Emotions and Subjective Crash Beliefs*

William N. Goetzmann[1]
Dasol Kim[2]
Robert J. Shiller[3]

May 2024

Abstract: Over the past two decades, respondents to the Shiller Investor Confidence Surveys assess the probability of a catastrophic stock market crash to be much higher that the historical frequency of such events. We decompose these crash probabilities into fundamental and subjective components and use a large language model to estimate the emotional content of respondent narratives. The subjective crash component is strongly associated with high negative affect. We use respondent location to test how news of unusual exogenous shocks affects crash belief formation. The results are consistent the risk-as-feelings hypothesis and suggest a path by which emotional response to news about salient events may play a role in the scale and variation in investor beliefs about rare disasters.

Keywords: Crash Beliefs, Emotion, Risk-as-Feelings, Investor Surveys

JEL codes: G00, G11, G23, E03, G02

[1] **William N. Goetzmann; Yale School of Management (corresponding author).** Email: william.goetzmann@yale.edu

[2] **Dasol Kim; Office of Financial Research, U.S. Department of Treasury.** Email: dasol.kim@ofr.treasury.gov

[3] **Robert J. Shiller Yale University.** Email: robert.shiller@yale.edu.

* This paper incorporates and updates some results from our 2017 NBER working paper *Crash beliefs from investor surveys* (No. w22143) National Bureau of Economic Research. We thank the U.S. National Science Foundation, Whitebox Advisors, and the International Center for Finance at the Yale School of Management for support with the survey data. We thank Francesco Fabozzi for research support, particularly with the LLM. We thank Nathan Novemsky and Ravi Dhar for helpful discussions. We thank Leigh Ann Clark, Sumithra Sudhir and Minhua Wan for help with the data. The authors take responsibility for all errors. The opinions expressed are those of the authors and do not necessarily reflect those of the Office of Financial Research or the U.S. Department of the Treasury. Portions of the paper were commenced prior to affiliation with the Office of Financial Research. All errors are ours alone.

1. Introduction

Emotion and mood have been shown to play a role in financial beliefs, preferences, and decisions.[1] Beliefs about extreme events are of particular interest to economic theory and empirical research.[2] Elevated disaster beliefs are a plausible explanation for the equity premium puzzle[3] and the dynamics of crash expectations have implications for time-varying risk premium models.[4] Crash risk is also associated with cross-sectional differences in asset returns.[5] Brunnermeier et. al. (2021) highlight the dynamics of investor beliefs as a promising channel for asset pricing research.

Beliefs about extreme or infrequent events are of great interest to behavioral research as well. Fear and dread are shown to influence subjective risk assessment of rare disasters with important implications for public policy.[6] Zajonc (1980) suggests that an affect channel for decision-making can operate separately and even supersede a cognitive channel. Johnson & Tversky (1983) found that emotional responses directly affect risk assessments for extreme events. Loewenstein et. al. (2001) term this phenomenon a risk-as-feelings reaction that can cause a "divergence of emotional responses from cognitive evaluation of risks." They argue that this potentially evolutionary mechanism is common in instances of decision-making under uncertainty.

In this paper we investigate the emotional component in investor beliefs about extreme crash probabilities. Using more than two decades of Robert Shiller's Stock Market Confidence Survey we first we first decompose crash beliefs into subjective and fundamental components based upon option prices.. The subjective crash probability measure has several interesting properties. It is highly volatile, typically right-skew, and nearly always positive. While the fundamental component is correlated with market-based indicators such as market returns, volatility and volatility-implied crash risk, the subjective component is not.

In order to better understand the subjective component in crash risk forecasts, we turn to the survey response narratives. We apply a large language model (LLM) to written narratives

---

[1] cf. Sanders (1993), Da et, al. (2015), Garcia (2013), Hirshleifer & Shumway (2003), Hirshleifer et al. (2020), Goetzmann et. al, (2014), Goetzmann et. al. (2017), Kamstra et. al. (2003), Yuan et. al. (2006), Gerrett et. al. (2005 & 2017), Loughran & Schultz (2004), Edmans et. al. (2007), Edmans et al. (2022), Taffler et. al. (2021), Griffith et. al. (2020) & Hasan et. al. (2023).
[2] cf. Barberis (2013) for a comprehensive review of research on the psychology of tail events, and their relationship to subjective probability.
[3] cf. Rietz (1988) and Barro (2006).
[4] cf. Gabaix (2012), Santa-Clara & Yan (2010), Berkman et. al., (2011), Bollerslev & Todorov (2011 & 2015), Bates (2000), Wachter (2013) and Tsai & Wachter (2015).
[55] cf. Kelly & Jiang (2014) and Gao & Song (2015).
[6] cf. Lichtenstein et al. (1978) and Slovic (2012).

solicited by the investor survey to construct high-frequency measures of sentiment and other dimensions of emotion that directly correspond to psychological models of affect. The model is fine-tuned using rich data on human-derived emotion annotations to narratives. This approach allows us to extract various dimensions of emotion from the narratives that would be otherwise difficult to identify. We find that investor sentiment is correlated with the subjective crash probability component and find strong evidence that it is related to non-fundamental variation. We then consider specific forms of affect to directly test the circumplex model of Russell (1980) that is widely used in the psychology literature. As predicted by this literature, the relationship is non-linear.

Finally, we test whether the seminal results of Johnson & Tversky (1983) hold in a real-life setting. To do this we use respondents' geo-location to conduct three natural experiments that conform to the laboratory conditions in Johnson & Tversky (1983) of priming conditions. We find that exposure and attention to news of rare, extreme events: i.e. nearby earthquakes, large lottery-winnings and excess COVID deaths is significantly associated with a higher or lower rare crash probability, depending on emotional valence. Taken together our results suggest that the risk-as-feelings channel has potentially relevant implications for asset pricing and investor market participation.

The balance of the paper is organized as follows. Section 2 reviews the relevant literature and highlights the contribution of the current paper. Section 3 describes the data used in the analysis. Section 4 presents evidence on subjective crash beliefs. Section 5 presents the results on investor sentiment tests. Section 6 presents the results on natural experiments. Section 7 concludes.

2. Literature review

2.1 Survey Research

Survey evidence on investor beliefs has played an increasingly important role in behavioral economics and asset pricing. Ben-David, Graham & Harvey (2007, 2013) use a survey of CFOs to estimate executive overconfidence and miscalibration of stock market distributions. Greenwood & Shleifer (2014) assemble and analyze aggregate measures of expected returns from a number of investor surveys. They find a common factor in subjective beliefs that is negatively correlated to standard valuation metrics. Nagel & Xu (2022) assemble expectations data from a variety of survey sources to explore the relationship between experience and market forecasts.

A promising line of survey research examines cross-sectional differences in survey responses. For example, Vissing-Jorgensen & Attanasio (2003) find considerable heterogeneity in implied risk aversion. Casella & Gulen (2019) document differences in beliefs between retail vs. individual investors. Biachi et. al. (2022) demonstrate that even professional macroeconomic forecasters manifest belief distortions. Cocco et. al. (2022) show that expectations among British households were influenced by specific financial circumstances.

Recent cross-sectional survey research also highlights the potential importance that rare disaster probabilities play in investor concerns. Our prior research showed that media sentiment explained cross-sectional variation in subjective crash probabilities.[7] Choi and Robertson (2020) survey financial decision-makers in US households and find strong evidence that rare disaster risk influences their decisions. Giglio et. al. (2021) survey a large panel of wealth investors and document a strong negative relationship between return expectations and subjective probability assessment of a rare disaster. The current paper contributes to this literature by exploiting the relationship between quantitative and narrative responses to explicitly test the role of emotion in belief formation.

2.2 Sentiment

There is a substantial literature in behavioral finance documenting the impact of sentiment in news and social media on investor beliefs and decisions, which in turn likely impacts asset prices.[8] In a seminal paper, Tetlock (2007) found that negative news sentiment predicted transitory price declines in the Dow. Da et. al. (2011) show that internet search for negative affective financial terms can explain mutual fund flows. As noted above, there is considerable evidence that weather and seasonally-induced mood variation may affect market prices and operation. More recently, Griffith et. al. (2020) document a relation between emotion-coded news and volatility, and Cuculiza et. al. (2021) show that analyst opinions are negatively affected by nearby terrorist attacks. Sias et. al. (2023) use the swine flu epidemic as an experimental setting for testing whether negative affect influences market participation, Kräussl & Mirgorodskaya, (2014) find a relation

[7] Goetzmann, Kim & Shiller (2016).
[8] Duxbury et. al. (2020) is a useful overview of the literature.

between media sentiment and future market volatility, and Hasan et al. (2023) show that emotion betas derived from sentiment analysis of the news explain cross-sectional variation in stocks. Yang et. al. (2024) construct a sentiment index from COVID-19 tweets and find a relation to equity index returns. Our innovation is to use respondents' own emotional state instead of relying on news or social media as an instrument for common sentiment. Our evidence suggests that the level and time-variation in investor risk beliefs have a significant individual emotional component that correlates to media sentiment but is not subsumed by it.

2.2 Risk as Feelings

Our theoretical foundation for the analysis of emotional content in crash forecasts is the risk-as-feelings framework. [9] An early paper in this literature, Johnson & Tversky (1983) [J&T] provide a particularly useful basis for identifying the influence of emotion on subjective rare disaster probabilities. Subjects primed with narratives about an unrelated rare event – for example, news of a fatal car crash – assigned a higher probability to a different category of rare event such as being struck by lightning. J&T argue that this referred subjective inference was due to an emotional response rather than a logical semantic association. One contribution to this literature is that test the seminal results of J&T in a real-life setting.

Much subsequent research in the risk-as-feelings literature has uncovered features of this phenomenon that are potentially relevant to the rare disaster literature. For example, Rottenstreich & Hsee (1999) find that the relationship between preference and affect-rich outcomes is highly non-linear. In addition, within the risk-as-feelings framework, appraisal-tendency theory makes specific predictions about the kinds of emotion that are likely to influence such things as crash probability assessment and investor choice. Lerner & Keltner (2000) for example show that anger and fear are similar in valence but can imply differing expectations about the future.[10] One implication of the appraisal-tendency hypothesis is the potentially adaptive role of emotion in directing judgement and action. Among other things, emotions may stimulate useful and protective reactions to threatening stimuli. Our use of a LLM to classify a range of emotions in subject narratives allows us to test whether emotions influence probabilistic beliefs. Our evidence

[9] cf. Johnson & Tversky (1983), Slovic & Peters (2006), Keller et. al. (2006), Slovic et al. (2007), Loewenstein et. al. (2001), Lerner & Keltner (2000).
[10] Tiedens & Linton (2001) link different emotional channels to heuristic vs, systematic processing.

lends support to recent research in the role of emotion in asset pricing tests (Hasan et al., 2023), and also suggests that emotion-driven excessive fear of a crash may be a barrier to widespread market participation.

2.3 Neuroscience

There is considerable evidence of a neurological relationship between emotion and risk-taking. Most of this literature seeks to directly link choice under uncertainty with neurological channels and the sequence of activities that lead up to it. According to Bossaerts (2021), Bechara & Damasio (2005) are the first to demonstrate a biological link between emotions and risky choice, and Bechara et. al. (1997) find that emotional engagement plays a crucial role in decision-making under uncertainty.

Experience likely creates the emotional context for risk-taking: Kuhnen & Knudsen (2005 & 2011) use brain imaging to show that anticipatory neurological activation -- conditioned on prior experience -- influences investor choice, and that affect mediates neurological channels relating to risk aversion. Bossaerts et al. (2023) use heart-rate and skin conductance to delve into the role of anticipatory vs. reactive affect in an experimental market setting and find that affect measures Granger-cause profitable trading. This is consistent with field experiments in Lo & Repin (2002) indicating that experienced traders exhibited greater emotional control and lower physical indications of anxiety and stress. Our potential contribution to this literature is to link emotion to the cognitive formation of explicit probabilities, as opposed to subject decisions.

2.4 Rare Disasters

Rietz's (1987) proposed solution to Mehra & Prescott (1985) is a seminal article in the rare disaster literature – ironically drafted just months before the biggest crash in US stock market history. The literature since has focused on empirically estimating probabilities of rare disasters, and on theoretical frameworks to study rare disaster expectations and their effect on asset prices and investor choice. In the empirical literature, Barro (2006) constructs a comprehensive global database of crashes and finds support for the Rietz hypothesis. Berkman et al (2011) examine geopolitical factors influencing market crashes. Bialkowski and Ronn (2016) highlight the risk to property rights during crashes. Goetzmann & Kim (2018) study global equity market crashes over the long durée.

A barrier to estimating extreme negative event probabilities is that the data are not unconditionally observed.[11] Recently Orlik & Veldkamp (2023) argue that this rarity can cause large fluctuations in beliefs about disaster risk. Jiang and Kelly (2015) address this limitation by estimating market crash risk from individual stock measures. An increasingly rich option-market has made it possible to estimate crash risk – Marin (2017) uses index options to measure crash probability variation, and in a series of papers cited above Bollerslev & Todorov and co-authors have been able to estimate refined, high-frequency crash risk indices. Our contribution to this empirical research is to demonstrate the role of emotion in the formation of probabilistic forecasts of rare disasters.

On the theory side, Tsai and Wachter's (2015) overview of the rare disaster literature is a comprehensive analysis of the theoretical literature on disaster risk and asset pricing – breaking out its static and time-varying effects. The literature on time-varying crash risk premia is particularly rich. These models are able to match several empirical regularities – not only the equity risk premium and the riskless rate puzzle, but also predictability in asset returns (Wachter, 2013). Dynamic models permit consideration of non-instantaneous crashes, clustered crashes, uncertainty and learning about crashes and more. The time-series' we develop in this paper is potentially useful in testing these model predictions. Indeed, as they are not derived from option market they add another interesting dimension to consider.

A natural question is how the subjective crash beliefs relate to representative agent models. Respondents to the Investor Confidence Survey are not necessarily indicative of the marginal investor in the option market or the equity market. Indeed, the spread between their beliefs and those implied by option prices is the topic of this paper. Favilukis' (2013) for example, shows that market participation can be procyclical. In a risk-as-feelings framework variation in beliefs are correlated to emotion and thus subjective crash probabilities may empirically proxy for risk aversion. When risk aversion truncates participation, the average risk aversion of a representative investor varies with participation, implying a time-varying equity premium.[12]

[11] cf. Brown et. al. (1995)

[12] See Merton (1980)

3. Data[13]

3.1 Investor Survey

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. Information about the ZIP codes of the respondents is readily available from 2007. The combined sample used in this paper contains 16,214 responses. The survey participants provide the date on which the questionnaire was completed.

