As

[28] While the crash narrativity metric is not designed as a volatility index, its predictive power for VIX suggests that a comparison to the results of Manela and Moreira (2017) are relevant. They construct a VIX proxy, NVIX by projecting VIX onto the space of financial news. NVIX significantly predicts market returns over longer horizons, with the power increasing for multiple-year intervals. We focus on daily return prediction.

such, we expect differential effects for individual investors, or investors with relatively poorer sophistication, allowing us to cleanly identify the effects of narratives.

For the outcome variable, we adjust the investor crash probabilities using information from the stock options market. We employ the methodology of Martin (2017) to estimate market-implied probabilities of a 15% decline in the S&P 500 over the next six months, which conforms to the parameters provided in the survey question. The survey data is adjusted using the option-implied crash probabilities, or π^{Adj}, by using the difference in log values based on information as of the survey response date.

Figure 5 displays the monthly averages of the crash probabilities for individual (red) and institutional (yellow) investors from the Shiller survey. Several observations stand out. First, the survey-based crash probabilities correspond closely with the option-implied crash probabilities (blue). The correlations between the individual (institutional) investor and option-implied crash probabilities are 47.1% (49.4%). Second, the survey-based crash probabilities are significantly larger than those based on options data, as noted in Goetzmann, Kim & Shiller (2024). Based on the respondent-level data, the difference is 14.1% (t-value = 66.55) and 12.1% (t-value = 57.40) for individual and institutional investors, respectively. Third, the individual investor crash probabilities are significantly higher than those of institutional investors based on parametric (t-value = 6.84) as well as non-parametric (z-value = 4.51) tests.

We estimate the following regression model separately for individual and institutional investors:

$$\begin{aligned}\pi^{Adj}{}_{i,t} = {} & \gamma_1 \times \text{'87 } Narrative_{t-1} + \gamma_2 \times StockMarketAttention_{t-1} \\ & + \gamma_3 \times \text{'87 } Narrative_{t-1} \times StockMarketAttention_{t-1} \\ & + \boldsymbol{\gamma_4} \times \boldsymbol{W_{i,t-1}} + \boldsymbol{\lambda}^{\mathbf{DOW}}{}_t + \boldsymbol{\lambda}^{\mathbf{Year \times Month}}{}_{\mathbf{t}} + \zeta_{i,t}.\end{aligned} \tag{3}$$

The control variables (**W**) include previous-day market returns, past 30-day market returns, past 30-day market returns volatility, the average survey-based crash probabilities over the past 30-days, and the previous-day VIX. In order to distinguish immediate effects associated with narratives, we control for past 30-day averages of *'87 Narrative* and *StockMarketAttention*. In addition to these controls, we also include two-fixed effects of the day-of-week and year-month levels. Robust standard errors clustered at the date level are used to calculate the test statistics.

We focus on the interaction term coefficient (γ_3) and compare the estimates for the individual and institutional investor samples. For direct comparisons, we also present estimates from a pooled model augmented with interaction terms on a dummy associated with investor type.

4.1. Results

We begin with basic results using a model estimated on the non-nested corpus. With this data, we find evidence that crash narratives have a significant, positive relation to investor crash assessments. Table 7 displays the results. Columns (1) and (2) only include the *'87 Narrative* and *StockMarketAttention* terms for the individual and institutional investor subsamples, respectively. For the individual investor subsample, the coefficient for un-interacted *'87 Narrative* is positive and statistically significant. The coefficient for the interaction term is positive and statistically significant. The results suggest that individual investor crash probability assessments are higher in general when more attention is paid to the market, and also when heightened attention to the market coincides with higher crash narrativity in financial media. In contrast, for institutional investors, the *'87 Narrative* interaction terms are statistically insignificant.

The differences in outcomes for individual versus institutional investors suggest that the effects of crash narratives that we detect may not be simply driven by fundamental conditions that coincide with the prevalence of narratives in financial media. To further investigate, we alter the specifications in Columns (1) and (2) in the following manner. First, the dependent variable is replaced with the difference between the survey crash probability and the option-implied crash probability survey responses. This adjustment takes advantage of the presumption that the prices of index options impound value-relevant, fundamental information about market crash probabilities. To the extent that out of the money options could reflect unrealistic beliefs, for example irrational crash concerns unrelated to media narrative, this is a conservative approach. Second, control variables are included that account for other market information. Columns (3) and (4) display the results for individual and institutional investors, respectively. The results for individual investors remain similar. While the coefficient for the un-interacted *'87 Narrative* terms becomes statistically insignificant, the interaction terms remain positive and statistically significant. The results are unchanged for the institutional investor subsample after including the control variables. Column (5) repeats the analysis using the pooled sample. The results are virtually identical.

4.2 Look-ahead Bias

While the methodology to this point mitigates some potential look-ahead biases, there may nonetheless be concerns given that the model is estimated over the full sample. In this section, we create a bias-free measure and compare it to the baseline measure. As a preview, we find little evidence that, in the context of our analysis, the results using the baseline methodology are sensitive look-ahead biases.

To construct a bias-free measure, we start by re-estimating the model for each period to ensure that no future information is used in the model training. Specifically, for date t, a separate model is estimated using only information up the previous month. For example, for June 2, 2005, we estimate a model based on articles published between January 1984 through May 2005 and then apply that model to articles published on June 2. We repeat this procedure for all articles published between January 2000 through December 2020. We start in January 2000 given that the survey data become more frequently available around that period. The same parameterizations described for the baseline model are used for these estimations. A total of 252 models are estimated using this procedure, from which the time-series of the bias-free measure is calculated.

Figure A1 plots the bias-free *'87 Narrative* measure alongside the baseline measure. The two measures are standardized for comparability. Visually, the two measures appear to be highly correlated, and there are few instances where there is divergence between the two. The Pearson correlation between the monthly measure values is 94.5%. At a daily frequency, the correlation is 85.5%.

One explanation for the high correlations is the stability in the model weights. To illustrate, Figure A2 displays the top 25 most similar terms to “crash” and their similarity scores from 2000 through 2020. While there is some fluctuation in the similarity scores, they remain at a relatively high level throughout the sample period. For reference, Table A2 displays the top 25 terms related to “crash” as of December 2020.

We next revisit the main results and assess how they differ using the bias-free measure. Table A3 displays the results. The results are similar to the baseline results. They provide evidence that the main results are unlikely to be due to look ahead bias.

Additionally, we are able to directly evaluate the magnitude of the bias and whether it has any effect. To do so, we decompose the baseline measure into bias-free components using a regression of the baseline measure onto the third-order polynomial transformation of the bias-free measure. We interpret the residual from this regression as the bias component, which we refer to as *Bias Residual*. We alter the main tests to use the *Bias Residual* in place of the *'87 Narrative* measure.

Table A4 displays the results. None of the *Bias Residual* components are statistically significant at the 10% level. This suggests that the results are unlikely to be affected significantly by look-ahead biases. This provides support that the methodology used to construct the crash narrativity measure is relatively robust look-ahead biases. These results are of course not generalizable to other kinds of machine-learning tests – particularly tests of market return predictability. Even the non-temporally restricted Doc2Vec model limits the leakage problem to some extent by controlling the source of embeddings -- restricting them to Wall Street Journal text. However it is reasonable to assume that future semantic associations will affect cosine similarities. The nesting approach mitigates this problem -- figure A2 demonstrates that the stability of associations is not an artifact of look-ahead bias.

4.3. Older Narratives

We next turn our attention to examining the effect of older narratives. To this end, we compare the results when using crash narratives based on the 1929 articles. Table A5 presents the results and is formatted similarly to Table 7. As with the crash attention tests, the coefficients for the *'29 Narrative* terms are mostly statistically insignificant across all the specifications, and none of the interaction terms are statistically significant. However, as pointed out above, this may also be due to differences in the perceived outcomes of the two crashes as well as other issues.

4.4. Term Frequency

As we discussed above, term frequency, not narrative structure, might be driving the results. We construct another measure that isolates the frequency of terms used in articles following the 1987 crash. This allows us to explicitly test against a null hypothesis that the structural relationship among words is irrelevant. *'87 Term Frequency* is the difference in cosine similarities between an article published on date *t* and articles published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987 based on term frequency-inverse document frequency. We use a weighted bag-of-words approach given that previous studies demonstrate their superior performance to unweighted approaches.

Table A6 presents the results and is formatted similarly to Table 7. The coefficients for the *'87 Term Frequency* interaction terms are statistically insignificant across all the specifications. Interestingly, the un-interacted term coefficient is positive and statistically significant for the individual investor subsample. In untabulated results, when augmenting the baseline models from Table 7 with *'87 Term Frequency* terms, the results remain significant and the coefficients are similar for the *'87 Narrative* terms. This suggests that term frequency by itself is not driving the results. It also suggests that methods using bag-of-words approaches like LDA-based topic models may have at least some power to capture narrative relationships.

4.5. Intra-article Salience

We next use an alternative approach to examine the role of salience. The tests rely on *StockMarketAttention* as a proxy for salience. However, it is to some extent correlated with market-based crash indicators, making the results difficult to interpret. The correlations are generally higher at a monthly rather than daily frequency, which our analysis focuses on. We use an alternative approach to study the role of salience that focuses on intra-article variation in narratives. Specifically, we use differences in information contained in the lede paragraph—the first 100 words of the article—versus other paragraphs. We assume that the information content is similar across the entire article, though readers are more likely to focus on the lede paragraph. Moreover, lede paragraphs in some cases include a dramatic thesis to engage readers to continue reading the article. Non-lede paragraphs may provide details supportive of this thesis. As such, non-lede paragraphs are more likely to capture fundamental conditions.

Because the machine learning techniques used to construct the narrative measure performs relatively poorly on shorter documents, we extract narrative information in the lede paragraph in the following manner. We start by constructing an analogous narrative measure, or *'87 Narrative*$^{Non\text{-}Lede}$ that excludes the first 100 words of each article. The average word count of articles in our sample is 622.9, and so we believe excluding the first 100 words will minimize degradation due to small document sizes. We then regress *'87 Narrative* onto the second-polynomial expansion of *'87 Narrative*$^{Non\text{-}Lede}$, then define *'87 Narrative*^{Lede} as the residual from this regression.

Table 8 displays the results. Columns (1) through (3) display the results using only *'87 Narrative*^{Lede} as the narrative measure for the individual investor subsample, institutional investor subsample, and pooled sample specifications, respectively. The results are quantitatively similar to those in Table 7. Because *'87 Narrative*^{Lede} is orthogonal to fundamental information that may be captured by non-lede paragraphs, by construction, the stability in the coefficients suggest that most of the effects are unlikely to be drive by potentially unobservable fundamental factors. Finally, we augment the model with *'87 Narrative*$^{Non\text{-}Lede}$ terms to evaluate the sources of the main effects based on the full decomposition. Columns (4) through (6) display the results and are formatted similarly to Columns (1) through (3). We show that the effects are concentrated in the *'87 Narrative*^{Lede} terms while the *'87 Narrative*$^{Non\text{-}Lede}$ terms are statistically insignificant.