There is existing research that uses data from the Shiller surveys. Greenwood and Shleifer (2015) find that the monthly investor confidence index constructed by aggregating information from the respondent-level Shiller survey is well-correlated to several other investor surveys and to mutual fund flows. Egan et. al. (2022) extract expectations from funds that track the S&P index and find that these corelate to investor flows. Barone-Adesi et. al. (2015) estimate behavioral pricing kernels from market data and find them to correlate well to investor sentiment surveys, including indices constructed from the Shiller survey data used in this paper. Goetzmann et. al. (2016) use the institutional investor responses from a telephone version of the survey about beliefs in market mispricing in order to study variation in investor mood. Their results are consistent with evidence derived from a different dataset of investor trading behavior.

3.2. Crash Probability Survey Question

In the current study, we use responses to the survey question:

[13] Parts of the data description in this section is taken or adapted with changes from the 2017 version of our NBER working paper Goetzmann et. al. (2016).

"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.[14] Thus it has the advantage of consistency throughout a period of 26 years, during which time the stock market, the macro-economy and the financial system has experienced considerable variation.

3.3. Options Data

For options market data, we use daily data on SPY options from OptionMetrics from 1996 to 2020. The S&P 500 option-implied expected returns and crash probabilities are calculated based on Martin (2017) using options data. For the expected returns, the variables in the analysis assumes a six-month tenor, consistent with the survey data. For the crash probabilities, we assume a six-month tenor and define a crash as of 20%. We also use a high-frequency crash metric derived from option from option surfaces: the daily estimate of the (risk-neutral) expected volatility of future returns over a weekly horizon that materialize below a lower threshold developed in Bollerslev et. al. (2015).[15]

3.4. Earthquake Data

Earthquake data from 1900 to 2020 is collected from the United States Geological Services (USGS). The data includes dates, magnitudes, and coordinate locations of each event. We match the earthquake data to the investor survey data using the centroid of the ZIP code location available for some of the survey respondents. Approximately 6.36% of the survey respondents experienced an earthquake with seismic magnitudes between 2.5 and 5.5 whose epicenter is within a 30-mile

[14] 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.

[15] We thank Viktor Todorov for providing this series.

radius within 30 days prior to the survey response, while 0.05% experienced those with magnitudes above 5.5.

3.5. Lottery Data

Data on winning stores (i.e., the store where the winning ticket is purchased) for Powerball and MegaMillions contests are hand collected from press releases from 2007 to 2020 that include information on the contest winners, jackpot size, and location. We treat multiple winners for a single lottery contest as separate events in the tests.

There are several challenges with using the location of the winning contestant. In some instances, the winner is allowed to remain anonymous in some states, so that her location is unavailable. Of those where state laws require the winner's identity to be disclosed, some winners choose to have the jackpot collected by an anonymous trust, which may be associated with an address that does not necessarily correspond with the residence of the winner. Winning tickets that are split across multiple individuals may have different locations, though generally cluster geographically.

To address issues related to ambiguity in winner locations, we rely on the location of the store selling the winning ticket, which is always reported. We match the lottery winner data to the investor survey data based upon whether the winning store is located within a 30-mile radius of the 5-digit ZIP code centroid of the investor. Only investors in states participating in either one of the contests at each point in time are included in the sample. Approximately 2.48% of the survey respondents are located within a 30-mile radius of the winning store within 30 days of the survey response.

3.6. COVID-19 Data

For the COVID-19 tests, we obtain county-level data from two primary sources. We obtain data on infections and deaths from the Johns Hopkins CSSE COVID-19 Tracking Project. To obtain data on county-level policy interventions, we use the updated version of the data collected by Killeen et al. (2020).[16]

[16] Data is available for download from: https://github.com/JieYingWu/COVID-19_US_County-level_Summaries/tree/master/data.

3.7. Market Data

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). The daily returns of each index are used to empirically measure market volatility and the occurrence of extreme events. We also use the returns to the indices on and before the day that the questionnaires are completed as a control for market trends that jointly influence media articles and investor heuristics. Market volatility implied by the VIX is obtained from Federal Reserve Economic Data (FRED). All data is collected daily from 1996-2020.

3.8. Media Data

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 presume that it corresponds reasonably well to the national edition. 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.[17] We use a sub-set of these corresponding to the time-interval of the current study.

Garcia (2014) documents a significant asymmetry in media reportage of past market returns – negative outcomes are reported more frequently in periodic columns of the Wall Street Journal.[18] This is consistent with evidence that both animals and humans are conditioned to give stronger weight to negative things, experiences and events (cf. Baumeister et. al., 2001 and Rozin and Royzman, 2001). Negative experiences engage greater cognitive effort (Ito et. al., 1998), have greater influence in evaluations (Ito et. al., 1998), are more likely to be taken as valid (Hilbig, 2009) increase arousal, and enhance the memory and comprehension of the event (Grabe and Kamhawi, 2006). These prior results lead us to expect that the availability bias – if it exists – should be asymmetric. Negative events should have a greater effect on probability assessments than positive events.

---

[17] 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.

[18] Garcia (2014) focuses on financial market columns from the New York Times and the Wall Street Journal from 1905 to 2007. These columns may not necessarily appear on the front page, where articles are more likely to be viewed by readers. For example, the "Abreast of the Market" column in the Wall Street Journal, which is used in Garcia (2014), appears on the front page of a section 40.8% of the time over our sample period.

For each article, sentiment is defined as a weighted function of the number of positive and negative words.[19,20] Media sentiment is measured by aggregating the sentiment of articles for each date. We classify words using the General Inquirer (GI) – a lexicon of positive and negative words that is widely used in text sentiment analysis. A number of prior studies use the GI to construct media sentiment measures for newspaper articles.[21] Other approaches include constructing lexica tailored to specific types of source documents. For example, Loughran and McDonald (2011) manually code common terms found in SEC regulatory filings into valence classes, as some terms may have different connotations in financial documents compared to other documents.

We construct daily scores that aggregates information across articles for each date. We define *News Sentiment* as the difference between the weighted positive and negative word counts across all articles for each date, scaled by the sum of the weighted positive and negative word counts across all articles for each date.[22]

4. Subjective Crash Beliefs

We start by examining the empirical properties of the survey-based crash probabilities.

Figure 1 graphs the average annual probabilities for the survey respondents and the market crash risk indicators: the option-implied probability of a six-month drop of at least 15% in the S&P

[19] We follow Loughran and McDonald (2011) in weighting word frequencies using the "tf–idf (term frequency-inverse document frequency)" method which accounts for a word's relative prevalence within and across documents. Using un-weighted word counts does not control for the fact that some words are simply more common than others (cf. Manning, Raghavan, and Schutze (2008)). Our use of the tf-idf weighting scheme is motivated by Loughran and McDonald's (2011) finding "…that this approach [tf-idf] produces regressions with better fit than the approaches using simple proportions." Specifically, the weighted word frequency for word w appearing in article a is calculated as the product of the log-scaled word frequency and the log-inverse document frequency: $\frac{1+\ln(n_{w,a})}{1+\ln(\overline{n_a})}\ln\left(\frac{A}{df_w}\right)$, where $n_{w,a}$ is the frequency for word w in article a, $\overline{n_a}$ is the average frequency for all words appearing in article a, A is the total number of articles used in the analysis, and df_w is the number of articles containing word w. Words that do not appear in the article ($n_{w,a} = 0$) are assigned a value of zero. See Loughran and McDonald (2016) for a complete discussion and survey of methods using word frequency in textual analysis to measure sentiment.

[20] The General Inquirer assigns approximately 10,000 words to 26 major and 182 minor categories, or tags. It aggregates categories from the Harvard IV-4 dictionary, Lasswell value dictionary, a social cognition dictionary from Semin and Fiedler (1988). The two largest categories are the positive and negative valence classes: the positive list includes 1,915 words, while the negative list includes 2,291 words.

[21] These studies include Tetlock (2007), Engelberg (2010), and Garcia (2013). There are a large number of studies that apply sentiment analysis to firm regulatory disclosures. Loughran and McDonald (2016) provides an overview of this literature.

[22] We also consider a number of other specifications: the weighted negative word count divided by the total weighted word count; the difference between the positive and negative word count divided by the total weighted word count; and the difference between the natural log of one plus the weight positive word count and the natural log of one plus the weighted negative word count. The results are not sensitive to these alternative specifications.

500 index, the annualized volatility of the daily DJIA, the largest negative return in each year (represented as a positive number on the figure), and the VIX implied volatility. The annual survey mean probability exhibits a high correlation with the market-based crash risk proxies. These periods also correspond to higher option-implied crash probabilities,[23] realized volatility, implied volatility and most extreme one-day DJIA percentage declines. The Pearson correlation between the average annual probabilities and the market-based proxies range between 43.3% to 57.3%. Appendix Table A1 displays correlations on the daily crash probabilities and the variables used in the analysis.

While the survey-based crash probabilities exhibit considerable time-series covariation with other proxies for market crash risk, it is unclear whether the levels at any given point in time are consistent with rational models. To assess this, we start by examining the summary statistics of the key variables used in the analysis displayed in Table 1. Over the sample period, the survey-based crash probabilities (π^{Survey}) have a sample mean of 19.3% with a standard deviation of 19.8%. There are challenges with comparing the probabilities from those inferred from historical market returns. It is well known that stock returns are fat-tailed and this log-normal model is not appropriate to estimate the probability of an extreme decline. The average daily standard deviation of the DJIA is about 1% and the two crashes of interest are 12 times and 20 times the daily standard deviation. This has motivated the use of fat-tailed distributions and mixed jump processes to describe stock market moves.[24]

A simple approach to estimating a baseline probability is to use the historical frequency of such events. Under the assumption of an i.i.d. distribution of daily returns, and using the number of trading days since October 23, 1929 through December 31, 1988 [taking the most conservative bounds] gives an average probability of an extreme crash over a six-month horizon of 1.7%. This decline to approximately 1% when the entire history of the DJIA is used. The average reported crash probability from the Shiller surveys is thus more than 10 times the conservative estimate.

A more refined approach is to directly compare the survey-based crash probabilities with the option-implied ones. Crash probabilities can be derived from index options that conform to some extent to the specifications indicated in the survey question. We use as a baseline probability

[23] The option-implied crash probability estimates are derived from the approach described in Martin (2017). Specifically, the probabilities are derived from data on S&P 500 index options expiring in six months or less.
[24] Cf. Gabaix (2012), Santa-Clara and Han (2010), Wachter (2013), Bollerslev and Todorov (2011).

estimate the option-implied probabilities from Martin (2017) based on a six-month horizon associated with an at least 15% drop in a market index. Table I reports summary statistics for the option-implied probabilities the mean value of 6.2% corresponds the historical frequency of overlapping six-month returns of large-cap stocks over the period June 1926 through December 2021.[25]

Santa-Clara & Han (2010) and Bollerslev et al. (2011) find that option-based metrics can be used to estimate high-frequency crash risk premia. The daily estimate of the (risk-neutral) expected volatility of future returns over a weekly horizon that materialize below a lower threshold from Bollerslev et al. (2015) is another fundamental metric for crash risk and is highly correlated with other option-based crash risk measures.[26] Over the period of our analysis, the correlation of this jump risk measure with the VIX is 67.8% and is 58.9% with the Martin crash probability measure – indicating they are not perfect substitutes but are all likely to be reflecting fundamental risks of large market moves.

One potential issue is that the Martin and the survey-based probability measures may differ due to wording. The wording of the survey question implies a one-day drop of 15%, rather than a 15% decline over the next six months. It would suggest that the option-implied measure may be systematically higher than the survey-based measure. This is not the case. The crash probability spread (CPS), which is the difference between the survey-based and option-implied crash probability measures, are positive throughout the sample period. In fact, we find that the survey-based crash probabilities are consistently higher than the option-implied crash probabilities across various sample splits.

Table 2 displays the results. Over the full sample period of 2000 to 2020, *CPS* is positive and statistically significant (*coefficient* = 13.2%, *t*-value = 40.08). In the subperiod splits, *CPS* remains positive and significant across the different periods. *CPS* stays positive during low volatility periods, and also remains positive when daily market returns or past month returns are high or low. Finally, the relationship in *CPS* across the various sample splits are also interesting. For example, while the option-implied crash probabilities decrease over the sample period (see Appendix Table A2), the survey-based crash probabilities increase (see Appendix Table A3). This provides an explanation for the growing spread over the sample period.

[25] We thank Ian Martin for providing code and data, as well as helpful discussions.
[26] We thank Viktor Todorov for providing this series.

Figure 2 displays the survey-based crash probability distribution against the option-implied crash probabilities over time. It indicates that the positive spread is not only due to extreme survey responses. The median crash probabilities are generally above the option-implied crash probabilities. The figure also shows that both the median and the right tail of the distribution are elevated during the second half of the sample period. Figure 3 graphs monthly average *CPS* from 1996 to 2020. The figure confirms an increasing pattern in the spread through the second half of the sample period. Because the survey-based and option-implied crash probabilities are correlated with each other, the spread does not exhibit the same correspondence with market-based crash indicators. For example, the correlations with VIX are 30.9% for survey-based crash probabilities and 87.4% for the option-implied crash probabilities (see Table A1). In contrast, the correlation between VIX and *CPS* is only -3.4%.

For the purposes of the analysis in the remainder of the paper, we focus on the spread of the survey-based over the option-implied probabilities. Our logic is that option prices are based on investment decisions in the index option markets, and thus reflect the time-varying beliefs and preferences of the marginal investor in these markets. Option investors are likely to take into account current fundamental factors affecting time-variation in the probability of extreme events. Thus, the option-implied probabilities provide a useful "market-based" benchmark for analysis of the survey responses. The fact that the mean implied crash probability of a negative 15% six-month return lies close to the historical frequency of such events suggests that we cannot reject the hypothesis that, on average, option market participants rationally assess crash probabilities.

Table 3 displays the results of regression models of the daily survey-based crash probabilities and *CPS*. We start by examining the relationship with option-based crash risk measures. Column (1) displays the results for the Martin measure. Column (2) presents the results for the LJV measure described in Bollerslev et al. (2015). Each of the specifications also include the first five daily lags of the explanatory variables, including for the dependent variable. The results confirm that the cumulative effects of each of both variables on the survey-based crash probabilities are positive and statistically significant. We next consider other fundamental measures. Columns (3) and (4) show the results for daily market returns and the VIX, respectively. These specifications also indicate a significant relationship. Columns (5) through (7) repeat the tests when using CPS as the dependent variable, which adjusts the crash probabilities for fundamentals. The results in these specifications are statistically insignificant. These findings suggest that focusing on the spread controls for variation in fundamental factors that rationally influence crash expectations.