4.6. Investor Narratives

The results from the previous section provides evidence that crash narratives propagated by the financial press affect on investor beliefs. While they may be most relevant for crash assessments, crash narratives may have broader effects than what are captured in those tests.

We next examine how crash narratives may influence the types of narratives investors use to describe their general beliefs about the stock market. We do so by constructing tests based on free responses that Shiller survey respondents are asked to provide. Specifically, investors are asked for general comments that may include an explanation for their responses to all the questions, including the crash assessment. We take the text responses and map them into document vectors based on the same model used for the WSJ articles. The document vectors are then used to estimate the similarity between the survey comments and newspaper articles that were published around the 1987 stock market crash. We call this measure *InvestorNarrative*.

We find evidence that crash narratives that appear in the financial press in the day prior are significantly related to individual comments about the market. We do so by examining lead-lag relationships between *'87 Narrative* and *InvestorNarrative*. Table 9 displays the results. Columns (1) and (2) are based on the individual investor subsample, while Columns (3) and (4) are based on the institutional investor subsample. Columns (2) and (4) include the full set of fixed effects and control variables used in Table 7. The table shows a positive and significant effect of *'87 Narrative* on *InvestorNarrative* the following day, and there is mostly no significant effect on other days. Moreover, these effects are only found in the individual investor subsample. All of the coefficients are statistically insignificant for the institutional investor subsample.

5. Folk Motifs and Crash Narratives

When is the financial press more likely to use higher narrativity? What is the relationship between *'87 Narrative* and general narrativity?

We consider narrativity using a measure that is unrelated to the stock market, financial press, and even the modern era. Using a similar approach to *'87 Narrative*, we calculate similarities in the semantic relationships between WSJ articles and classic folktales published more than a century prior. Specifically, we use the complete corpus of the Tales of the Brothers Grimm Volumes 1 and 2, first published in 1812. The collection is based on oral tradition of known fairy tales in German-speaking countries at the time and is cited as one of the founding texts of Western culture.[29] The version we obtain is an updated English translation by Margaret Hunt in 1884 (Grimm et al., 2014). There are 210 stories across both volumes. We use the Doc2Vec model trained on the WSJ corpus to populate the feature vectors for each of the stories. The feature vector is then organized based on the TMI motif categories of each story. For each WSJ article, cosine similarities are calculated for all motif categories.[30] We define *FolkMotif* as the average of the first

[29] Cf. Auden (1944).
[30] Cf. Würzbach (1997) for a discussion of motifs.

principal component of these cosine similarities across all articles published on date *t*. We argue that *FolkMotif* captures the pure narrativity of articles, allowing us to evaluate the prevalence of narratives more generally.

Figure 6 displays monthly averages of *FolkMotifs*. *'87 Narratives* is also displayed for comparison. Both series are standard normalized. The correlation between the two measures is high—83.4% at a monthly frequency and 65.9% at a daily frequency. The high degree of correlation affirms our interpretation of *'87 Narratives*, namely that it is likely capturing narrativity rather than only terms associated with adverse fundamental conditions. The figure also suggests that, historically, the financial press is more likely to use high narrativity during periods of extreme market stress and uncertainty. Interestingly, as with *'87 Narratives*, the narrativity of the financial press has increased in the post-2015 sample period. This period also coincides with greater prevalence of volatility spikes, raising questions about whether greater narrativity in media accounts may be influencing investor behavior, or whether higher narrativity levels are reflective of other dynamics.

We formally examine whether narrativity corresponds with *'87 Narrative*. We adapt the model from Equation (2) by using *FolkMotif* as the dependent variable. All specifications include the first five lags of *FolkMotif* to account for serial correlations and restrict variations to its innovations. Table 10 displays the results. Column (1) displays the regression model results with only *'87 Narrative*, showing a positive and significant relationship. Column (2) features the interaction terms with *StockMarketAttention*. Both the interacted and un-interacted *'87 Narrative* term coefficients are positive and statistically significant, suggesting that the narrativity of *'87 Crash* increases during periods of high investor attention. The un-interacted *StockMarketAttention* term coefficient is also positive and significant. This suggests that the financial press is more likely to use narrativity in newspaper articles when investor attention is higher, independent of whether they use crash narratives. One potential explanation is that the financial press serves broader readership during such periods, and journalists may rely on narrativity to make financial news more accessible.

Column (3) displays the results for the market-based crash indicators, excluding the *'87 Narrative* terms. It presents the sum of the coefficients over the lagged terms. *p*-values from the χ^2 statistics on the sum of the control variable coefficients are also shown. All but market turnover is statistically insignificant. This is notable as it suggests that the financial press does not generally use high narrativity to describe market conditions *per se*, although the significance of turnover, which is indicative of investor activity may suggest narrative-related trading. Column (4) displays the full model. Here, the results on the *'87 Narrative* terms are similar, and even stronger, than those of Column (2). Moreover, all the control variable coefficients become statistically insignificant at the 10% level. This provides evidence that the effects of crash narratives that we document in prior sections are unlikely to be due to unobservable fundamental heterogeneity. It also supports the validity of *FolkMotif* as an instrument for narrativity.

We revisit the tests in Equation (3) by replacing the *'87 Narrative* terms with *FolkMotif*. This allows us to directly assess the effects of narrativity on investor beliefs. Table 11 shows the results. We find an analogous effect in the individual investor subsample. Column (1) shows that the *FolkMotif* interaction term is positive and statistically significant. We isolate components in *FolkMotif* related to crash narratives using a regression model that projects *FolkMotif* on the second-order polynomial transformation of *'87 Narrative*. This allows us to directly compare the effects of crash-related versus non-crash-related narratives. We find that only the component associated with the crash narrative measure is statistically significant. As with our other results, these results are only significant for the individual investor subsample and suggest that the only forms of narratives that impact investor crash concerns are those associated with stock market crashes.

6. Conclusion

In this paper we explore the relationship of media and investor narratives to investor attention and the formation of crash expectations at times when the probability of a major market crash is higher. We find that attention to the dynamics of the stock market is episodic and varies with measures of market uncertainty. We test the hypothesis that in such periods, interest in narratives is also high – perhaps because stories are fundamentally concerned with state changes intermediated by events. In such times, investors undoubtedly look to the financial press to understand the causes and consequences of market dynamics. We show that, in these circumstances, history plays an important role in press narratives by bringing attention to past "rare disasters" in U.S. stock market history. In this sense, it is a collective analogue to the powerful models of individual memory explored by recent research as a filter in decision-making under uncertainty. Collective and individual memory processes are not mutually exclusive – indeed stories may also be a mechanism to perpetuate a particularly important event in individual memory.

We focus on the narrative aspect of this recall by constructing a crash narrative similarity score for the daily flow of stock market news in *The Wall Street Journal* since 1987. The variable interacts with attention to the market to explain individual predictions of a catastrophic stock market crash – but not so with institutional investor predictions, which are presumably better-informed opinions or perhaps less susceptible to influence by reference to (or recall of) a limited number of events in U.S. capital market history.

We address the challenge of separating narrative topics from "narrativity" in a couple of ways – first by the use of a natural language processing approach that captures higher-order relationships across documents, and second by using text in which topics are entirely orthogonal to financial themes. These folk motifs derive from archetypal tales that many people experience as children. We find that when public attention turns to the stock market, the resonance between financial news and fairy tales predicts higher individual crash forecasts. As the tale of the Andaman Islanders suggests, such behavior is not necessarily irrational. Stories can encode wisdom about a rare disaster in a memorable narrative.

References

Andersen, T.G., Fusari, N. and Todorov, V., 2020. The pricing of tail risk and the equity premium: Evidence from international option markets. *Journal of Business & Economic Statistics*, *38*(3), pp.662-678.

Auden, W.H. 1944. In praise of the Brothers Grimm; Margaret Hunt's Revised Translation Offers New Delight to Child and Adult. *New York Times*. Nov. 12, 1944.

Barberis, N., 2013. The psychology of tail events: progress and challenges. *American Economic Review*, 103(3), pp.611-16.

Bertsch, C., Hull, I. and Zhang, X., 2021. Narrative fragmentation and the business cycle. *Economics Letters*, *201*, p.109783.

Bhatia, S., 2017. Associative judgment and vector space semantics. *Psychological review*, *124*(1), p.1.

Bollerslev, T. and Todorov, V., 2011. Tails, fears, and risk premia. *The Journal of Finance*, *66*(6), pp.2165-2211.

Bordalo, P., Conlon, J.J., Gennaioli, N., Kwon, S.Y. and Shleifer, A., 2021. *Memory and probability* (No. w29273). National Bureau of Economic Research.

Bordalo, P., Gennaioli, N. and Shleifer, A., 2012. Salience theory of choice under risk. *The Quarterly journal of economics*, *127*(3), pp.1243-1285.

Bordalo, P., Gennaioli, N. and Shleifer, A., 2013. Salience and asset prices. *American Economic Review*, *103*(3), pp.623-28.

Bordalo, P., Gennaioli, N. and Shleifer, A., 2017. Memory, attention, and choice. *The Quarterly journal of economics*.

Bruner, J., 1991. The narrative construction of reality. *Critical inquiry*, *18*(1), pp.1-21.

Bybee, J.L., 2023. The ghost in the machine: Generating beliefs with large language models. *arXiv preprint arXiv:2305.02823*.

Chahrour, R., Nimark, K. and Pitschner, S., 2021. Sectoral media focus and aggregate fluctuations. *American Economic Review*, *111*(12), pp.3872-3922.

Da, Z., Engelberg, J. and Gao, P., 2011. In search of attention. *The journal of finance*, *66*(5), pp.1461-1499.

Da, Z., Engelberg, J. and Gao, P., 2015. The sum of all FEARS investor sentiment and asset prices. *The Review of Financial Studies*, *28*(1), pp.1-32.

Dehghani, M., Boghrati, R., Man, K., Hoover, J., Gimbel, S.I., Vaswani, A., Zevin, J.D., Immordino-Yang, M.H., Gordon, A.S., Damasio, A. and Kaplan, J.T., 2017. Decoding the neural representation of story meanings across languages. *Human brain mapping*, *38*(12), pp.6096-6106.

Dierckx, T., Davis, J. and Schoutens, W., 2021. Quantifying News Narratives to Predict Movements in Market Risk. *Data Science for Economics and Finance*, p.265.

Enke, B., Schwerter, F. and Zimmermann, F., 2020. Associative memory and belief formation (No. w26664). National Bureau of Economic Research.

Flynn, J.P. and Sastry, K., 2022. The macroeconomics of narratives. *Available at SSRN 4140751*.Garry, J., 2017. *Archetypes and Motifs in Folklore and Literature: A Handbook: A Handbook*. Routledge.