One issue to consider is that the phrasing of the survey question may make a crash salient and lead to a heightened probability assessment. The term "catastrophic" and the highlighting of the two crash dates may themselves prime a response biased towards higher probability. The high-sentiment term "catastrophic" could itself elicit an emotional response. By the same token, highlighting two crashes out of a century or more of data could trigger an availability heuristic. There are several other questions in the surveys – some with positive and some with negative sentiment; all about the stock market. These may also prime an investor response. These stimuli make it potentially difficult to identify the marginal influence of sentiment and other non-fundamental shocks on probability assessments.

Another feature of the question is that it relies partially on a narrative about an event occurring in other countries and spreading to the U.S. This may also have confounding effects. Construal theory (cf. Trope & Wakslak, 2007) suggests that psychological distance in time and space can prime higher-level, abstract reasoning which can "guide prediction." Finally, Experimental evidence suggests that people rely on numerical and narrative evidence in assessing probabilities, and the relative degree of reliance may depend on numeracy (Dieckmann et. al. 2009). All of these suggest that the key question of interest may be biased, or inconsistent with the respondent's beliefs, and that there could be heterogeneity in the responses.

One way to assess this potential issue is to examine the internal validity of the crash probabilities with the subject's other survey responses. Given the affective and narrative features of the question, prior research suggests that we should find cross-sectional differences among respondents based on their numerical sophistication and perhaps other factors. If the high base-line probability assessments are due solely to framing factors within the questionnaire, this would suggest that direct priming may be a source of extreme bias about the probabilities – an interesting fact in itself. That is, the crash probabilities may reflect severity due to the wording of the question rather than their actual beliefs. This would imply inconsistency between the response to the crash probability question and the responses to other questions where the wording is unlikely to be confused with severity. Moreover, confusion due to question wording as well as noise, possibly related to respondent error, may limit the explanatory power of the statistical models used in subsequent sections. We assess the former in two ways.

First, we estimate a model that examines whether the crash probabilities can be explained by responses to the questions related to stock market valuation. The wording of these questions may be less likely to elicit extreme responses.[27] The key explanatory variables include the expected percentage change in the DJIA over the next six months (*Expected Returns*); a dummy based upon whether stock prices when compared to true fundamental value is too high (*Too High Valuation*); a dummy based upon whether stock prices when compared to true fundamental value is too low (*Too Low Valuation*); a dummy based upon whether the investor is inclined to sell stocks overall (*Sell*); and a dummy based upon whether the investor is inclined to buy stocks overall (*Buy*). Appendix Table A4 displays the results. The explanatory variables are mostly significant, and their coefficients are signed in a manner consistent with the crash probability responses.[28]

Of particular interest in Table A4 is the relatively high explanatory power and coefficient estimate on the variable "Sell Recommendation." It indicates that the high estimate of a crash probability is associated with the a recommendation to exit the market. Overall, these results suggest that the crash probability responses overall are internally consistent with responses to the other survey questions.

Second, we directly test whether specific words in the survey question influences the crash probability responses. Different versions of the survey question that alter keywords that may potentially prime severity in the question phrasing are randomly presented to 500 subjects using Amazon's Mechanical Turk.[29] We estimate a regression model where the dependent variable is the crash probability response and the explanatory variables are dummies associated with keywords, along with date fixed effects.[30] We do not find any statistically significant differences in the crash probabilities across the different versions. The crash probabilities for the versions that exclude the words "catastrophic" and replaces the word "crash" with "decline" (*estimate* = 0.007,

[27] We use three questions from the survey for these tests: (i) "How much of a change in percentage terms do you expect in 6 months [for the Dow Jones Industrial Average]?", (ii) "Stock prices in the United States, when compared with measures of true fundamental value or sensible investment value, are [too high, too low, or about right].", and (iii) "Are you inclined to buy stocks overall, or sell stocks overall, or hold steady?"
[28] In untabulated results, we also find similar effects when examining the spread between the survey responses over option-implied crash probability estimates based upon S&P 500 options data. This suggests that idiosyncratic variation in the survey responses relative to fundamental market conditions are also internally consistent.
[29] Specifically, we consider four versions of the survey question: the original question; one that excludes the words "catastrophic" and replaces the word "crash" with "decline"; one that replaces "like that of October 28, 1929 or October 19, 1987" with "of 20%"; and one that excludes the words "including the case that a crash occurred in the other countries and spreads to the U.S.".
[30] The survey was conducted over three dates: July 14, July 15, and July 16, 2016.

t-value = 1.22); replace "like that of October 28, 1929 or October 19, 1987" with "of 20%" (*estimate* = 0.013, *t-value* = 0.57); and exclude the words "including the case that a crash occurred in the other countries and spreads to the U.S." (*estimate* = 0.010, *t-value* = 0.63) are not significantly different from those of the original question, and the estimated magnitudes are negligible. Together, these results provide evidence that the crash probability responses are not significantly driven by confusion over the wording of the survey question.

5. Investor Sentiment Tests

In this section, we examine to what extent is time variation in the survey-based crash probabilities driven by non-fundamental factors. Towards this end, we utilize other data from the survey as well as generative modelling techniques in order to identify various dimensions of affect and to construct a high-frequency measure of investor sentiment.

5.1. Measurement

Our identification strategy relies on access to high-frequency measures of non-fundamental factors. This is challenging. For example, most proxies of investor sentiment are only available at the monthly-level or lower frequencies. Coarser frequencies increase the difficulty for teasing out the influence of fundamental factors. We attempt to overcome these challenges in part by developing high-frequency measures of investor sentiment using narrative data from the survey. Specifically, the survey asks respondents to provide an open-ended commentary to describe the stock market conditions or provide an explanation for their survey predictions.

We feed the responses into a large language model (LLM) in order to quantify various dimensions of emotion, including sentiment. Our approach employs a model that is tailored to capture emotional qualities in textual data. Specifically, we utilize an open-source Mistral model (Mistral-7B-Instruct-v0.1),[31] which is a generative language model trained for conversational output. This model contains 7.3 billion parameters, which is relatively large compared to BERT (340 million parameters) but smaller than GPT-4 (speculated to be over 1 trillion parameters). Despite its compactness relative to some LLMs, there are studies that document at least comparable performance of the Mistral model compared to larger models (Jiang et al., 2023).

[31] For a complete description of the model, refer to Jiang et al. (2023).

Moreover, the advantage of using the Mistral model is that it is designed for the purpose of being fine-tuned on other data. GPT-4, in contrast, is not open-source and so is less adaptable to our purpose. We apply fine-tuning using Low-Rank Adaptation (LoRA) on the "Affective Text" dataset from Strapparava & Mihalcea (2007).[32] The advantage of using this data for the fine-tuning is that it features much more fine-grained emotion annotator scores compared to others that only feature coarse binary scores.

To demonstrate the effectiveness of the LLM adopted for the analysis, we consider several examples. Each example is assigned a specific emotion score on a [0, +100] scale. The sentiment is also evaluated on a [-100, +100] scale. None of the examples were used to fit the model.

Example 1: "Who are you calling fat?"

This example is associated with very negative sentiment (Sentiment Score = -54). It also scored highly on the anger dimension (Anger Score = 47). Outside of anger, the text does not seem to reflect other negative emotions, such as fear or sadness. Consistently, the model assigns the lowest scores possible for those emotions (Fear Score and Sadness Score = 0). The model also assigns the lowest score possible for optimism (Optimism Score = 0).

Example 2: "Squirrel jumps boy in park; rabies suspected."

This example is also associated with very negative sentiment (Sentiment Score = -61). However, it received the lowest score possible for the anger dimension (Anger Score = 0). It received a moderately low sadness score (Sadness Score = 17), but a very high fear score (Fear Score = 61). As with Example 1, the optimism score is low (Optimism Score = 0).

Example 3: "Global Sludge Ends in Tragedy for Ivory Coast."

[32] Specifically, the data provides manually coded annotations of news headlines to scores for six emotions: anger, disgust, fear, joy, sadness, surprise. Human subjects are asked to annotate headlines extracted from news web sites with each emotion label from a [0,+100] scale. Each headline is independently labelled by six subjects.

Similar to the Examples 1 and 2, this example also received a very negative sentiment score (Sentiment Score = -74). It received the lowest possible score for anger (Anger Score = 0), and a moderately low score for fear (Fear Score = 12). However, it scored highly on sadness (Sadness Score = 74).

Example 4: “Scientists discover miracle in the depths.”

For completeness, we consider an example with a very positive sentiment score (Sentiment Score = 67). The example received the lowest possible score for anger, fear, and sadness (Anger Score, Fear Score, and Sadness Score = 0). In contrast, the optimism score is high (Optimism Score = 50).

At the least, the examples demonstrate the emotional dimensionality beyond sentiment captured by the LLM approach. To evaluate the external validity of the model, we compare the emotion scores to those produced by GPT-4. The average correlation between the two sets of scores is 63.3%. Moreover, the main results from the analysis are not sensitive to using the GPT-4 scores. Using GPT-4 to score the examples above yields qualitatively similar results.[33]

To start, we construct two different investor sentiment measures based upon the model responses. First, we construct a daily index. It is based upon a seven-day, backward-looking moving average, which is used in order to account for the sparsity in survey responses during some periods with sufficient text to generate scores. For periods when no survey data is available within a seven-day period, we use the most recent daily value available. Second, we use the individual scores for the respondent-level tests. We evaluate the comparability of the model results to those based upon traditional dictionary-based approaches, and find qualitatively similar results.

The primary advantage of the LLM approach is that it allows us to measure other dimensions of affect that have not be as readily studied. A large literature studies sentiment analysis with many different approaches that offer at least comparable levels of performance. There is evidence even within this domain on the relatively strong performance of large language models over dictionary-based approaches (Kant et al., 2018; Chang et al., 2023). However, much

[33] Using GPT-4, for examples 1 / 2 / 3 / 4, the Sentiment Scores are -60 / -70 / -90 / 80, the Anger Scores are 80 / 20 / 50 / 10, the Fear Scores are 20 / 80 / 70 / 20, the Sadness Scores are 30 / 10 / 90 / 10, and the Optimism Scores are 30 / 10 / 5 / 90.

less work has been conducted in quantifying other dimensions of affect. This is not the first study to use LLM to quantify affect. For example, in early work Lindquist (2015) considers the relationship between language and affect, and the extent to which natural language processing can be effective in emotion detection. More recently, various studies have tested a broad range of large language models and document their effectiveness in detecting emotional stimuli (Amin et al., 2023; Li et al., 2023).

5.2. Empirical Specification

Our initial specification focuses on time-series variation in subjective crash probabilities and its relationship with the investor sentiment measures. A key challenge is that the two are likely to be jointly related to fundamental factors. For example, there may be similarity in language used by investors and media outlets to describe stock market conditions.

We address this issue in three ways. First, rather than focusing on survey-based crash probabilities directly, we instead use the crash probability spread measure examined in the previous section. While the survey-based crash probability levels are correlated with fundamental factors, we have already shown that the spread is not. Second, we directly account for language used to describe fundamental factors by controlling for an analogous measure of sentiment but based upon news articles about the stock market. Similar to the option-implied crash probabilities used in the spread measure, the news sentiment measure may also be capturing non-fundamental factors in addition to fundamental factors. In this regard, our approach is conservative as it will restrict the tests to variation orthogonal to what is captured by the news sentiment measure. Finally, we include a host of other control variables that correspond with observable fundamental factors.

The following is the baseline regression model:

$$CPS_t = \alpha_1 \times Average\ Investor\ Sentiment_t + \sum_{j=0}^{5} \alpha_{2,j} \times News\ Sentiment_{t-j} + \sum_{j=0}^{5} \boldsymbol{\alpha_{3,j}} \times \boldsymbol{X_{t-j}} + \sum_{j=1}^{5} \alpha_{3,j} \times CPS_{t-j} + \varepsilon_t \quad (1)$$

The dependent variable is the daily spread between the average survey crash probability across responses over the past seven days and the option-implied crash probability.[34]

[34] We also consider alternative specifications. We compute the survey crash probability only using information for date *t*, and using the most recent value available if there are no survey responses available. In another specification,

$Average\ Investor\ Sentiment_t$ is the daily investor sentiment measure described above. $News\ Sentiment$ as well as its first five lags are included in the model. The other control variables ($\boldsymbol{X}$) that correspond with other fundamental factors are the daily market returns and VIX, as well as the first five lags of each variable. The first five lags of the dependent variable are also included in order to limit the influence of any other factors correlated with its lagged values. Given the construction of the dependent variable and in order to mitigate the influence of serial correlation, we use Hansen and Hodrick (1980) standard errors to assess statistical significance.

5.3. Baseline Results

Table 4 displays the results from the baseline specification. We iteratively include the different components of the baseline model in order to better understand the contribution of each to the overall results. Column (1) displays the results for the survey-based crash probabilities with only the lagged dependent variable terms. The *Average Investor Sentiment* coefficient is negative and statistically significant (*estimate* = -0.561, *t*-value = 7.79). That is, more positive investor sentiment is associated with lower crash probabilities. Column (2) displays the results when only including *News Sentiment* in the model. As expected, the coefficient is also negative and statistically significant (*estimate* = -0.125, *t*-value = 2.50). Column (3) includes both measures simultaneously in the same model. The *Average Investor Sentiment* coefficient remains virtually unchanged while the *News Sentiment* coefficient attenuates by more one quarter. However, both terms remain significant.

We next examine the specifications where the dependent variable is the crash probability spread measure. Column (4) shows that the *Average Investor Sentiment* coefficient remains significant while the *News Sentiment* coefficient becomes statistically insignificant at the 10% level. This confirms our intuition that the spread measure is effective in accounting for fundamental factors. Column (5) includes all the other control variables. The coefficient on *Average Investor Sentiment* actually grows larger in absolute magnitude and remains significant. In the last model, we examine an even stronger specification. It regresses the crash probability spread measure on an analogous spread measure for sentiment – the difference between the investor and news sentiment. Column (6) displays the results. The sentiment spread measure is

we orthogonalize the survey crash probability to the option-implied crash probability by using the residuals from a regression. The main results are insensitive to both specifications.

negative and statistically significant. We find similar results when using similar specifications but for the respondent-level data. Appendix Table A5 displays the results. The pattern in the results is strikingly similar. The statistical significance of the *News Sentiment* term also disappears for the crash probability spread specifications. Likewise, the coefficients for the *Investor Sentiment* terms are quite stable across all the specifications.

An important observation from the previous section is that a large amount of time variation in the spread is explained by the right tail of the survey-based crash probability distribution. We next assess to what extent the effects are related to spikes in the spreads. Quantile regressions are estimated to assess the degree of non-linear dependence in the effects. We include specifications for the 10th, 25th, 50th, 75th, and 90th percentiles. For these tests, we use the disaggregated, respondent level data. Table 5 shows the results. While the effects remain negative and significant for the 50th percentile specification, the magnitude exponentially increases as the percentile increases. For the 90th percentile specification, the *Investor Sentiment* coefficient is almost four times larger compared to that of the 50th percentile specification. This suggests that the effects of sentiment are much stronger during periods when the spread spikes.