Gennaioli, N. and Shleifer, A., 2010. What comes to mind. *The Quarterly journal of economics*, *125*(4), pp.1399-1433.

Gentzkow, M., Kelly, B. and Taddy, M., 2019. Text as data. *Journal of Economic Literature*, *57*(3), pp.535-74.

Glasserman, P. and Lin, C., 2023. Assessing look-ahead bias in stock return predictions generated by gpt sentiment analysis. *arXiv preprint arXiv:2309.17322*.

Greifeneder, R., & Bless, H. 2007. Relying on accessible content versus accessibility experiences: The case of processing capacity. *Social Cognition*,*25*: 853-881.

Grimm, Jacob, Grimm, Wilhelm and Zipes, Jack. *The Original Folk and Fairy Tales of the Brothers Grimm: The Complete First Edition*, Princeton: Princeton University Press, 2014.

Halawi, Danny, Fred Zhang, Chen Yueh-Han, and Jacob Steinhardt, "Approaching Human-Level Forecasting with Language Models.

Jha, M., Liu, H. and Manela, A., 2022. Does finance benefit society? A language embedding approach. *A Language Embedding Approach (June 1, 2022)*.

Kabir, K.L., Alam, F.F. and Islam, A.B., 2020. Word embeddings for semantic resemblance of substantial text data: A comparative study. In *Smart Systems and IoT: Innovations in Computing* (pp. 303-312). Springer, Singapore.

Kahana, M.J., 2012. *Foundations of human memory*. OUP USA.

Kim, H.K., Kim, H., Cho, S.: Bag-of-concepts: comprehending document representation through clustering words in distributed representation. Neurocomputing.

Larsen, V.H., Thorsrud, L.A. and Zhulanova, J., 2021. News-driven inflation expectations and information rigidities. *Journal of Monetary Economics*, *117*, pp.507-520.

Lopez-Lira, A. and Tang, Y., 2023. Can chatgpt forecast stock price movements? return predictability and large language models. *arXiv preprint arXiv:2304.07619*.

Manela, A. and Moreira, A., 2017. News implied volatility and disaster concerns. *Journal of Financial Economics*, *123*(1), pp.137-162.

Martin, I., 2017. What is the Expected Return on the Market?. *The Quarterly Journal of Economics*, *132*(1), pp.367-433.

Morag, D. and Loewenstein, G., 2021. Narratives and Valuations. *Available at SSRN 3919471*.

Mullainathan, S., 2002. A memory-based model of bounded rationality. *The Quarterly Journal of Economics*, *117*(3), pp.735-774.

Ottati, V. C., & Isbell, L. M. 1996. Effects on mood during exposure to target information on subsequently reported judgments: An on-line model of misattribution and correction. *Journal of Personality and Social Psychology*,*71*: 39-53.

Pandya, V., 2005. " When Land Became Water": Tsunami and the Ongees of Little Andaman Island. *Anthropology News*, *46*(3), pp.12-13.

Piper, A., Bagga, S., Monteiro, L., Yang, A., Labrosse, M. and Liu, Y.L., 2021. Detecting Narrativity Across Long Time Scales. *Proceedings http://ceur-ws. org ISSN*, *1613*, p.0073.

Propp, V.I., 1968. *Morphology of the Folktale* (Vol. 9). University of Texas Press.

Sarkar, S.K. and Vafa, K., 2024. Lookahead bias in pretrained language models. *Available at SSRN*.

Sedikides, C. 1995. Central and peripheral self-conceptions are differentially influenced by mood: Tests of the differential sensitivity hypothesis. *Journal of Personality and Social Psychology*, *69*: 759-777.

Siemer, M., & Reisenzein, R. 1998. Effects of mood on evaluative judgments: Influence of reduced processing capacity and mood salience. *Cognition & Emotion*, *12*: 783-805.

Shiller, R.J., 2017. Narrative economics. *American economic review*, *107*(4), pp.967-1004.

Thompson, S. 1989 [1955-1958; 1932-1936]. Motif-Index of Folk Literature. Bloomington: Indiana University Press.

Tversky, A., & Kahneman, D. 1983. Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90, 293–315.

van Binsbergen, J.H., Hua, S. and Wachter, J.A., 2020. Is the United States a lucky survivor: A hierarchical Bayesian approach. *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper*.

Wachter, J.A., 2013. Can time-varying risk of rare disasters explain aggregate stock market volatility? *The Journal of Finance*, *68*(3), pp.987-1035.

Wachter, J.A. and Kahana, M.J., 2019. *A retrieved-context theory of financial decisions* (No. w26200). National Bureau of Economic Research.

Welch, I., 2016. The (time-varying) importance of disaster risk. *Financial Analysts Journal*, *72*(5), pp.14-30.

Würzbach, N., 1997. Theory and practice of compiling a motif index, with the child corpus as example. *Journal of folklore research*, *34*(3), pp.251-258.

Figure 1

Article Counts Citing 1987 and 1929 Crashes

This figure displays monthly counts of articles that reference the 1987 and 1929 stock market crashes.

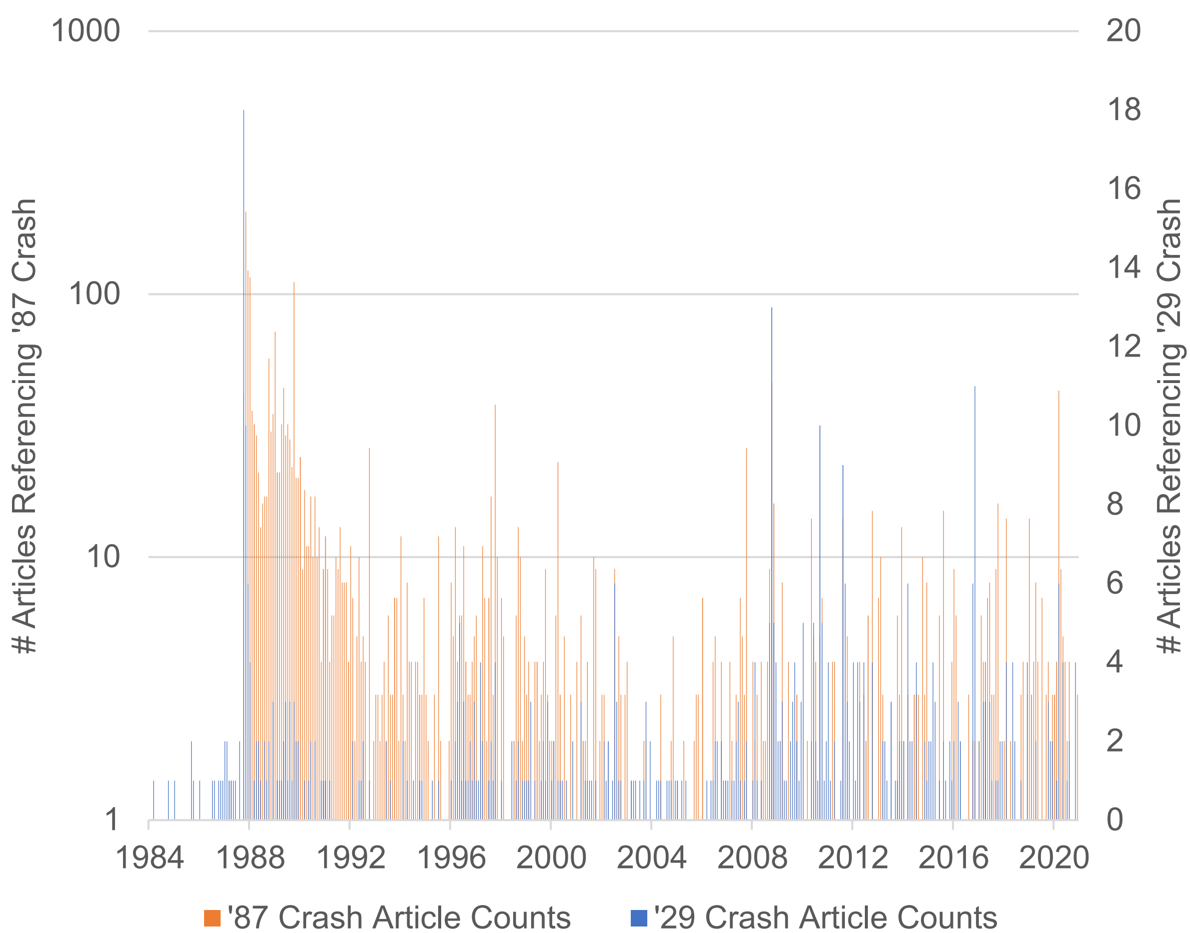

Figure 2

Crash Narratives and Market-based Crash Indicators

The figure displays monthly averages of *'87 Narrative*, survey-based crash probabilities, and other market-based crash indicators from January 1987 through December 2020. *'87 Narrative* is the cosine similarity of articles published on day *t* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987. The survey-based crash probabilities are from the Shiller survey. The market-based crash indicators include: minimum daily market return for a given month, the volatility of the daily market returns, the average daily option-based crash probabilities, and the average daily VIX. All measures are standard normalized.

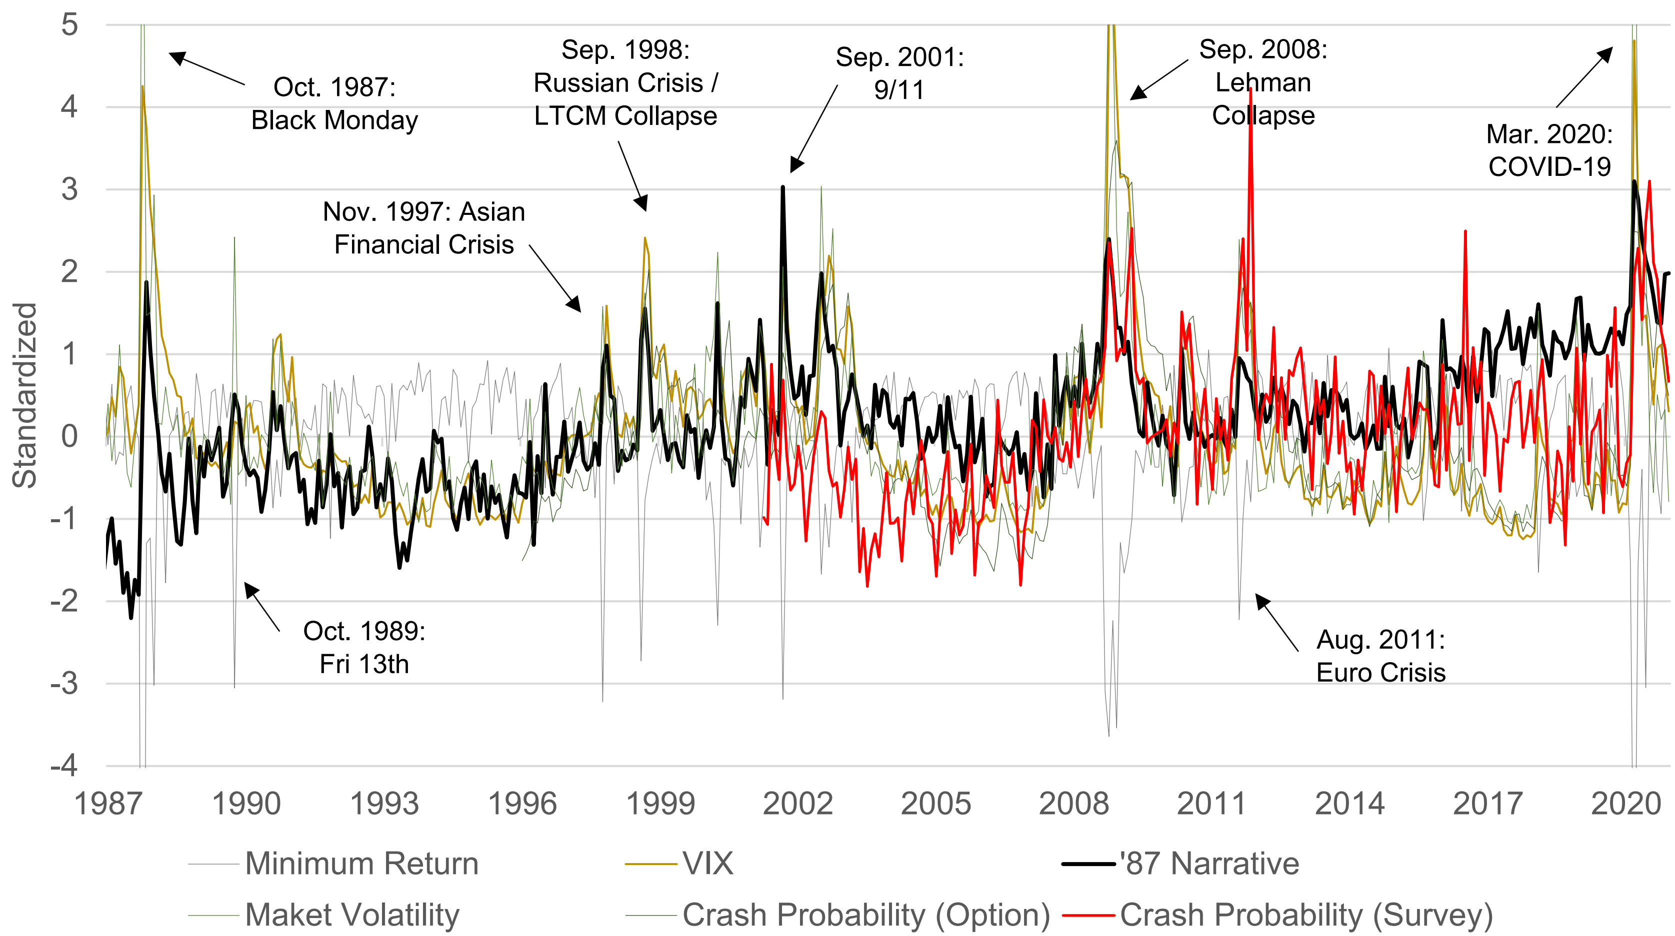

Figure 3

Crash Attention

The figure displays monthly averages of *'87 Narrative* and *Crash Attention* from January 2005 through December 2020. *'87 Narrative* is the cosine similarity of articles published on day *t* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987. *Crash Attention* is the Google Search Index for the search term "stock market crash." All measures are standard normalized.

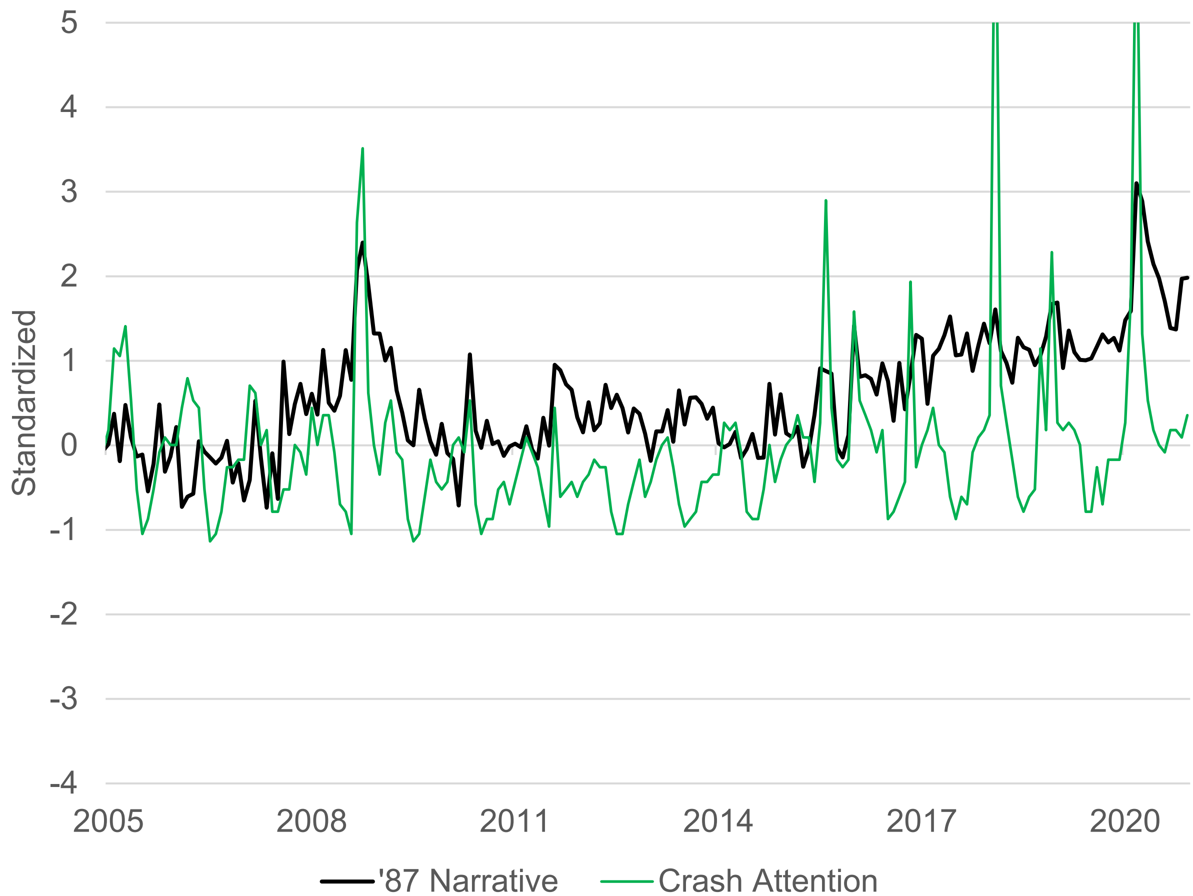

Figure 4

1929 Crash Narratives

The figure displays monthly averages of *'87 Narrative* and *'29 Narrative* from January 1987 through December 2020. *'87 Narrative* is the cosine similarity of articles published on day *t* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987. *'29 Narrative* is the cosine similarity of articles published on day *t* with those published between October 24 – November 8, 1929 minus the cosine similarity with those published between October 7 – 23, 1929. All measures are standard normalized.

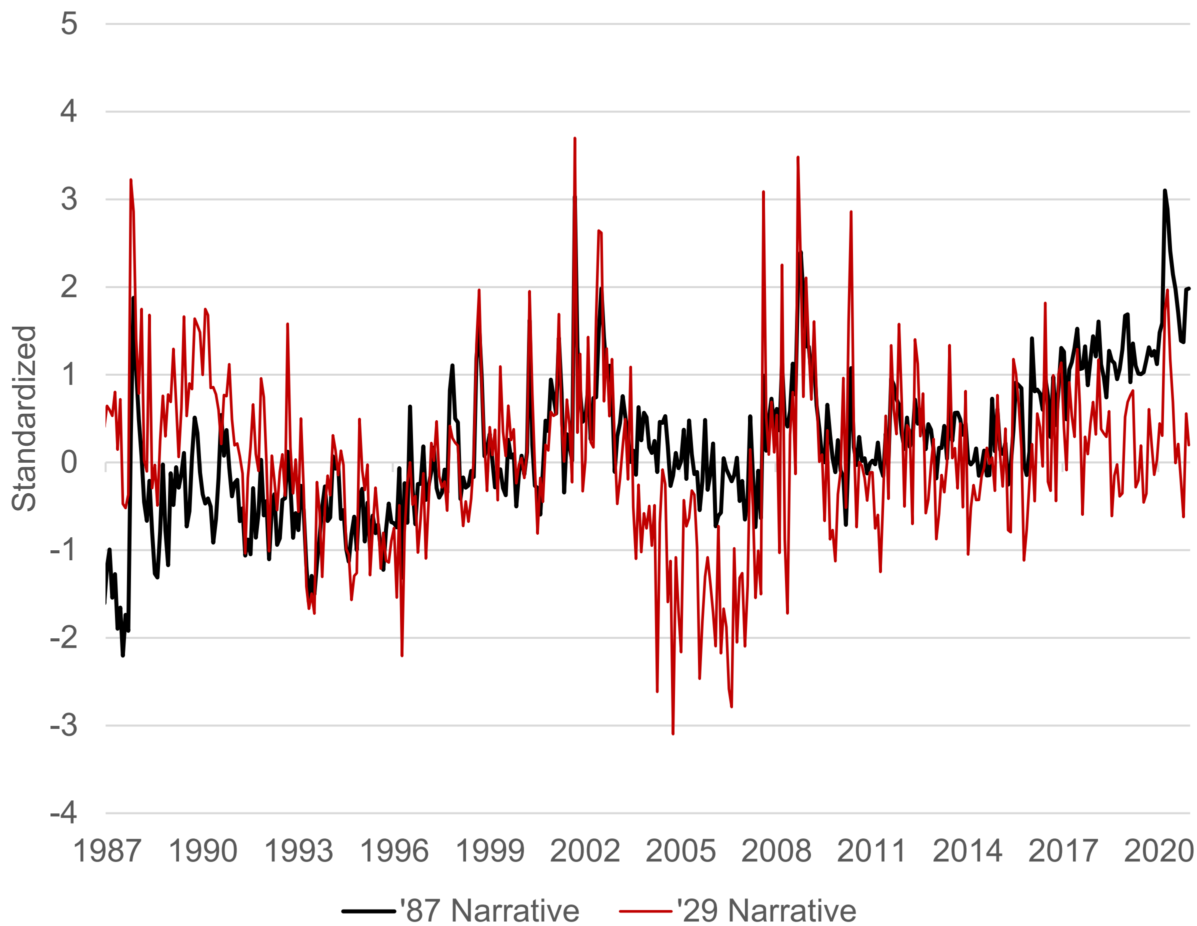

Figure 5

Survey-based Crash Probabilities

The figure displays the survey-based crash probabilities based on individual and institutional investor subsamples as well as the option-implied crash probabilities from January 1996 through December 2020.