The option-implied crash probabilities also exhibit spikes throughout the sample, though to a lesser extent. We repeat the exercise to identify whether there is non-linear dependence for sentiment in our fundamental factor related to crash risks. Namely, we rerun the quantile regressions using the daily option-implied crash probability series. Appendix Table A6 displays the results. The coefficient is negative for all of the specifications, but not all are statistically significant. The coefficients are statistically significant for the 10th, 50th, and 90th percentiles, but are not for the other specifications. Moreover, the sentiment measure coefficients are not monotonically increasing or decreasing in the quantile. This suggests that the pattern found in Table 5 differs from that for fundamental factors.

5.4. Affect Models

The previous section shows a strong relationship between investor sentiment and the crash probabilities. In this section, we dig deeper to better understand the role of affect and to what extent it is driving the results.

We start by extending the baseline specification to directly evaluate predictions from the seminal study Russell (1980). The study presents a circumplex model that represents affect in terms of two dimensions: valence and arousal. The valence dimension, or sentiment, captures the polarity of the emotional affect, ranging from attractive to averse. The arousal dimension captures the intensity of the emotional affect. For example, sadness is considered to be a low-arousal emotion while disgust is considered to be a high-arousal emotion. The implication of the model is that arousal serves as a moderating factor for valence, or sentiment. The effects of positive sentiment are expected to be stronger when arousal is higher.

A key challenge is measuring dimensions of affect unrelated to sentiment. We use the LLM approach described earlier in this section to quantify arousal. A prompt that simply asks for a numerical arousal score is used to obtain the measure used in the analysis. We further investigate the sensitivity of using more detailed prompts and alternative LLMs. For example, we consider prompts that use a much more detailed definition of arousal: "Arousal: Captures the intensity or activation associated with an emotion (e.g., calm vs. excited). It interacts with valence to create different emotional states and plays a role in motivation and attention." We find similar results when using the detailed prompt. This suggests that the LLM is able to associate the term arousal with the definition commonly used in the psychology literature. We also find qualitatively similar results using arousal measures derived from GPT-4.

Table 6 presents the results. Column (1) shows the results where the arousal score is included as an explanatory variable. The coefficient is positive and statistically significant (estimate = 1.713, t-value = 8.12). That is, comments that include language that rate higher on the arousal scale are generally associated with higher crash probabilities. When adding the sentiment score in Column (2), the arousal score coefficient attenuates somewhat but remains statistically significant. Column (3) introduces the interaction term between sentiment and arousal. The coefficient is negative and statistically significant (value = -0.469, t-value = -4.69). In other words, the effect of sentiment is amplified by the arousal factor. Finally, the results are robust to using the crash probability spread as the dependent variable, as shown in Column (4).

The results are also economically significant. Consider the difference in marginal effects for a one standard deviation decline in the sentiment score for the arousal score at the mean and one standard deviation above the mean. The effects when the arousal score is at the mean is an increase in the crash probability spread of 2.74%, which represents 14.1% of the total sample variation. The effects when the arousal score is one standard deviation above the mean is 5.39%, which represents 27.7% of the total sample variation. In other words, the effects almost double when accounting for arousal.

Our results are consistent with the predictions of the circumplex model. We next examine the effects of specific emotions on the crash probabilities. As before, we use a LLM approach in quantifying each emotion. The emotions considered are: anxiety, fear, sadness, disgust, optimism, and excitement. We also consider the first principal component of these emotion scores. Appendix Table A7 displays the factor weights. The principal component weights correspond with negative valence and take on the expected signs. In particular, the first principal component tends to load more on emotions with higher arousal for the emotions that generally associated with negative sentiment. However, for the emotions generally associated with positive sentiment, the factor weights are somewhat comparable, though "optimism" has a relatively larger weight in absolute terms compared to "excitement." The eigenvalues are 3.20, 1.12, and 0.96 for the first, second, and third principal components, respectively.

Table 7 shows the results. The coefficient signs generally correspond with the expected directions based upon the sentiment tests. Some of the variation in the coefficients may be related to the difficulty in capturing certain emotions using the LLM approach. For example, "disgust" corresponds with negative valence but high arousal, yet the coefficient is relatively smaller in absolute magnitude. In contrast, "sadness" has negative valence with low arousal, yet the coefficient is relatively larger or is comparatively large in absolute magnitude to the other emotion scores. Using the principal component score rather than the specific emotions may alleviate measurement issues to some extent. Consistently, that specification has the largest factor loading as well as the highest R^2 across the specifications.

5.5. Robustness Checks

In this section, we consider robustness checks on the baseline specification. We first consider whether the moving average specification used to construct *CPS* and the *Investor Sentiment* measures may complicate interpretation of the results. To address this issue, we use the disaggregated survey data, rather than the time-series data, and examine the relationship for the investor's own sentiment versus the sentiment of other investors. Appendix Table A8 displays the results. It shows that each have a significant effect on the crash probabilities. Even when considering the relationship with the sentiment of other investors on the previous day, the effect is still negative and statistically significant.

We also examine the sensitivity of the results to our choice to measure sentiment using the LLM approach. We consider two alternative sentiment measures: (i) a measure based upon the methodology of Loughran and McDonald (2011) and (ii) a measure based upon positive and negative valence terms similar to what is used in Tetlock (2007). Appendix Table A9 displays the results. It shows that the results are similar when using these alternative measures, suggesting that they are not sensitive to the choice of sentiment measure.

Finally, we examine to what extent other qualities associated with the survey respondent or writing can be attributable to the results. Appendix Table A10 shows the results. First, we examine the effects of the extent to which the survey responses reflect investor sophistication through the language used. The survey respondents may be more articulate when expressing more strongly valanced views. We use the LLM approach to score the survey text data based upon sophistication. Column (1) displays the results. It shows that the results are statistically insignificant. Second, we examine how the effects vary based upon the coherency of the responses. Column (2) shows the results. The coefficient is negative and statistically significant (*estimate* = -0.104, *t*-value = 2.88), similar to the results on investor sentiment. It suggests that lower language coherency in the survey responses is associated with higher spread levels. In Column (3), we include all three measures in the same model. The only coefficient that remains significant is associated with the *Investor Sentiment* term. The results suggest that the results are unlikely to be driven by other qualities associated with the investor and the text.

6. Natural Experiments

In the previous section, we provided evidence that non-fundamental shocks explain at least some of the time variation in the survey-based crash probabilities. This section considers cross-sectional tests using three natural experiments to identify a specific source of non-fundamental variation: availability bias.

Experimental evidence from the social and cognitive psychology literature demonstrates the influence of availability and affect heuristics in decision-making. Lichtenstein et al. (1978) find that individuals overestimate the probability of highly publicized causes of death, while underestimating those that are less publicized, and attribute their findings to availability. Johnson and Tversky (1983) document global effects associated with availability mediated by affect. They show that individuals who read sad newspaper articles about a specific cause of death or disaster gave higher risk assessments on mortality rates on those that are unrelated compared to the control group. Similarly, they show that individuals who read happy articles recounting fortunate events unrelated to death reported lower risk assessments on mortality rates for different causes of death compared to the control group.

They interpret their findings as evidence that individuals may more rely on readily available affective impressions in decision-making, particularly when the domain of judgment is complex or when mental resources are limited. In other words, "availability may work not only through ease of recall or imaginability, but because remembered and imagined images come tagged with affect" (Slovic et al., 2004).

In this section, we examine the relationship of investor crash probability assessments to plausibly exogenous negative and positive rare events unrelated to aggregate market activity. We use specific predictions from the social and cognitive psychology literature to distinguish whether investor crash beliefs may be conditioned by availability and affect biases. Stock market crashes, though rare, have been highly publicized events carrying unambiguously negative connotations.

Likewise, earthquake and large lottery jackpot events are also rare and highly publicized, though the former (latter) is generally associated with negative (positive) affect. These events are also unlikely to have economic relevance for a future stock market crash. We assess whether investor proximity to these events influences crash probabilities in a manner that is consistent with the experiments of Johnson and Tversky (1983). We also use cross-sectional variation across regions and investor types to detect whether there were differential effects on crash probability estimates of information associated with COVID-19 during the early days of the pandemic. The advantage of these tests is that they focus on variation across counties related to public health as well as related policy interventions while properly controlling for national trends, which could be related to aggregate market risks.

There are other judgement heuristics that may be applicable. Simulation biases may arise when individuals are asked to make predictions or generate counterfactuals on uncommon events. Kahneman and Tversky (1982) show that individuals may overweight scenarios associated with adverse affect.[35] The salience of these scenarios corresponds with the low redundancy and high

[35] Cf. Aktar et al. (2012).

causal significance. Bordalo et al. (2022) show that older individuals underestimated risks related to the COVID-19 pandemic while younger individuals overestimated risks, positing the use of simulation heuristics as an explanation. Namely, they argue that individuals selectively recall and use past experiences in evaluating novel risks. In our context, investors may either have personal experiences with major stock market crashes, or narrative accounts of them from other sources, that inform subjective probability estimates of a stock market crash. Simulation describes a mechanism that governs availability. However, our tests cannot be used directly to detect simulation biases in crash probability estimates given that the set of instruments we consider are completely exogenous to the occurrence of stock market crashes.

We exploit the ZIP code location of a subset of the survey respondents to identify regional events that plausibly make rare events more cognitively available. We use the occurrence of the event in the past month for investors located within 30 miles of the epicenter for earthquakes and the location of the winning store for lottery jackpots. While the timing of the earthquakes and lottery winners are exogenous to current market conditions, the events should be salient to individuals located nearby. Moderate earthquakes can sometimes be felt over long distances depending on a number of factors, though they are generally more readily detectable closer to the epicenter. Figure 4 displays the geographic distribution of the earthquake events considered in our analysis that occur during the sample period. It is comparable to the geographic distribution of earthquakes historically.

Recent studies suggest that proximity to lottery winners can affect both beliefs and actions, as neighbors of lottery winners tend to exhibit higher levels of consumption and bankruptcy risk (Kuhn et al., 2011; Agarwal, Mikhed, & Scholnick, 2018). Figure 5 displays the geographic distribution of lottery winners. Interestingly, the geographic distributions of earthquakes and lotter winners are quite distinct from each other, providing greater credibility to the analysis. Finally, we consider information associated with the investor's county of residence using county-level data on COVID-19. All of these events are also likely to be reported in the local news.

We expect investor sophistication to be higher on average in the institutional investor subsample, or that assessing crash probabilities are more challenging for individual investors. Experimental evidence from the social psychology literature would suggest that the influence of availability and affect biases should be pronounced in individuals lacking expertise in the domain of judgment, or individual investors. As such, we perform the tests separately for the individual and institutional investor subsamples. Figure 6 displays the annual average crash probabilities across institutional and individual investors. For a majority of the sample period, the crash probabilities of individual investors are higher than those of institutional investors. Our tests assess whether susceptibility to biases may be an explanation for the differences.

5.1. Experimental Validation

Before turning to the tests, we start by assessing whether the events we are focusing on indeed prime the attention of local investors. Karlsson et. al. (2008) show that attention heightens the effect of information in decision-making and Barber and Odean (2008) highlight the salient role of attention in investor behavior and Sicherman et. al. (2015) connect this to portfolio allocation decisions. We measure attention using weekly, disaggregated internet search volume data from Google Trends. The data provide high-frequency data on search volume indices (SVI) associated with various search terms. The disaggregated form that we use breaks out *SVI* geographically. Specifically, the data is obtained at the designated market area (DMA)-level. We measure the distance of each survey respondent from each event based upon its distance from the DMA centroid. For the tests, we use *SVI*s based upon the search terms “earthquake” and “lottery”.

In our first set of tests, we assess the relationship between the proximity to the event and *SVI*. Table 8 presents the results. Panel A and B display the findings associated with earthquakes and lottery winners, respectively. Column (1) displays the specification where the key explanatory variable is the natural log of one plus the number of miles between the nearest earthquake or lottery winner that occurred within the past month. Columns (2) through (4) display the specifications where indicator variables associated with distance are used instead. In all the specifications, date and DMA fixed effects are included where indicated. The control variables include the one-week lagged dependent variable as well as the natural log of one plus the historical frequency of earthquakes (Panel A) or lottery winners (Panel B). The standard errors are double clustered on the date and DMA levels.

We start with the tests using the earthquake events. The coefficient on the distance measure is negative and statistically significant (*estimate* = -0.047, *t*-value = 3.92). That is, greater proximity is associated with higher attention, or *SVI*, levels. The indicator variable specifications yield consistent results. The coefficients for the indicator variable associated with a within-30 mile range is positive and statistically significant across all the specifications. The results remain significant when adding the control variables. Finally, the results grow slightly stronger when including the DMA and date fixed effects in the final specification.

The test results for the lottery events are quite similar. Both the continuous and indicator variable distance measures are significant and are signed consistently with the earthquake results. One difference is that the results for the lottery events appear to be stronger based upon the indicator variable specifications. The results are also more sensitive to the inclusion of the control variables and fixed effect terms. Regardless, the results indicate that both earthquake and lottery events generate a strong, localized effect on internet search volumes.

The second set of tests examine time-variation in attention around the timing of the events. Table 9 displays the results. The key explanatory variables in these tests are indicator variables associated with whether there was an earthquake (Panel A) or lottery winner (Panel B) within a 30-mile radius of the centroid of the DMA during week *t*. The control variables and fixed effects terms are identical to the tests from Table 8.

For the earthquake events, there is a sharp positive effect the week following the earthquake. The effects become insignificant for the second week following the event. There is no significant effect during the weeks prior, which squares with the fact that earthquakes are difficult to predict. The lack of significance for the week of the earthquake could be associated with the timing of the earthquake and the estimation window for the *SVI* measure. Regardless, the results suggest a slightly delayed response in internet search volumes.

For the lottery events, the results are mostly significant in the weeks prior as well as after the announcement of the lottery winner. However, the effects are most acute in the week of and week following the announcement of the lottery winner. The significant effects in the weeks prior could be due to greater discussion of the increasing jackpot in the weeks prior to the winning one. The delayed response in the internet search volumes could be due to additional information that is released in the days following the announcement of a winning jackpot.