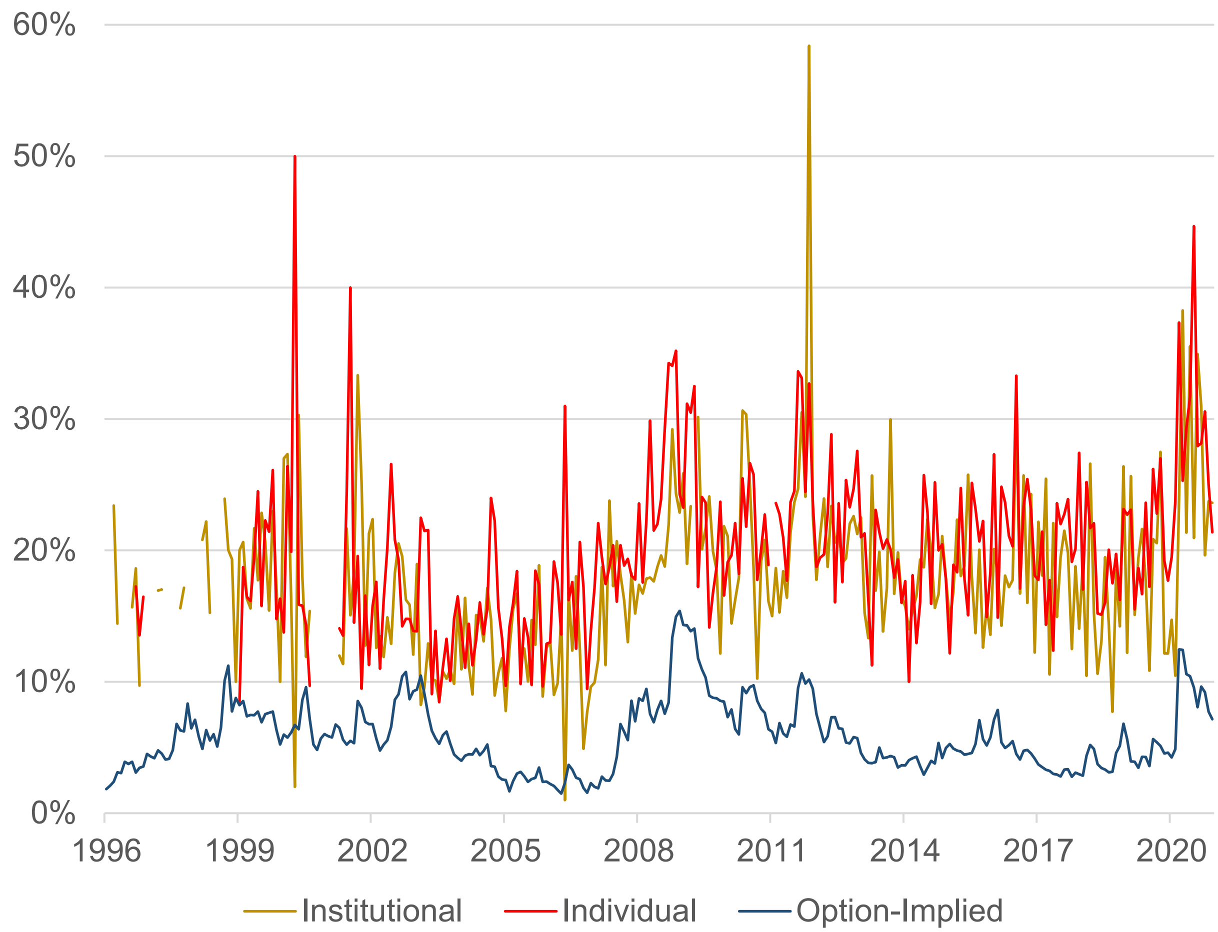

Figure 6

Folk Motifs

The figure displays monthly averages of *'87 Narrative* and *Folk Motif* from January 1987 through December 2020. *'87 Narrative* is the cosine similarity of articles published on day *t* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987. *FolkMotif* is the first principal component of the cosine similarities of articles published on day *t* with folktales by motif categories. All measures are standard normalized.

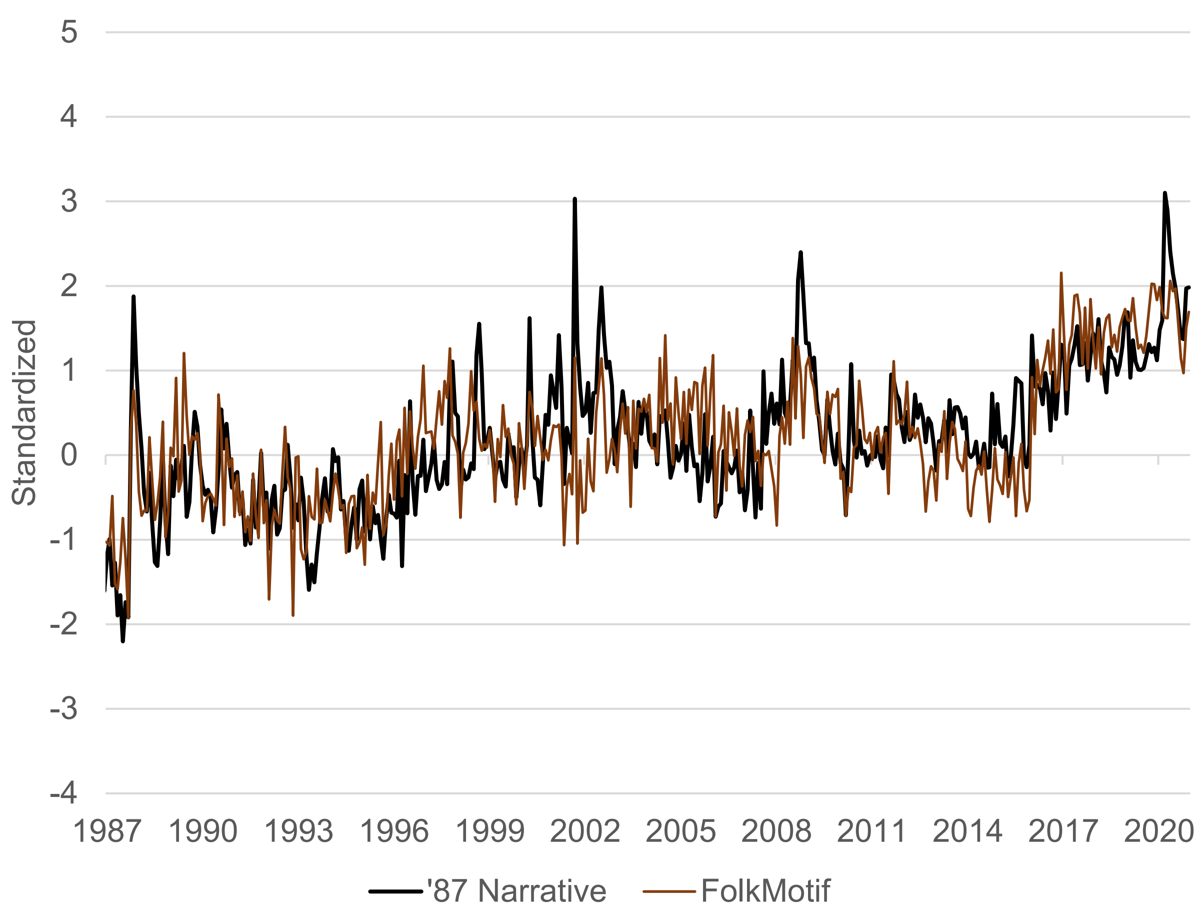

Table 1

Summary Statistics

The table displays summary statistics for variables for the full data (Panel A) and the survey data (Panel B) tests.

Panel A: Full Data					
	Mean	Standard Deviation	25th Pctl.	50th Pctl.	75th Pctl.
'87 Narrative	0.484	0.675	0.028	0.508	0.826
'29 Narrative	-0.065	0.850	-0.504	-0.116	0.444
StockMarketAttention	7.215	1.703	5.926	7.538	8.300
CrashAttention	0.976	1.365	0.000	0.000	2.079
VIX	19.375	8.633	13.350	16.580	22.810
π^{Option}	0.056	0.028	0.036	0.052	0.070
R_M	0.027	0.978	-0.349	0.049	0.517
$(R_M)^2$	0.011	0.025	0.000	0.002	0.009
Turnover	2.318	0.312	2.110	2.278	2.527
FolkMotif	0.382	0.956	-0.288	0.286	1.105

Panel B: Survey Data					
	Mean	Standard Deviation	25th Pctl.	50th Pctl.	75th Pctl.
π	0.167	0.153	0.049	0.095	0.223
π^{Adj}	0.109	0.152	-0.004	0.062	0.183
'87 Narrative	0.575	0.754	0.057	0.525	0.973
'29 Narrative	-0.018	0.977	-0.538	-0.039	0.538
StockMarketAttention	6.986	1.761	5.331	7.315	8.213
R_M	0.001	0.011	-0.004	0.001	0.006
Past Month Market Returns	0.011	0.058	-0.011	0.019	0.045
Market Return Volatility	0.010	0.007	0.006	0.008	0.012
Past Month π	0.200	0.055	0.165	0.196	0.232
VIX	19.538	8.992	13.490	17.110	22.500
FolkMotif	0.360	1.033	-0.306	0.246	0.944

Table 2

Correlation Matrix

The table displays Pearson correlation coefficients for *'87 Narrative* and market-based crash indicators. Panel A displays correlations using previous day market factors. Panel B displays correlations based on the one-month changes in the variables.

Panel A: Pearson Correlations on *'87 Narrative$_{t+1}$*

	(1)	(2)	(3)	(4)	(5)
	'87 Narrative$_{t+1}$	VIX_t	π^{Option}_t	R^M_t	$(R^M_t)^2$
VIX_t	28.94%				
π^{Option}_t	26.23%	85.26%			
R^M_t	-8.72%	1.89%	-7.04%		
$(R^M_t)^2$	28.14%	43.78%	41.72%	-5.49%	
$Turnover_t$	33.83%	18.29%	35.59%	-2.04%	26.11%

Panel B: Pearson Correlations on 30-day Changes

	(1)	(2)	(3)	(4)	(5)
	Δ'87 Narrative$_t$	ΔVIX_t	$\Delta\pi^{Option}_t$	$R^M_{t-30,t}$	$\Delta(R^M_t)^2$
ΔVIX_t	25.20%				
$\Delta\pi^{Option}_t$	15.88%	54.89%			
$R^M_{t-30,t}$	-2.31%	-9.50%	-4.17%		
$\Delta(R^M_t)^2$	11.54%	26.47%	21.96%	-3.85%	
$\Delta Turnover_t$	9.75%	22.67%	17.00%	2.65%	31.66%

Table 3

Crash Narratives and Market-based Crash Indicators

The dependent variables are *'87 Narrative*, defined as the cosine similarity of articles published on day *t+1* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987. Only data from January 1988 are used for these tests. Specifications that include the option-implied crash probabilities only use data from January 1996 due to data availability. Two-way fixed effects for month and day-of-week are included, but not reported. Newey-West standard errors are reported in parentheses. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable:	'87 $Narrative_{t+1}$	'87 $Narrative_{t+1}$	'87 $Narrative_{t+1}$	'87 $Narrative_{t+1}$	'87 $Narrative_{t+1}$	'87 $Narrative_{t+1}$
VIX_t	0.253***					0.171***
	(0.013)					(0.018)
π^{Option}_t		0.051***				-0.018***
		(0.004)				(0.005)
R^M_t			-0.042***			-0.032***
			(0.008)			(0.006)
$(R^M_t)^2$				7.071***		2.347***
				(0.377)		(0.340)
$Turnover_t$					0.005***	0.004***
					(0.001)	(0.001)
Month and Day-of-week FEs	YES	YES	YES	YES	YES	YES
N	12,784	9,858	12,784	12,784	12,784	9,858
R^2	14.0%	13.3%	6.6%	10.9%	22.6%	21.3%

Table 4

Crash Narratives and Attention

The dependent variables are *Crash Attention*, or the Google search volume index for "stock market crash," and *Crash Attention (Dummy)*, or the dummy coded as one on days where the search volume index for "stock market crash" is non-zero, and zero otherwise. The main explanatory variables include: *'87 Narrative* is the cosine similarity of articles published on day *t* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987; *StockMarketAttention* is the Google search volume index for stock market-related terms on day *t*; and the interaction between *'87 Narrative* and *StockMarketAttention*. The control variables are included where indicated, but not reported. They include values for days *t-4* through *t* for the *Crash Attention,* VIX, option-implied crash probabilities, turnover, market returns, and market returns squared. Two-way fixed effects for month and day-of-week are included where indicated, but not reported. Newey-West standard errors are reported in parentheses. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Dependent Variable:	Crash Attention$_{t+1}$	Crash Attention$_{t+1}$	Crash Attention$_{t+1}$	Crash Attention (Dummy)$_{t+1}$
'87 Narrative$_t$	8.790***	5.739**	5.618**	6.971***
	(2.446)	(2.709)	(2.847)	(1.032)
StockMarketAttention$_t$		0.124***	0.104***	0.137***
		(0.015)	(0.016)	(0.007)
'87 Narrative$_t$ × StockMarketAttention$_t$		4.334***	4.515***	2.078***
		(1.192)	(1.216)	(0.440)
Lagged *Crash Attention* Terms	YES	YES	YES	YES
Control Variables	NO	NO	YES	YES
Month and Day-of-week FEs	NO	NO	YES	YES
N	6,268	6,268	6,264	6,265
R^2	19.9%	21.0%	22.3%	28.2%

Table 5

Market Volatility

The dependent variables are the daily VIX. The main explanatory variable is *'87 Narrative,* defined as the cosine similarity of articles published on day *t* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987. Column (1) display the results for the OLS regressions while Columns (2) – (4) display the results for the quantile regressions. The quantile used for Columns (2), (3), and (4) are the 50th, 75th and 90th percentiles, respectively. The control variables include values for days *t-5* through *t* for the VIX, turnover, market returns, and market returns squared. Two-way fixed effects for month and day-of-week are included where indicated, but not reported. Newey-West standard errors are reported in parentheses in Column (1), and robust standard errors are reported in parentheses in Columns (2) – (4). Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Estimator:	OLS	Q(0.5)	Q(0.75)	Q(0.9)
Dependent Variable:	VIX_{t+1}	VIX_{t+1}	VIX_{t+1}	VIX_{t+1}
'87 $Narrative_t$	0.005**	0.005**	0.009***	0.012***
	(0.002)	(0.002)	(0.003)	(0.004)
'87 $Narrative_{t-1}$	-0.005**	-0.005**	-0.004	-0.002
	(0.002)	(0.002)	(0.003)	(0.004)
'87 $Narrative_{t-2}$	-0.002	-0.001	0.000	-0.004
	(0.002)	(0.002)	(0.002)	(0.004)
'87 $Narrative_{t-3}$	0.003	0.002	0.004	0.011***
	(0.002)	(0.002)	(0.003)	(0.004)
'87 $Narrative_{t-4}$	-0.004**	-0.006**	-0.005**	-0.008***
	(0.002)	(0.003)	(0.002)	(0.003)
Control Variables	YES	YES	YES	YES
Month and Day-of-week FEs	YES	YES	YES	YES
N	4,275	4,275	4,275	4,275

Table 6

Market Returns

The dependent variables are daily market returns. The main explanatory variable is *'87 Narrative,* defined as the cosine similarity of articles published on day *t* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987. Column (1) display the results for the OLS regressions while Columns (2) – (4) display the results for the quantile regressions. The quantile used for Columns (2), (3), and (4) are the 50th, 75th and 90th percentiles, respectively. The control variables include values for days *t-5* through *t* for the VIX, turnover, market returns, and market returns squared. Two-way fixed effects for month and day-of-week are included where indicated, but not reported. Newey-West standard errors are reported in parentheses in Column (1), and robust standard errors are reported in parentheses in Columns (2) – (4). Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Estimator:	OLS	Q(0.5)	Q(0.25)	Q(0.1)
Dependent Variable:	R^M_{t+1}	R^M_{t+1}	R^M_{t+1}	R^M_{t+1}
'87 Narrative$_t$	-0.017	-0.018	-0.026	-0.029
	(0.019)	(0.016)	(0.024)	(0.033)
'87 Narrative$_{t-1}$	-0.003	0.000	0.002	-0.013
	(0.019)	(0.017)	(0.023)	(0.033)
'87 Narrative$_{t-2}$	0.035*	0.020	0.029	0.017
	(0.020)	(0.016)	(0.022)	(0.032)
'87 Narrative$_{t-3}$	-0.002	-0.010	0.022	0.067**
	(0.019)	(0.016)	(0.021)	(0.030)
'87 Narrative$_{t-4}$	0.024	0.029**	0.024	0.001
	(0.020)	(0.014)	(0.025)	(0.033)
Control Variables	YES	YES	YES	YES
Month and Day-of-week FEs	YES	YES	YES	YES
N	4,275	4,275	4,275	4,275

Table 7

Investor Crash Probabilities

The dependent variables are the survey-based crash probabilities, both adjusted and unadjusted using option-implied crash probabilities. The analysis is performed using only the subsample of individual investors, institutional investors, or both, as indicated in the second row. The main explanatory variables include: *'87 Narrative* is the cosine similarity of articles published on day *t* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987; *StockMarketAttention* is the Google search volume index for stock market-related terms on day *t*; and the interaction between *'87 Narrative* and *StockMarketAttention*. The control variables include lagged values of the VIX, turnover, market returns, market returns squared, past 30-day market returns, past 30-day market returns volatility, the average survey-based crash probabilities over the past 30 days, the average *'87 Narrative* over the past 30 days, and the average *StockMarketAttention* over the past 30 days. Two-way fixed effects for month and day-of-week are included where indicated, but not reported. Robust standard errors clustered at the date level are reported in parentheses. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Investor Subsample:	Indiv.	Inst.	Indiv.	Inst.	All
Dependent Variable:	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$
'87 Narrative$_t$	1.529***	0.861***	-0.084	-0.257	0.231
	(0.306)	(0.308)	(0.394)	(0.382)	(0.347)
StockMarketAttention$_t$	0.006***	0.003***	0.001	0.006**	0.004**
	(0.001)	(0.001)	(0.003)	(0.003)	(0.002)
'87 Narrative$_t$ × StockMarketAttention$_t$	0.697***	0.074	0.504***	-0.133	0.566***
	(0.162)	(0.156)	(0.180)	(0.157)	(0.169)
Institutional$_i$					-0.025***
					(0.003)
Institutional$_i$ × '87 Narrative$_t$					-0.002
					(0.002)
Institutional$_i$ × StockMarketAttention$_t$					-0.768*
					(0.427)
Institutional$_i$ × "87 Narrative$_t$ × StockMarketAttention$_t$					-0.740***
					(0.215)
Control Variables	NO	NO	YES	YES	YES
Year-Month and Day-of-week FEs	NO	NO	YES	YES	YES
N	7,271	6,284	7,270	6,282	13,555
R^2	1.0%	0.3%	6.2%	6.9%	3.9%

Table 8

Intra-article Salience

The dependent variables are the survey-based crash probabilities adjusted using option-implied crash probabilities. The analysis is performed using only the subsample of individual investors, institutional investors, or both, as indicated in the second row. *'87 Narrative^{Lede}* is the residual cosine similarity of *'87 Narrative* when using only articles excluding the first 100 words, and *'87 Narrative$^{Non\text{-}Lede}$* is the predicted cosine similarity of *'87 Narrative* when using only articles excluding the first 100 words. The control variables are identical to those used in Table 7. Robust standard errors clustered at the date level are reported in parentheses. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Investor Subsample:	Indiv.	Inst.	All	Indiv.	Inst.	All
Dependent Variable:	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$
'87 Narrative$^{Lede}_t$	-0.268	-0.249	0.221	-0.350	-0.196	-0.108
	(0.428)	(0.442)	(0.411)	(0.442)	(0.454)	(0.436)
StockMarketAttention$_t$	0.002	0.006*	0.005*	0.002	0.006*	0.005*
	(0.004)	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)
'87 Narrative$^{Lede}_t$ × StockMarketAttention$_t$	0.462**	-0.136	0.562***	0.402**	-0.126	0.426**
	(0.192)	(0.170)	(0.192)	(0.187)	(0.183)	(0.187)
Institutional$_i$			-0.025***			-0.025***
			(0.004)			(0.004)
Institutional$_i$ × '87 Narrative$^{Lede}_t$			-0.002			-0.001
			(0.002)			(0.002)
Institutional$_i$ × StockMarketAttention$_t$			-0.845*			-0.230
			(0.511)			(0.511)
Institutional$_i$ × ''87 Narrative$^{Lede}_t$ × StockMarketAttention$_t$			-0.777***			-0.563**
			(0.259)			(0.241)

Table 8 (cont.)