6.2. Empirical Specifications and Results

6.2.1. Earthquakes

We begin with tests that update our prior findings on the relationship between nearby earthquakes and survey crash probabilities. Our tests focus on moderate earthquakes that are likely to be felt by of individuals located close to the epicenter but unlikely to be associated with structural damage. Earthquakes of stronger magnitudes may have direct effects on economic conditions, which in turn can conceivably affect stock market conditions. However, the empirical evidence on this link is mixed. Ferreira and Karali (2015) show that stock markets do not exhibit a significant reaction within five days of strong earthquakes. In contrast, Shiller (2000b) points to the Tokyo stock market reaction to the earthquake in Kobe, Japan on January 17, 1995, which measured at 7.2 on the Richter scale. The Nikkei index fell by 8 percent overall within ten days of the earthquake, though significant price movements did not transpire until after one week. This market reaction far exceeded an official estimate of the economic damage, which was eventually set at approximately $100 billion. While circumstantial, the delay in and magnitude of the market reaction to the Kobe earthquake suggest a mix of fundamental and sentiment factors.

To address potential confounding effects of the economic impact of stronger earthquakes, we distinguish between moderate magnitude, or earthquakes with a magnitude between 2.5 up to 5.5, and strong magnitude, or earthquakes with a magnitude above 5.5. The cutoffs are based upon information from the USGS, which classifies earthquakes with magnitudes above 2.5 as physically detectable, and earthquakes with magnitudes above 5.5 as inflicting at least minor damage to buildings and other structures.

Using the investor survey and earthquake data, we estimate the following regression model:

$$CPS_{i,t} = \beta_1 \times Moderate\ Earthquake_{i,t-30,t-1} + \beta_2 \times Severe\ Earthquake_{i,t-30,t-1} + \beta_3 \times Earthquake\ Frequency_{i,t-30} + \boldsymbol{\beta} \times \boldsymbol{X_{t-1}} + \boldsymbol{\tau_t} + \boldsymbol{\delta_t} + \epsilon_{i,t} \quad (2)$$

The dependent variable is *CPS*. $Moderate\ Earthquake$ is a dummy variable associated with whether a moderate earthquake occurred within 30-miles of the investor sometime within the past 30 days. Similarly, $Severe\ Earthquake$ is a dummy variable associated with whether a severe earthquake occurred within 30 miles of the investor within the past 30 days. $Earthquake\ Frequency$ is the natural log of one plus the number of earthquakes per year that occurred within a 30 miles radius of the investor from 1900 to 2006. While the earthquake events are unlikely to be related to market conditions, we nonetheless include a number of control variables ($\boldsymbol{X}$): previous day market return, past month market returns, past month average investor survey crash probability, past month market return volatility, and VIX. We also include month ($\boldsymbol{\tau}$) and day-of-week ($\boldsymbol{\delta}$) fixed effects to account for potential seasonality in the crash probability responses. Robust standard errors clustered on the ZIP code and date levels are used to assess statistical significance.

Table 10 presents the results. Columns (1) and (2) display the results for the individual investor subsample, while Columns (3) and (4) present the results for the institutional investor subsample. Column (5) present the results for the pooled sample. Columns (1) and (3) only include the earthquake related terms in the model for comparison. We find that moderate magnitude earthquakes have a positive and statistically significant association with investor crash probabilities, but only for individual investors. The coefficients on the severe magnitude earthquake terms are all statistically insignificant. The results remain mostly unchanged after controlling for the market-related control variables, which is consistent with our assumption that market conditions are uncorrelated with the earthquake events. Untabulated robustness checks show the results remain significant when eliminating either the top 1st sample percentile in terms of earthquake frequency or removing investors located in California from the sample.

The economic magnitudes of the earthquake effect are sizable for the individual investor subsample. Based upon Column (2), respondents located in close proximity to an earthquake report crash probabilities 2.1 percentage points higher than those that do not. The marginal effect is 16.9% of the sample mean and 13.0% of the total sample variation for *CPS*. While the findings suggest that the high average response by individual investors may be attributable to some extent to behavior factors related to availability and affect biases, it does not explain why the average response for institutional investors are also high.

6.2.2. Lottery Winners

We next present tests on the neighbors of lottery jackpot winners. We adapt Equation (1) by replacing the earthquake terms for the lottery measure:

$$CPS_{i,t} = \gamma_1 \times Neighboring\ Lottery\ Winner_{i,t-30,t-1}$$
$$+\gamma_2 \times Lottery\ Winner\ Frequency_{i,t-30} + \boldsymbol{\gamma} \times \boldsymbol{X_{t-1}} + \boldsymbol{\tau_t} + \boldsymbol{\delta_t} + \zeta_{i,t} \quad (3)$$

$Neighboring\ Lottery\ Winner$ is a dummy variable associated with whether a winning lottery ticket for MegaMillions or Powerball was purchased at a store located within a 30-mile radius of the investor sometime in the 30 days prior to the survey response. $Lottery\ Winner\ Frequency$ is the number of times a winning lottery ticket was purchased within a 30-mile radius of the investor previously. The control variables ($\boldsymbol{X}$) and the fixed effects terms are identical to those in Equation (1). We exclude investors residing in states that do not offer either contest at the time that the survey is filled. Robust standard errors clustered on the ZIP code and date levels are used to assess statistical significance.

Table 11 presents the results. The table is formatted similarly to Table 10. We find that neighbors of lottery winners have a negative and significant association with investor crash probabilities, but only for individual investors. The significance and magnitude of the estimates remain unchanged when removing the control variables. Based upon Column (2), respondents located in close proximity to a lottery winner report crash probabilities 2.5 percentage points lower than those that do not. The marginal effect is 15.4% of the total sample variation for the crash probability spread. Untabulated robustness checks show that the results are not sensitive to expanding the sample to states not participating in the lottery.

6.2.3. COVID-19 Pandemic

Finally, we present tests on the COVID-19 pandemic. As with the previous section, we adapt Equation (1) to include the COVID-19 public health and policy intervention terms. These tests focus on the 2019 through 2020 sample period.

$$\begin{aligned} CPS_{i,t} = {} & \theta_1 \times Change\ in\ County\ Cases\ (p.c.)_{i,t-7,t-1} \\ & + \theta_2 \times County\ Cases\ (p.c.)_{i,t-1} \\ & + \theta_3 \times Change\ in\ County\ Deaths\ (p.c.)_{i,t-7,t-1} \\ & + \theta_4 \times County\ Deaths\ (p.c.)_{i,t-1} \\ & + \boldsymbol{\theta_5} \times \boldsymbol{County\ COVID\ Restrictions_{i,t-1}} \\ & + \boldsymbol{\theta} \times \boldsymbol{X_{t-1}} + \boldsymbol{\mu_t} + \boldsymbol{\delta_t} + \boldsymbol{\gamma_i} + \psi_{i,t} \end{aligned} \tag{4}$$

The main explanatory variables are: $Change\ in\ County\ Cases\ (p.c.)_{i,t-7,t-1}$ is the one-week change in county-level COVID cases per capita, $County\ Cases\ (p.c.)_{i,t-1}$ is the total county-level COVID cases per capita, $Change\ in\ County\ Deaths\ (p.c.)_{i,t-7,t-1}$ is the one-week change in COVID deaths per capital, $County\ Deaths\ (p.c.)_{i,t-1}$ is the total county-level COVID deaths per capita, and various county-level COVID restrictions. $\boldsymbol{County\ COVID\ Restrictions_{i,t-1}}$ is a vector that includes: stay-at-home orders, retail establishment closures, gathering restrictions, and school closures. The control variables ($\boldsymbol{X}$) are identical to those in Equation (1). In addition to day-of-week fixed effects, we also include year-month and state levels fixed effects. The year-month fixed effects account for nationwide factors that affect aggregate economic and stock market conditions. We include state fixed effects in order to account for statewide public health policies. We use robust standard errors clustered on the ZIP code and date levels to calculate them.[36]

Table 12 presents the results for individual investors. Columns (1) through (5) enters each set of explanatory variables individually. Columns (6) and (7) display the results with all the explanatory variables in the same model. Columns (1) through (6) displays the results without the control variables and using $\pi_{i,t}$ as the dependent variable. The results show that the change in county-level COVID cases and total county-level COVID deaths are positive and statistically significant at the 1% level. Additionally, the county-level COVID restrictions on school closures is negative and statistically significant at the 1% level. Column (7) includes the control variables and uses CPS as the dependent variable. In this specification, both changes in COVID cases and total COVID deaths remain statistically significant while the school restriction term becomes statistically insignificant. The results are economically significant as well. An one-standard deviation increase in changes in COVID cases (total COVID deaths) increases CPS by 1.2 (1.7) percentage points, which represent 7.2 (10.3) percent of the total sample variation in CPS.

Table 13 presents the results for institutional investors. Unlike the tests on the individual investor subsample, only changes in COVID deaths and COVID restrictions on school closures are statistically significant. Interestingly, the coefficient on changes in COVID deaths is negative, indicating that an increase in the number of COVID deaths decreases CPS. However, the results lose statistical significance after inclusion of the control variables to the 10% level. An one-standard deviation increase in changes in COVID deaths decreases CPS by 1.3 percentage points, which represent 8.2 percent of the total sample variation in CPS.

Overall, all three of the natural experiments considered in our analysis provide evidence consistent with availability and affect biases and are consistent with the hypothesis that availability is mediated by emotion not semantic association. While we do not anticipate that they represent

[36] We do not bootstrap *p*-values for these tests given that the likelihood of repeated winners for a ZIP code is low within our sample.

systematic variation in crash probabilities, they provide evidence on the plausibility of broader factors influencing investor crash beliefs and why the survey-based crash probabilities are large.

7. Conclusion

The high, sustained subjective crash probabilities in the Shiller Investor Confidence Survey are consistent with the rare disaster explanation for the equity risk premium. We turn to the risk-as-feelings literature for models to explore emotional factors potentially influencing probabilities. Advances in machine-learning allow us to use the narrative content of the Investor Confidence Survey to test the laboratory findings about emotions and beliefs with real-world data that have potentially interesting economic implications. We find strong evidence that the non-fundamental probability component of predictions about a catastrophic stock market crash are associated with negative emotional valence that has a specific polarity contrasting optimism with anxiety and related emotions.

Our comparison of individual sentiment with media sentiment provide a strong methodological justification for the now-widespread use of the latter as an instrument for the former. While this relationship has been presumed, analysis of the survey narratives supports the presumption. At the same time the survey evidence also suggests an important heterogeneity related to emotion which in turn can have implications for variation in equity market participation and trade.

The temporal variation in the subjective component of crash beliefs is also relevant to the asset pricing literature. It is substantially more volatile than the fundamental, option-derived crash beliefs. This suggests that irrational factors – or at least some factors not manifest in option markets – cause large variation in average and median beliefs.

Finally, our results also raise questions about whether preferences and beliefs can be empirically separated – and indeed whether they operate independently in the mind of investors. Respondents report a high probability of a crash when they are imagining the extreme disutility they would experience from that event if it came to pass. This convolution presents further modeling possibilities for behavioral research.

References

Agarwal, S., Mikhed, V. and Scholnick, B., 2018. Does the Relative Income of Peers Cause Financial Distress? *Evidence from Lottery Winners and Neighboring Bankruptcies* (No. 18-16).

Akhtar, S., Faff, R., Oliver, B. and Subrahmanyam, A., 2012. Stock salience and the asymmetric market effect of consumer sentiment news. *Journal of Banking & Finance*, 36(12), pp.3289-3301.

Barber, B.M. and Odean, T., 2008. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies*, 21(2), pp.785-818.

Barberis, N., 2013. The psychology of tail events: Progress and challenges. *The American Economic Review*, 103(3), pp.611-616.

Barone-Adesi, G., Mancini, L. and Shefrin, H., 2015. Sentiment, Risk Aversion, and Time Preference. *Swiss Finance Institute Research Paper*, (12-21).

Barro, R.J., 2006. Rare disasters and asset markets in the twentieth century. *The Quarterly Journal of Economics*, *121*(3), pp.823-866.

Bates, D.S., 2000. Post‘'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*, 94(1), pp.181-238.

Baumeister, R.F., Bratslavsky, E., Finkenauer, C. and Vohs, K.D., 2001. Bad is stronger than good. *Review of general psychology*, 5(4), p.323.

Bechara, A. and Damasio, A.R., 2005. The somatic marker hypothesis: A neural theory of economic decision. *Games and economic behavior*, *52*(2), pp.336-372.

Bechara, A., Damasio, H., Tranel, D. and Damasio, A.R., 1997. Deciding advantageously before knowing the advantageous strategy. *Science*, *275*(5304), pp.1293-1295.

Ben-David, I., Graham, J.R. and Harvey, C.R., 2007. *Managerial overconfidence and corporate policies* (No. w13711). National Bureau of Economic Research.

Ben-David, I., Graham, J.R. and Harvey, C.R., 2013. Managerial miscalibration. *The Quarterly journal of economics*, *128*(4), pp.1547-1584.

Berkman, H., Jacobsen, B. and Lee, J.B., 2011. Time-varying rare disaster risk and stock returns. *Journal of Financial Economics*, 101(2), pp.313-332.

Bianchi, F., Ludvigson, S.C. and Ma, S., 2022. Belief distortions and macroeconomic fluctuations. *American Economic Review*, *112*(7), pp.2269-2315.

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

Bollerslev, T., Todorov, V. and Xu, L., 2015. Tail risk premia and return predictability. *Journal of Financial Economics*, 118(1), pp.113-134.

Bossaerts, P., 2021. How neurobiology elucidates the role of emotions in financial decision-making. *Frontiers in Psychology*, *12*, p.697375.

Bossaerts, P., Fattinger, F., Rotaru, K. and Xu, K., 2023. Emotional engagement and trading performance. *Management Science*.

Brown, S.J., Goetzmann, W.N. and Ross, S.A., 1995. Survival. *The Journal of Finance*, *50*(3), pp.853-873.

Brunnermeier, M., Farhi, E., Koijen, R.S., Krishnamurthy, A., Ludvigson, S.C., Lustig, H., Nagel, S. and Piazzesi, M., 2021. Perspectives on the future of asset pricing. *The review of financial studies*, *34*(4), pp.2126-2160.

Chang, Y., Wang, X., Wang, J., Wu, Y., Yang, L., Zhu, K., Chen, H., Yi, X., Wang, C., Wang, Y. and Ye, W., 2023. A survey on evaluation of large language models. ACM Transactions on Intelligent Systems and Technology.

Choi, J.J. and Robertson, A.Z., 2020. What matters to individual investors? Evidence from the horse's mouth. *The Journal of Finance*, *75*(4), pp.1965-2020.

Cocco, J.F., Gomes, F. and Lopes, P., 2020. Evidence on expectations of household finances. *Available at SSRN 3362495*.

Cuculiza, C., Antoniou, C., Kumar, A. and Maligkris, A., 2021. Terrorist attacks, analyst sentiment, and earnings forecasts. *Management science*, *67*(4), pp.2579-2608.

Egan, M., MacKay, A. and Yang, H., 2022. Recovering investor expectations from demand for index funds. *The Review of Economic Studies*, *89*(5), pp.2559-2599.