'87 $\text{Narrative}^{NonLede}_t$				0.683	-0.940*	0.641
				(0.502)	(0.493)	(0.465)
'87 $\text{Narrative}^{NonLede}_t$ × $\text{StockMarketAttention}_t$				0.216	0.008	0.312
				(0.272)	(0.266)	(0.265)
Institutional_i × '87 $\text{Narrative}^{NonLede}_t$						-1.424**
						(0.601)
Institutional_i × "87 $\text{Narrative}^{NonLede}_t$ × $\text{StockMarketAttention}_t$						-0.365
						(0.401)
Control Variables	YES	YES	YES	YES	YES	YES
Year-Month and Day-of-week FEs	YES	YES	YES	YES	YES	YES
N	7,270	6,282	13,555	7,270	6,282	13,555
R^2	7.7%	8.4%	4.6%	7.7%	8.5%	4.7%

Table 9

Investor Narratives

The dependent variables are the cosine similarities between the *Investor Narrative* measure based on the free-response comments of the Shiller survey and the *'87 Narrative* based on Wall Street Journal articles published between October 20 – 23. The analysis is performed using only the subsample of individual investors or institutional investors, as indicated in the second row. The main explanatory variables are lead and lag values of *'87 Narrative*. The control variables are identical to those of Table 7. Robust standard errors clustered at the date level are reported in parentheses. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Investor Subsample:	Indiv.	Indiv.	Inst.	Inst.
Dependent Variable:	Investor $Narrative_{i,t}$	Investor $Narrative_{i,t}$	Investor $Narrative_{i,t}$	Investor $Narrative_{i,t}$
'87 $Narrative_{t+3}$	-0.028	-0.039	-0.018	-0.026
	(0.046)	(0.046)	(0.046)	(0.047)
'87 $Narrative_{t+2}$	0.055	0.046	0.035	0.032
	(0.043)	(0.043)	(0.049)	(0.049)
'87 $Narrative_{t+1}$	0.023	0.004	-0.010	-0.012
	(0.044)	(0.045)	(0.052)	(0.053)
'87 $Narrative_{t}$	0.012	-0.002	-0.009	-0.008
	(0.043)	(0.045)	(0.049)	(0.050)
'87 $Narrative_{t-1}$	0.126***	0.112***	-0.017	-0.027
	(0.039)	(0.041)	(0.034)	(0.035)
'87 $Narrative_{t-2}$	0.007	0.015	-0.008	0.002
	(0.039)	(0.040)	(0.038)	(0.039)
'87 $Narrative_{t-3}$	-0.037	-0.027	-0.013	-0.006
	(0.042)	(0.042)	(0.044)	(0.045)
Control Variables	NO	YES	NO	YES
Year-Month and Day-of-week FEs	NO	YES	NO	YES
N	4,411	4,408	3,246	3,232
R^2	4.7%	4.8%	4.4%	4.7%

Table 10

Folk Motifs and Crash Narratives

The dependent variables are *FolkMotif*, or the first principal component of the cosine similarities of articles published on day *t* with folktales by motif categories. The control variables include values for days *t-4* through *t* for the *FolkMotif,* VIX, option-implied crash probabilities, market returns, market returns squared, and turnover. The sum of the coefficients for each of the control variables is reported where indicated. Two-way fixed effects for month and day-of-week are included where indicated, but not reported. Newey-West standard errors are reported in parentheses. The *p*-values for the χ^2 statistics on the sum of the control variable coefficients are reported in brackets. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Dependent Variable:	$FolkMotif_{t+1}$	$FolkMotif_{t+1}$	$FolkMotif_{t+1}$	$FolkMotif_{t+1}$
‘87 $Narrative_t$	0.096***	0.066***		0.098***
	(0.017)	(0.019)		(0.021)
$StockMarketAttention_t$		0.066***		0.098***
		(0.016)		(0.020)
‘87 $Narrative_t$ x $StockMarketAttention_t$		0.073***		0.075***
		(0.013)		(0.013)
Lagged Option Crash Probability Coefficients			-0.033	-0.016
			[0.50]	[0.12]
Lagged VIX Coefficients			-0.014	-0.013
			[0.08]	[0.07]
Lagged Market Return Coefficients			-0.024	-0.057
			[0.25]	[1.49]
Lagged Market Return^2 Coefficients			0.012	-0.038
			[0.19]	[2.38]
Lagged Turnover Coefficients			0.062**	-0.011
			[5.27]	[0.13]
Lagged *FolkMotif* Controls	YES	YES	YES	YES
Month and Day-of-week FEs	NO	NO	YES	YES
N	4,275	4,275	4,275	4,275
R^2	16.6%	17.7%	15.2%	18.3%

Table 11

Folk Motifs and Investor Crash Probabilities

The dependent variables are the survey-based crash probabilities adjusted using option-implied crash probabilities. The analysis is performed using only the subsample of individual investors or institutional investors, as indicated in the second row. *FolkMotif* is the first principal component of the cosine similarities of articles published on day *t* with folktales by motif categories. $FolkMotif^{Crash}$ is the predicted values of a regression that projects *FolkMotif* on the second-order polynomial expansion of *'87 Narratives*, while $FolkMotif^{Non\text{-}crash}$ is the residual. The control variables are identical to those used in Table 7. Robust standard errors clustered at the date level are reported in parentheses. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Investor Subsample:	Indiv.	Inst.	Indiv.	Inst.
Dependent Variable:	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$
$FolkMotif_t$	0.002	-0.002		
	(0.003)	(0.003)		
$StockMarketAttention_t$	0.002	0.005		
	(0.004)	(0.003)		
$FolkMotif_t \times StockMarketAttention_t$	0.003**	-0.002		
	(0.001)	(0.001)		
$FolkMotif_t^{Crash}$			-0.002	-0.003
			(0.005)	(0.005)
$FolkMotif_t^{Crash} \times StockMarketAttention_t$			0.006**	-0.001
			(0.002)	(0.002)
$FolkMotif_t^{Non\text{-}crash}$			0.004	-0.002
			(0.004)	(0.003)
$FolkMotif_t^{Non\text{-}crash} \times StockMarketAttention_t$			0.001	-0.002
			(0.002)	(0.001)
Control Variables	YES	YES	YES	YES
Year-Month and Day-of-week FEs	YES	YES	YES	YES
N	7,270	6,282	7,270	6,282
R^2	7.59%	8.46%	7.69%	8.46%

Figure A1

Bias-free '87 Narrative Comparison

The figure displays monthly averages of *'87 Narrative* using a model trained on the full sample and the rolling estimation procedure described in the text. The rolling estimation procedure involves estimating a new model only using articles available up through the previous month. The '87 Narrative measure is then constructed using the bias-free model. All measures are standard normalized.

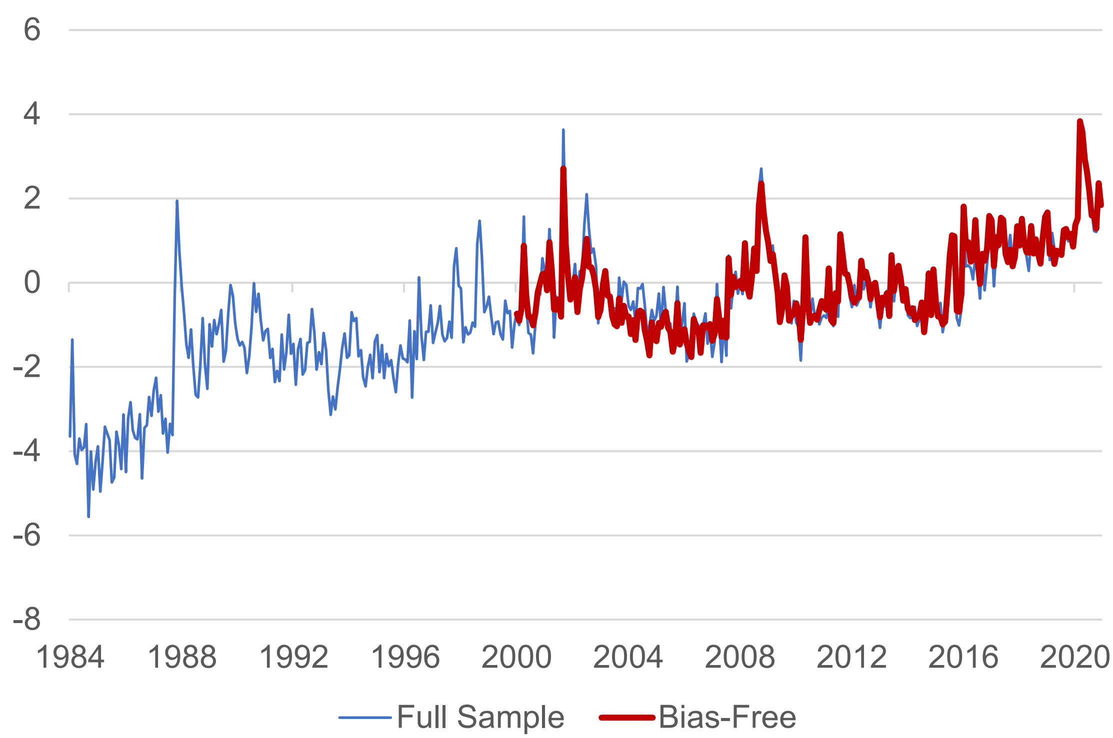

Figure A2

Crash Embeddings over Time

The figure shows the top 25 similarity scores for terms associated with the term "Crash" for every year from 2000 to 2020 using the rolling estimation procedure described in the text. Only the terms that do not explicitly include variations of the term "Crash" are included in this figure.