Engelberg, J.E. and Parsons, C.A., 2011. The causal impact of media in financial markets. *The Journal of Finance*, 66(1), pp.67-97.

Da, Z., Engelberg, J. and Gao, P., 2011. In Search of Attention. *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. *Review of Financial Studies*, 28(1), pp.1-32.

Dieckmann, N.F., Slovic, P. and Peters, E.M., 2009. The use of narrative evidence and explicit likelihood by decisionmakers varying in numeracy. *Risk Analysis*, *29*(10), pp.1473-1488.

Duxbury, D., Gärling, T., Gamble, A. and Klass, V., 2020. How emotions influence behavior in financial markets: a conceptual analysis and emotion-based account of buy-sell preferences. *The European Journal of Finance*, *26*(14), pp.1417-1438.

Ferreira, S. and Karali, B., 2015. Do earthquakes shake stock markets? *PloS one*, 10(7), p.e0133319.

Fisher, I., 1930. The theory of interest. *New York*.

Gabaix, X., 2012. Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance. *The Quarterly journal of economics*, *127*(2), pp.645-700.

Gao, G. and Song, Z., 2015. Rare disaster concerns everywhere. Available at SSRN 2606345.

Garcia, D., 2013. Sentiment during recessions. *The Journal of Finance*, *68*(3), pp.1267-1300.

Garcia, D., 2014. The kinks of financial journalism. Available at SSRN 2517591.

Garrett, I., Kamstra, M.J. and Kramer, L.A., 2005. Winter blues and time variation in the price of risk. *Journal of Empirical Finance*, *12*(2), pp.291-316.

Kamstra, M.J., Kramer, L.A., Levi, M.D. and Wermers, R., 2017. Seasonal asset allocation: Evidence from mutual fund flows. *Journal of Financial and Quantitative Analysis*, *52*(1), pp.71-109.

Giglio, S., Maggiori, M., Stroebel, J. and Utkus, S., 2021. Five facts about beliefs and portfolios. *American Economic Review*, *111*(5), pp.1481-1522.

Goetzmann, W.N. and Kim, D., 2018. Negative bubbles: What happens after a crash. *European Financial Management*, *24*(2), pp.171-191.

Goetzmann, W.N., Kim, D., Kumar, A. and Wang, Q., 2014. Weather-induced mood, institutional investors, and stock returns. *Review of Financial Studies*, 28(1), pp. 73-111.

Goetzmann, W.N., Kim, D. and Shiller, R.J., 2016. *Crash beliefs from investor surveys* (No. w22143). National Bureau of Economic Research.

Goetzmann, W.N. and Zhu, N., 2005. Rain or shine: where is the weather effect?. *European Financial Management*, *11*(5), pp.559-578.

Grabe, M.E. and Kamhawi, R., 2006. Hard wired for negative news? Gender differences in processing broadcast news. *Communication Research*, *33*(5), pp.346-369.

Greenwood, R. and Shleifer, A., 2014. Expectations of returns and expected returns. *Review of Financial Studies*, 27(3), pp.714-746.

Bin Hasan, S., Kumar, A. and Taffler, R., 2023. Anxiety, Excitement, and Asset Prices. *University of Miami Business School Research Paper*, (3902654).

Hilbig, B.E., 2009. Sad, thus true: Negativity bias in judgments of truth. *Journal of Experimental Social Psychology*, 45(4), pp.983-986.

Jiang, A.Q., Sablayrolles, A., Mensch, A., Bamford, C., Chaplot, D.S., Casas, D.D.L., Bressand, F., Lengyel, G., Lample, G., Saulnier, L. and Lavaud, L.R., 2023. Mistral 7B. arXiv preprint arXiv:2310.06825.

Johnson, E.J. and Tversky, A., 1983. Affect, generalization, and the perception of risk. *Journal of personality and social psychology*, 45(1), p.20.

Kahneman, D. and Tversky, A., 1982. The psychology of preferences. *Scientific American*, *246*(1), pp.160-173.

Kamstra, M.J., Kramer, L.A. and Levi, M.D., 2003. Winter blues: A SAD stock market cycle. *American economic review*, *93*(1), pp.324-343.

Kant, N., Puri, R., Yakovenko, N. and Catanzaro, B., 2018. Practical text classification with large pre-trained language models. arXiv preprint arXiv:1812.01207.

Karlsson, N., Loewenstein, G. and Seppi, D., 2009. The ostrich effect: Selective attention to information. *Journal of Risk and uncertainty*, *38*(2), pp.95-115.

Keller, C., Siegrist, M. and Gutscher, H., 2006. The role of the affect and availability heuristics in risk communication. *Risk analysis*, *26*(3), pp.631-639.

Killeen, B.D., Wu, J.Y., Shah, K., Zapaishchykova, A., Nikutta, P., Tamhane, A., Chakraborty, S., Wei, J., Gao, T., Thies, M. and Unberath, M., 2020. A county-level dataset for informing the United States' response to COVID-19. arXiv preprint arXiv:2004.00756.

Kräussl, R. and Mirgorodskaya, E., 2014. News media sentiment and investor behavior (No. 492). *Center for Financial Studies* (CFS).

Kuhn, P., Kooreman, P., Soetevent, A. and Kapteyn, A., 2011. The effects of lottery prizes on winners and their neighbors: Evidence from the Dutch postcode lottery. *American Economic Review*, *101*(5), pp.2226-47.

Kuhnen, C.M. and Knutson, B., 2005. The neural basis of financial risk taking. *Neuron*, *47*(5), pp.763-770.

Kuhnen, C.M. and Knutson, B., 2011. The influence of affect on beliefs, preferences, and financial decisions. *Journal of Financial and Quantitative Analysis*, *46*(3), pp.605-626.

Lerner, J.S. and Keltner, D., 2000. Beyond valence: Toward a model of emotion-specific influences on judgement and choice. *Cognition & emotion*, *14*(4), pp.473-493.

Lerner, J.S., Li, Y., Valdesolo, P. and Kassam, K.S., 2015. Emotion and decision making. *Annual review of psychology*, *66*, pp.799-823.

Li, C., Wang, J., Zhang, Y., Zhu, K., Hou, W., Lian, J., Luo, F., Yang, Q. and Xie, X., 2023. Large language models understand and can be enhanced by emotional stimuli. arXiv preprint arXiv:2307.11760.

Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M. and Combs, B., 1978. Judged frequency of lethal events. *Journal of experimental psychology: Human learning and memory*, *4*(6), p.551.

Lindquist, K.A., MacCormack, J.K. and Shablack, H., 2015. The role of language in emotion: Predictions from psychological constructionism. Frontiers in psychology, 6, p.121301.

Loewenstein, G.F., Weber, E.U., Hsee, C.K. and Welch, N., 2001. Risk as feelings. *Psychological bulletin*, *127*(2), p.267.

Loughran, T. and McDonald, B., 2011. When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-K's. *The Journal of Finance,* 66, pp. 35-65.

Nagel, S. and Xu, Z., 2022. Asset pricing with fading memory. *The Review of Financial Studies*, *35*(5), pp.2190-2245.

Manning, C.D., Raghavan, P. and Schutze, H., 2008. An information to information retrieval. *Cambridge University Press,* Cambridge.

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

Newey, W.K. and West, K.D., 1994. Automatic lag selection in covariance matrix estimation. *The Review of Economic Studies*, *61*(4), pp.631-653.

Rietz, T.A., 1988. The equity risk premium a solution. *Journal of monetary Economics*, *22*(1), pp.117-131.

Rozin, P. and Royzman, E.B., 2001. Negativity bias, negativity dominance, and contagion. *Personality and social psychology review*, 5(4), pp.296-320.

Russell, J.A., 1980. A circumplex model of affect. Journal of personality and social psychology, 39(6), p.1161.

Santa-Clara, P. and Yan, S., 2010. Crashes, volatility, and the equity premium: Lessons from S&P 500 options. *The Review of Economics and Statistics*, 92(2), pp.435-451.

Semin, G.R. and Fiedler, K., 1988. The cognitive functions of linguistic categories in describing persons: Social cognition and language. *Journal of personality and Social Psychology*, *54*(4), p.558.

Seo, S.B. and Wachter, J.A., 2013. Option prices in a model with stochastic disaster risk (No. w19611). *National Bureau of Economic Research*.

Strapparava, C. and Mihalcea, R., 2007, June. Semeval-2007 task 14: Affective text. In *Proceedings of the fourth international workshop on semantic evaluations (SemEval-2007)* (pp. 70-74).

Shiller, R.J., 1987. Investor behavior in the October 1987 stock market crash: Survey evidence.

Shiller, R.J.,1990. Speculative prices and popular models. *The Journal of Economic Perspectives*, pp.55-65.

Shiller, R. J., 2000a. Measuring bubble expectations and investor confidence. *Journal of Psychology and Financial Markets*, 1, pp. 49-60.

Shiller, R. J., 2000b. Irrational Exuberance. *Princeton University Press*, Princeton.

Sicherman, N., Loewenstein, G., Seppi, D.J. and Utkus, S.P., 2015. Financial attention. *The Review of Financial Studies*, *29*(4), pp.863-897.

Slovic, P. and Peters, E., 2006. Risk perception and affect. *Current directions in psychological science*, *15*(6), pp.322-325.

Slovic, P., Finucane, M.L., Peters, E. and MacGregor, D.G., 2007. The affect heuristic. *European journal of operational research*, *177*(3), pp.1333-1352.

Taffler, R.J., Agarwal, V.I.N.E.E.T. and Obring, M.A.X.I.M.I.L.I.A.N., 2021. Narrative Economics and Market Bubbles. *European Financial Management, European Financial Management Association*, *21*(4).

Tetlock, P.C., 2007. Giving content to investor sentiment: The role of media in the stock market. *The Journal of Finance*, 62(3), pp.1139-1168.

Tiedens, L.Z. and Linton, S., 2001. Judgment under emotional certainty and uncertainty: the effects of specific emotions on information processing. *Journal of personality and social psychology*, *81*(6), p.973.

Trope, Y., Liberman, N. and Wakslak, C., 2007. Construal levels and psychological distance: Effects on representation, prediction, evaluation, and behavior. *Journal of consumer psychology*, *17*(2), pp.83-95.

Tsai, J. and Wachter, J.A., 2015. Disaster risk and its implications for asset pricing (No. w20926). *National Bureau of Economic Research*.

Tversky, A. and Kahneman, D., 1973. Availability: A heuristic for judging frequency and probability. *Cognitive psychology*, 5(2), pp.207-232.

Vissing-Jørgensen, A. and Attanasio, O.P., 2003. Stock-market participation, intertemporal substitution, and risk-aversion. *American Economic Review*, *93*(2), pp.383-391.

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.

Welch, I., 2015. The Importance of Disaster Risk. Available at SSRN 2596909.

Yuan, K., Zheng, L. and Zhu, Q., 2006. Are investors moonstruck? Lunar phases and stock returns. *Journal of Empirical Finance*, *13*(1), pp.1-23.

Yuan, Y., 2015. Market-wide attention, trading, and stock returns. *Journal of Financial Economics*, *116*(3), pp.548-564.

Yang, T., Yu, T.R. and Zhao, H., 2024. Uncovering the relationship between incidental emotions toward a disaster and stock market fluctuations: Evidence from the US market. *Decision Support Systems*, p.114213.

Zajonc, R.B., 1980. Feeling and thinking: Preferences need no inferences. *American psychologist*, *35*(2), p.151.

Figure 1
Average Annual Crash Probabilities from 1996-2020

This figure displays the average annual probabilities from 1996-2020 for the survey respondents of a crash in the next six months on the scale of 10/19/1987 or 10/28/1929. Also displayed are the option-implied probability of a drop of at least 15% in the S&P 500 index over the next six months, annualized volatility of the daily DJIA, the largest negative return in each year (right axis) and the VIX (divided by 100).

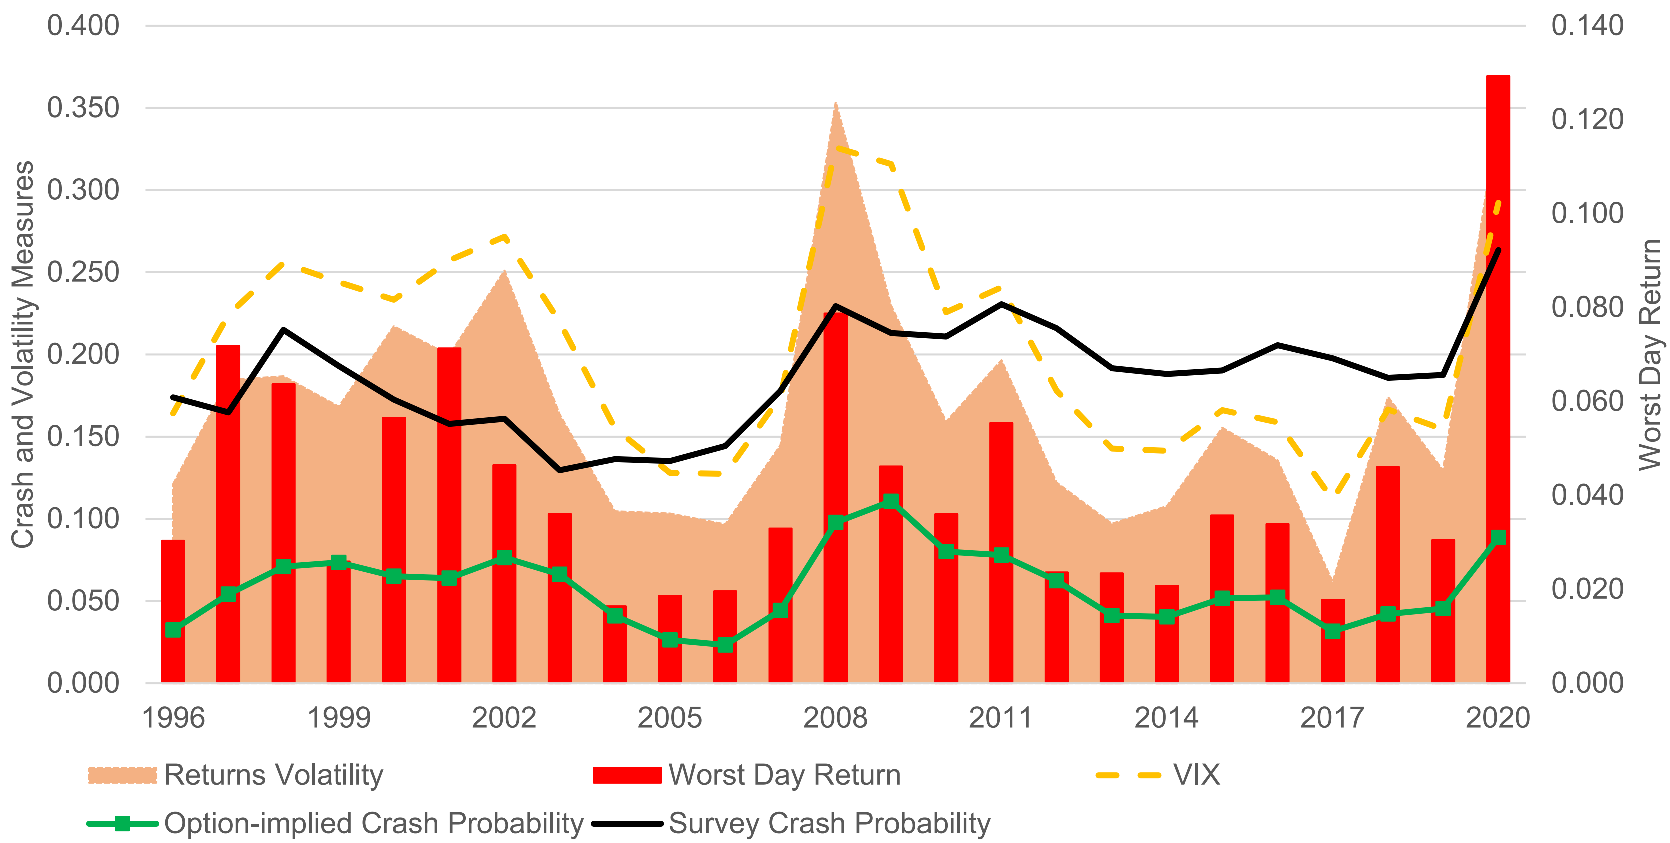

Figure 2
Crash Probability Distribution

This figure displays the 50th and 90th percentiles of the monthly probabilities from 1996-2020 for the survey respondents of a crash in the next six months on the scale of 10/19/1987 or 10/28/1929 are displayed in black, solid and dashed, respectively. The option-implied probability of a drop of at least 15% in the S&P 500 index over the next six months are displayed in red.