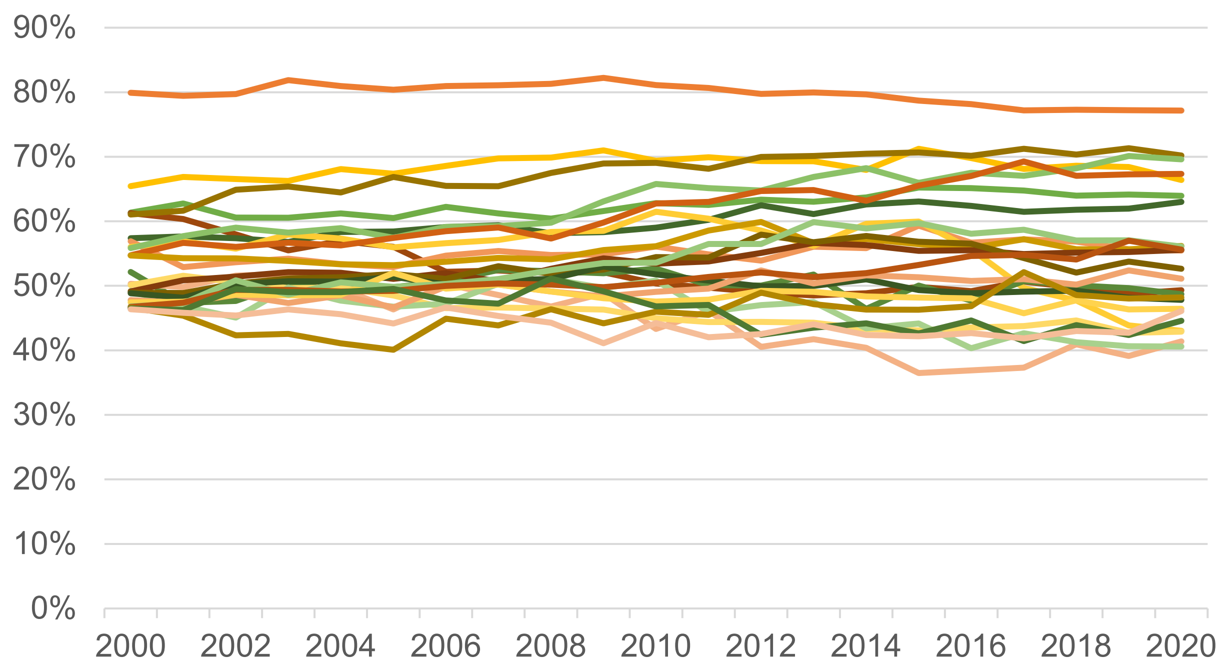

Table A1

’29 Narrative Measure

The dependent variables are *Crash Attention*, or the Google search volume index for “stock market crash,” and *Crash Attention (Dummy)*, or the dummy coded as one on days where the search volume index for “stock market crash” is non-zero, and zero otherwise. The main explanatory variables include: *’29 Narrative* is the cosine similarity of articles published on day *t* with those published between October 24 – November 8, 1929 minus the cosine similarity with those published between October 7 – 23, 1929; *StockMarketAttention* is the Google search volume index for stock market-related terms on day *t*; and the interaction between *’29 Narrative* and *StockMarketAttention*. The control variables are included where indicated, but not reported. They include values for days *t-4* through *t* for the *Crash Attention,* VIX, option-implied crash probabilities, turnover, market returns, and market returns squared. Two-way fixed effects for month and day-of-week are included where indicated, but not reported. Newey-West standard errors are reported in parentheses. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Dependent Variable:	Crash Attention_{t+1}	Crash Attention_{t+1}	Crash Attention_{t+1}	Crash Attention $(\text{Dummy})_{t+1}$
'29 Narrative_t	3.610**	1.720	1.110	1.169
	(1.793)	(2.044)	(2.069)	(0.742)
$\text{StockMarketAttention}_t$		0.137***	0.120***	0.151***
		(0.014)	(0.015)	(0.006)
'29 Narrative_t × $\text{StockMarketAttention}_t$		0.296	0.086	-0.180
		(0.923)	(0.911)	(0.327)
Lagged *Crash Attention* Terms	YES	YES	YES	YES
Control Variables	NO	NO	YES	YES
Month and Day-of-week FEs	NO	NO	YES	YES
N	6,268	6,268	6,264	6,265
R^2	19.8%	20.9%	22.1%	27.7%

Table A2

"Crash" Embeddings

The table displays top 25 similarity scores for terms associated with the term "Crash" as of December 2020. Only the terms that do not explicitly include variation of the term "Crash" are included in this table.

Term	Similarity Score
1987	77.2%
1929	70.2%
meltdown	67.4%
collapse	63.9%
plunge	63.0%
swoon	59.8%
bubble	59.1%
bursting	58.7%
debacle	56.2%
aftermath	56.1%
downturn	55.6%
depression	55.6%
plunges	55.4%
crisis	54.4%
gyrations	54.3%
panic	53.3%
bust	52.8%
rout	52.6%
tumble	52.6%
downdraft	52.1%
flash	51.8%
burst	51.7%
slump	51.1%
implosion	50.6%
turmoil	50.2%

Table A3

Bias-free ’87 Narrative Measure

The dependent variables are the survey-based crash probabilities, both adjusted and unadjusted using option-implied crash probabilities. The analysis is performed using only the subsample of individual investors and institutional investors, as indicated in the second row. The main explanatory variables include: *’87 Narrative*^{BiasFree} is the cosine similarity using the rolling estimation procedure of articles published on day *t* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987; *StockMarketAttention* is the Google search volume index for stock market-related terms on day *t*; and the interaction between *’87 Narrative*^{BiasFree} and *StockMarketAttention*. The rolling estimation procedure involves estimating a new model only using articles available up through the previous month. The control variables are identical to those used in Table 7. Robust standard errors clustered at the date level are reported in parentheses. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Investor Subset:	Indiv.	Indiv.	Inst.	Inst.
Dependent Variable:	$\pi_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi_{i,t+1}$	$\pi^{Adj}_{i,t+1}$
'87 Narrative$^{BiasFree}_t$	0.385***	0.134*	0.285***	-0.057
	(0.079)	(0.078)	(0.075)	(0.067)
StockMarketAttention$_t$	0.046***	0.009	0.056***	0.029
	(0.017)	(0.033)	(0.014)	(0.032)
'87 Narrative$^{BiasFree}_t$ x StockMarketAttention$_t$	1.204***	0.673**	-0.021	-0.192
	(0.328)	(0.287)	(0.274)	(0.213)
Control Variables	NO	YES	NO	YES
Year-Month and Day-of-week FEs	NO	YES	NO	YES
N	3,719	3,713	4,101	4,088
R^2	0.84%	5.61%	0.68%	6.27%

Table A4

Bias Residual Measure

The dependent variables are the survey-based crash probabilities, both adjusted and unadjusted using option-implied crash probabilities. The analysis is performed using only the subsample of individual investors and institutional investors, as indicated in the second row. The main explanatory variables include: *Bias Residual* is the residual from the regression model projecting *'87 Narrative* on the cubic polynomial transformation of *'87 Narrative*^{BiasFree}; *StockMarketAttention* is the Google search volume index for stock market-related terms on day *t*; and the interaction between *Bias Residual* and *StockMarketAttention*. The rolling estimation procedure involves estimating a new model only using articles available up through the previous month. The control variables are identical to those used in Table 7. Robust standard errors clustered at the date level are reported in parentheses. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Investor Subset:	Indiv.	Indiv.	Inst.	Inst.
Dependent Variable:	$\pi_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi_{i,t+1}$	$\pi^{Adj}_{i,t+1}$
Bias Residual$_t$	-0.626	-1.022	0.401	0.368
	(1.093)	(0.893)	(0.995)	(0.736)
StockMarketAttention$_t$	0.021	0.004	0.038***	0.029
	(0.017)	(0.034)	(0.015)	(0.032)
Bias Residual$_t$ x StockMarketAttention$_t$	3.817	1.353	2.819	-0.227
	(3.619)	(2.975)	(3.353)	(2.466)
Control Variables	NO	YES	NO	YES
Year-Month and Day-of-week FEs	NO	YES	NO	YES
N	3,719	3,713	4,101	4,088
R^2	0.19%	5.54%	0.24%	6.26%

Table A5

Investor Crash Probabilities and '29 Narratives Measure

The dependent variables are the survey-based crash probabilities, both adjusted and unadjusted using option-implied crash probabilities. The analysis is performed using only the subsample of individual investors, institutional investors, or both, as indicated in the second row. The main explanatory variables include: *'29 Narrative* is the cosine similarity of articles published on day *t* with those published between October 24 – November 8, 1929 minus the cosine similarity with those published between October 7 – 23, 1929; *StockMarketAttention* is the Google search volume index for stock market-related terms on day *t*; and the interaction between *'29 Narrative* and *StockMarketAttention*. The control variables are identical to those used in Table 7. Robust standard errors clustered at the date level are reported in parentheses. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Investor Subsample:	Indiv.	Inst.	Indiv.	Inst.	All
Dependent Variable:	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$
'29 Narrative$_t$	0.498**	0.765***	-0.083	0.204	-0.061
	(0.226)	(0.223)	(0.253)	(0.237)	(0.236)
StockMarketAttention$_t$	0.006***	0.002**	0.001	0.006**	0.004**
	(0.001)	(0.001)	(0.003)	(0.003)	(0.002)
'29 Narrative$_t$ × StockMarketAttention$_t$	0.282**	0.185*	0.123	-0.019	0.116
	(0.116)	(0.107)	(0.116)	(0.106)	(0.115)
Institutional$_i$					-0.022***
					(0.003)
Institutional$_i$ × '29 Narrative$_t$					-0.003
					(0.002)
Institutional$_i$ × StockMarketAttention$_t$					0.311
					(0.315)
Institutional$_i$ × "29 Narrative$_t$ × StockMarketAttention$_t$					-0.178
					(0.156)
Control Variables	NO	NO	YES	YES	YES
Year-Month and Day-of-week FEs	NO	NO	YES	YES	YES
N	7,271	6,284	7,270	6,282	13,555
R^2	0.6%	0.3%	6.0%	6.9%	3.8%

Table A6

’87 Term Frequency

The dependent variables are the survey-based crash probabilities, both adjusted and unadjusted using option-implied crash probabilities. The analysis is performed using only the subsample of individual investors, institutional investors, or both, as indicated in the second row. The main explanatory variables include: *’87 TermFrequency* is the cosine similarity of articles published on day *t* with those published between October 20 – 23, 1987 minus the cosine similarity with those published between October 5 – 9, 1987 based on the term frequency-inverse document frequency; *StockMarketAttention* is the Google search volume index for stock market-related terms on day *t*; and the interaction between *’87 TermFrequency* and *StockMarketAttention*. The control variables are identical to those used in Table 7. Robust standard errors clustered at the date level are reported in parentheses. Statistical significance for the 10%, 5% and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Investor Subsample:	Indiv.	Inst.	Indiv.	Inst.	All
Dependent Variable:	$\pi_{i,t+1}$	$\pi_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$	$\pi^{Adj}_{i,t+1}$
’87 TermFrequency$_t$	0.616***	0.195	0.363*	-0.096	0.395**
	(0.160)	(0.156)	(0.188)	(0.184)	(0.172)
StockMarketAttention$_t$	0.004***	0.002**	0.001	0.006**	0.004
	(0.001)	(0.001)	(0.003)	(0.003)	(0.003)
’87 TermFrequency$_t$ × StockMarketAttention$_t$	0.110	0.105	-0.037	0.067	0.056
	(0.076)	(0.070)	(0.087)	(0.080)	(0.080)
Institutional$_i$					-0.022***
					(0.003)
Institutional$_i$ × ’87 TermFrequency$_t$					0.000
					(0.002)
Institutional$_i$ × StockMarketAttention$_t$					-0.515**
					(0.216)
Institutional$_i$ × ’87 TermFrequency$_t$ × StockMarketAttention$_t$					-0.053
					(0.100)
Control Variables	NO	NO	YES	YES	YES
Year-Month and Day-of-week FEs	NO	NO	YES	YES	YES
N	7,271	6,284	7,270	6,282	13,555
R^2	0.7%	0.1%	7.6%	8.4%	4.5%