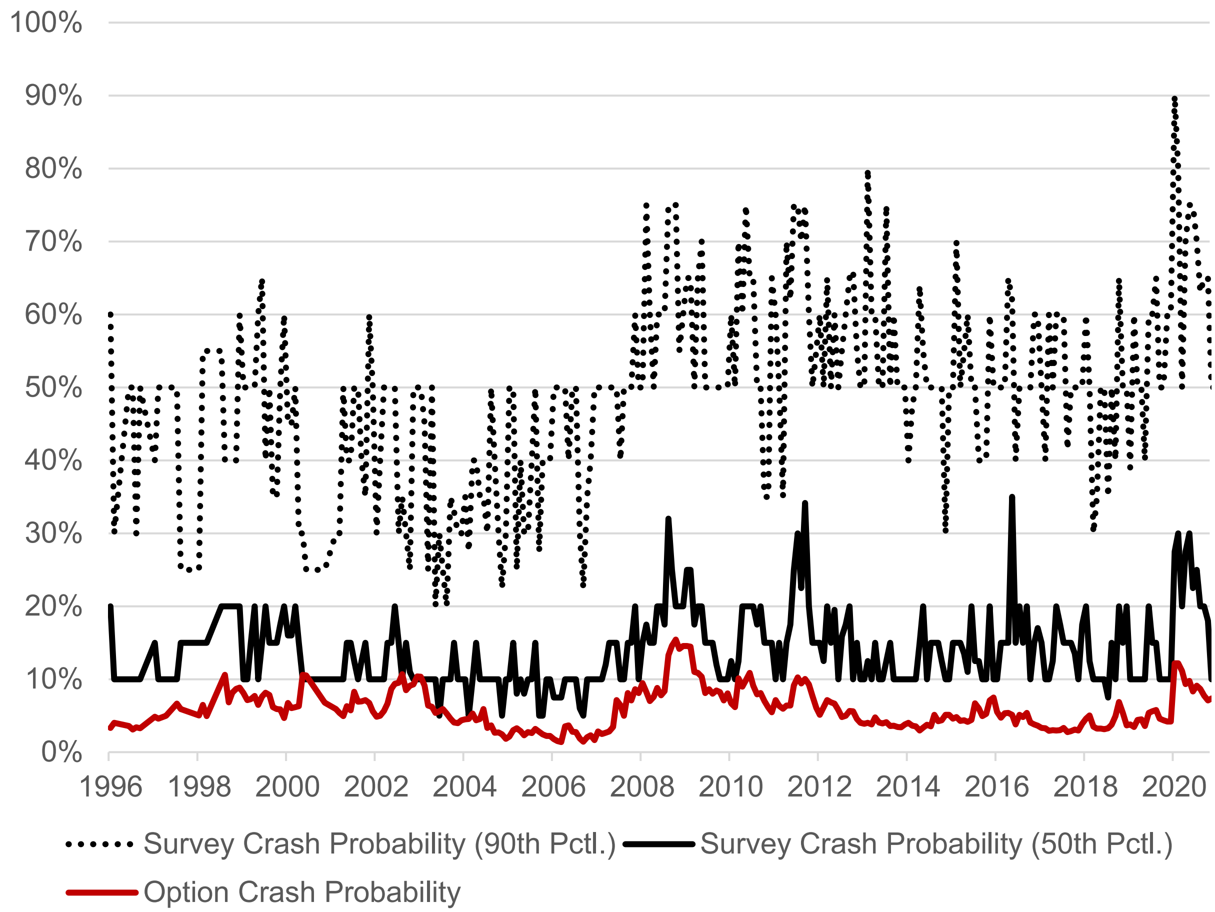

Figure 3
Crash Probability Spread

This figure displays the average monthly crash probability spread, calculated as the difference in (a) the probabilities for the survey respondents of a crash in the next six months on the scale of 10/19/1987 or 10/28/1929 and (b) the option-implied probability of a drop of at least 15% in the S&P 500 index over the next six months.

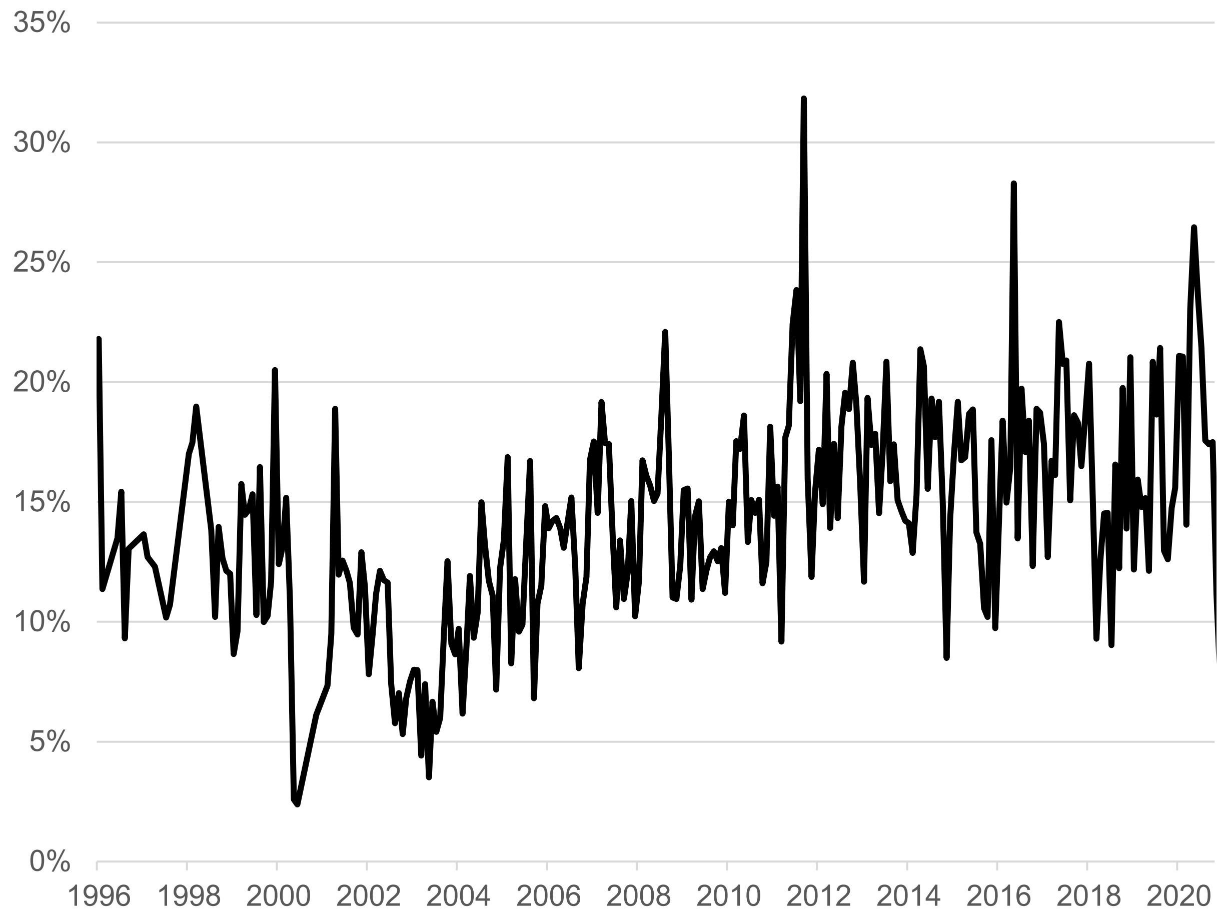

Figure 4
Moderate Earthquake Geographical Distribution

This figure displays the choropleth map of the average frequency of earthquakes from 2007 through 2020 of counties within a 30 mile radius of the epicenter of an earthquake, where darker shades correspond with greater frequency.

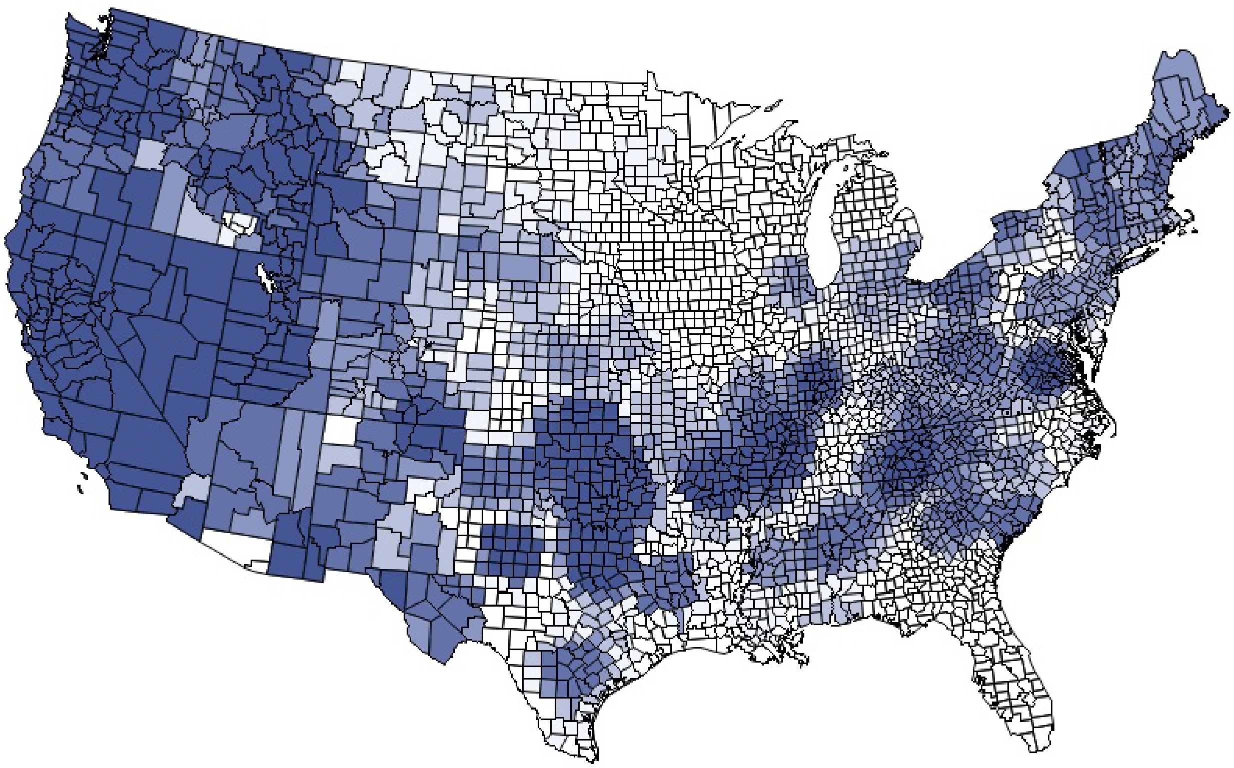

Figure 5
Lottery Winner Geographic Distribution

This figure displays the choropleth map of the average frequency of lottery winners from 2007 through 2020 of counties within a 30 mile radius of a lottery winner, where darker shades correspond with greater frequency.

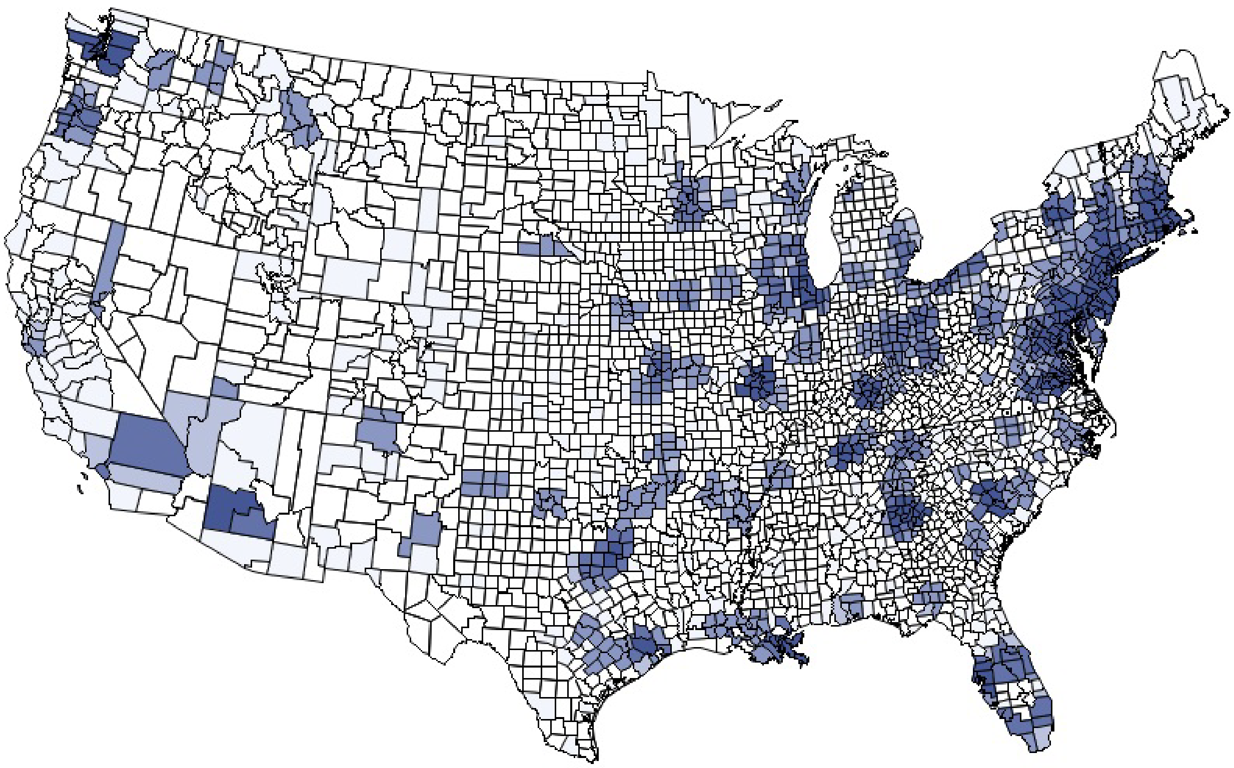

Figure 6
Individual versus Institutional Investor Crash Probabilities

This figure displays the annual average probabilities from 2000-2020 for individual (red) and institutional (blue) investor survey respondents of a crash in the next six months on the scale of 10/19/1987 or 10/28/1929.

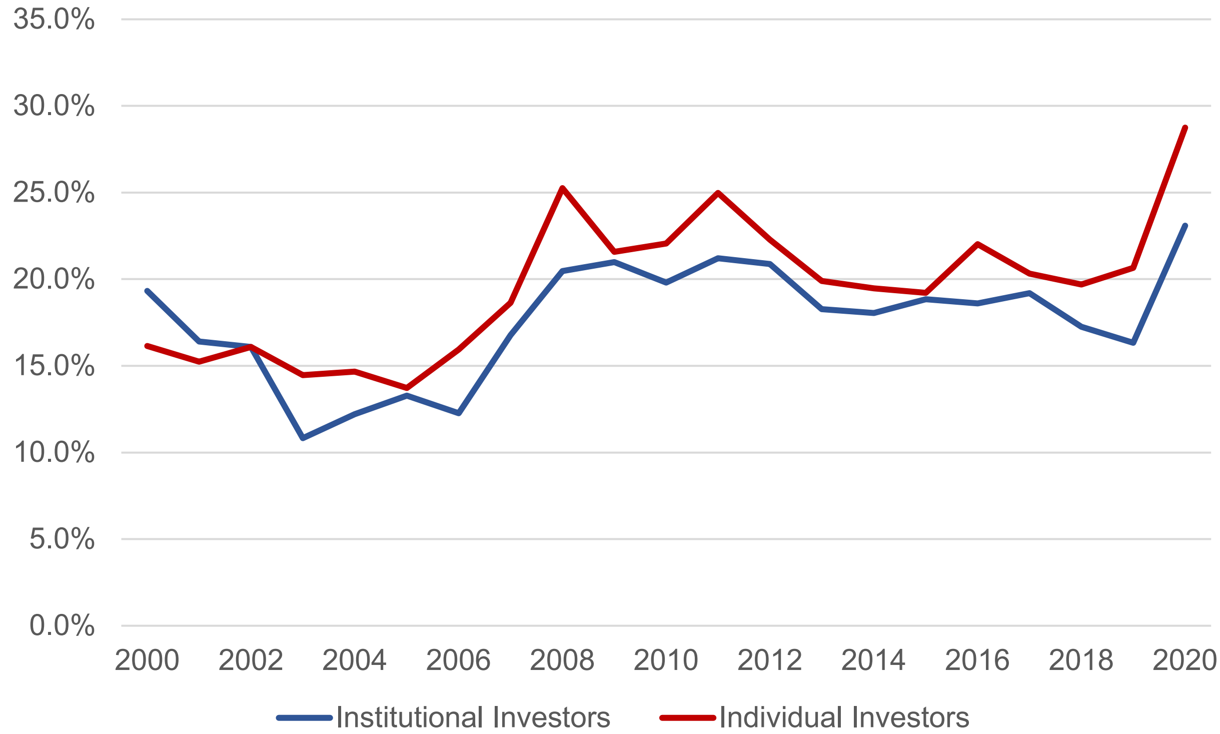

Table 1
Summary Statistics

The table displays variable descriptions (Panel A) and summary statistics (Panel B) of the key variables used in the analysis. The variables are collected from Robert Shiller's Investor Confidence Surveys (ICS); data obtained from Martin (2017) (M17); the Center for Research on Security Prices (CRSP); the Chicago Board of Options Exchange (CBOE); the United States Geological Services (USGS); press releases from Powerball and MegaMillions (PBMM); COVID-19 policy intervention data (COV); and Proquest (PRO).

Panel A: Summary Statistics

Variable Name	**Description**	**Source**
π^{Survey}_t	The 7-day moving average of the crash probability reported by the survey respondents on date *t*.	ICS
π^{Option}_t	The option-implied crash probability at the close of date *t*.	M17
CPS_t	The spread between the average 7-day moving average of the crash probability reported by the survey respondents on date *t* and the option-implied crash probability at the close of date *t*.	ICS, M17
R^M_t	Total return on date *t* based upon the CRSP-VW (NYSE/AMEX/Nasdaq/Arca) index.	CRSP
VIX_t	The closing value as of date *t* of the Volatility Index, which is based upon S&P 500 index options.	CBOE
Nearby Moderate $Earthquake_{i,t}$	Dummy variable that takes value one if the survey respondent is located within a 30-mile radius of an earthquake with a seismic magnitude from 2.5 to 5.5 within the past 30 days, zero otherwise.	USGS
Nearby Severe $Earthquake_{i,t}$	Dummy variable that takes value one if the survey respondent is located within a 30-mile radius of an earthquake with a seismic magnitude above 5.5 within the past 30 days, zero otherwise.	USGS
Historical Earthquake $Frequency_{i,t}$	Annualized frequency of earthquakes with seismic magnitude between 2.5 and above whose epicenter is located within a 30-mile radius of the survey respondent from 1900-2006.	USGS
Nearby Lottery $Winner_{i,t}$	Dummy variable that takes value one if the survey respondent is located within a 30-mile radius of a store that sold a winning Powerball of MegaMillions ticket within the past 30 days, zero otherwise.	PBMM

Table 1 (cont.)

Historical Lottery Frequency$_{i,t}$	Total frequency of Powerball and MegaMilions winners that is located within a 30-mile radius of the survey respondent as of date t.	USGS
Change in County Cases (p.c.)$_{i,t}$	Percentage change in the number of county-level COVID-19 cases-per-capital over the past week.	JHU
County Cases (p.c.)$_{i,t-7}$	The number of county-level COVID-19 cases-per-capital in the past week.	JHU
Change in County Deaths (p.c.)$_{i,t}$	Percentage change in the number of county-level COVID-19 deaths-per-capital over the past week.	JHU
County Deaths (p.c.)$_{i,t-7}$	The number of county-level COVID-19 deaths-per-capital in the past week.	JHU
County Stay-at-Home Order$_{i,t}$	Dummy variable that takes value one if the survey respondent resides in a county that has an active stay-at-home order as of date t.	COV
County Retail Establishment Order$_{i,t}$	Dummy variable that takes value one if the survey respondent resides in a county that has an active retail establishment restriction order as of date t.	COV
County Gatherings Order$_{i,t}$	Dummy variable that takes value one if the survey respondent resides in a county that has an active gathering restriction order as of date t.	COV
County School Order$_{i,t}$	Dummy variable that takes value one if the survey respondent resides in a county that has an active school restriction order as of date t.	COV
Investor Sentiment$_t$	Sentiment of the textual responses by the survey respondent *i* on date *t*.	ICS
News Sentiment$_t$	The sentiment of news articles related to the stock market on date *t*.	PRO

Panel B: Summary Statistics

	N	Mean	Standard Deviation	25th Percentile	50th Percentile	75th Percentile
π^{Survey}_t	16,177	19.26%	19.75%	5.00%	10.00%	25.00%
π^{Option}_t	16,177	6.16%	2.88%	3.95%	5.62%	7.92%
CPS$_t$	16,177	13.11%	19.54%	-0.72%	6.32%	20.26%
R^M_t	16,177	0.04%	1.24%	-0.46%	0.09%	0.60%
VIX$_t$	16,177	20.454	9.795	13.850	18.160	23.790
Nearby Moderate Earthquake$_{i,t}$	11,698	8.19%	27.42%	0.00%	0.00%	0.00%

Table 1 (cont.)

Nearby Severe Earthquake$_{i,t}$	11,698	0.11%	3.33%	0.00%	0.00%	0.00%
Historical Earthquake Frequency$_{i,t}$	11,698	126.902	516.571	0.000	1.000	8.000
Nearby Lottery Winner$_{i,t}$	11,698	2.87%	16.70%	0.00%	0.00%	0.00%
Historical Lottery Frequency$_{i,t}$	11,698	3.463	7.727	0.000	1.000	3.000
Change in County Cases (p.c.)$_{i,t}$	1,162	0.00%	0.02%	0.00%	0.00%	0.00%
County Cases (p.c.)$_{i,t-7}$	1,162	2.29%	9.72%	0.00%	0.00%	1.46%
Change in County Deaths (p.c.)$_{i,t}$	1,162	0.001	0.006	0.000	0.000	0.000
County Deaths (p.c.)$_{i,t-7}$	1,162	0.252	0.482	0.000	0.000	0.181
County Stay-at-Home Order$_{i,t}$	1,162	11.10%	31.43%	0.00%	0.00%	0.00%
County Retail Establishment Order$_{i,t}$	1,162	12.56%	33.16%	0.00%	0.00%	0.00%
County Gatherings Order$_{i,t}$	1,162	34.42%	47.53%	0.00%	0.00%	100.00%
County School Order$_{i,t}$	1,162	34.68%	47.62%	0.00%	0.00%	100.00%
Investor Sentiment$_{t}$	12,808	0.000	1.000	-0.984	-0.186	1.011
News Sentiment$_{t}$	12,808	0.000	1.000	-0.633	0.009	0.620

Table 2
Crash Probability Spread

The table displays average *CPS* values across various sample splits: sample period splits, splits across VIX quintiles, daily returns quintiles and past month return quintiles. Robust standard errors are displayed in the parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Sample Period Splits:	2000 - 2020	2000 - 2004	2005 - 2009	2010 - 2014	2015 - 2020
	13.187***	9.650***	11.592***	15.025***	15.933***
	(0.329)	(0.516)	(0.491)	(0.545)	(0.644)
Volatility Ranking Splits:	Low VIX_t	2	Middle VIX_t	4	High VIX_t
	14.748***	14.164***	11.943***	11.951***	13.132***
	(0.186)	(0.176)	(0.177)	(0.178)	(0.213)
Daily Returns Ranking Splits:	Low R^M_t	2	Middle R^M_t	4	High R^M_t
	12.672***	13.181***	13.605***	13.197***	13.281***
	(0.188)	(0.187)	(0.182)	(0.185)	(0.202)
Past Month Returns Ranking Splits:	Low $R^M_{t-30,t-1}$	2	Middle $R^M_{t-30,t-1}$	4	High $R^M_{t-30,t-1}$
	13.488***	13.356***	13.747***	13.160***	12.186***
	(0.181)	(0.184)	(0.196)	(0.185)	(0.196)

Table 3
Fundamental Factors

The table presents the results of regression models where the dependent variables are π^{Survey} and CPS_t. The explanatory variables include π^{Option} at date *t* as well as its first five lags, the *LJV* measure from Bollerslev et al. (2015) at date *t* as well as its first five lags, the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags. The first five lags of the dependent variable are also included. Only the sum of the regression coefficients associated with each variable is displayed. Hansen-Hodrick standard error corrections are used to construct χ^2 statistics, displayed in brackets. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent Variable:	π^{Survey}_t	π^{Survey}_t	π^{Survey}_t	π^{Survey}_t	CPS_t	CPS_t	CPS_t
$\Sigma_{j=0}^{5} \pi^{Option}_{t-j}$	0.081***						
	[32.090]						
$\Sigma_{j=0}^{5} LJV_{t-j}$		0.079***					
		[37.740]					
$\Sigma_{j=0}^{5} R^M_{t-j}$			-24.320***		-7.421		-0.303
			[9.040]		[0.900]		[0.000]
$\Sigma_{j=0}^{5} VIX_{t-j}$				0.027***		-0.003	-0.004
				[33.500]		[0.510]	[0.810]

Table 3 (Cont.)

$\Sigma_{j=1}^{5}\,\pi^{Survey}_{t-j}$	0.894***	0.981***	0.906***	0.895***			
	[13559.770]	[14502.690]	[15964.200]	[14009.950]			
$\Sigma_{j=1}^{5}\,CPS_{t-j}$					2.333***	2.403***	2.421***
					[13277.850]	[13250.880]	[13298.070]
N	6,570	6,570	6,570	6,570	6,570	6,570	6,570
R^2	82.23%	82.28%	82.17%	82.25%	79.10%	79.08%	79.12%

Table 4
Investor Sentiment

The table presents the results of regression models where the dependent variables are π^{Survey} and CPS_t. Average Investor $Sentiment_t$ is the seven-day moving average of the daily average investor sentiment measure constructed from the survey response text. News $Sentiment_t$ is the daily average sentiment measure constructed from newspaper articles. Sentiment $Spread_t$ is the difference between the investor and news sentiment measures. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed include the first five lags of News Sentiment, market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. The first five lags of the dependent variable are also included but not displayed. Hansen-Hodrick standard error are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable:	$\pi^{Survey}{}_t$	$\pi^{Survey}{}_t$	$\pi^{Survey}{}_t$	CPS_t	CPS_t	CPS_t
Average Investor $Sentiment_t$	-0.561***		-0.554***	-0.422***	-0.473***	
	(0.072)		(0.072)	(0.067)	(0.074)	
News $Sentiment_t$		-0.125**	-0.089*	-0.026	-0.040	
		(0.050)	(0.050)	(0.051)	(0.053)	
Sentiment $Spread_t$						-0.209***
						(0.042)
Lagged π^{Survey} Terms	YES	YES	YES	NO	NO	NO
Lagged CPS Terms	NO	NO	NO	YES	YES	YES
Month FEs	YES	YES	YES	YES	YES	YES
Day-of-week FEs	YES	YES	YES	YES	YES	YES
Control Variables	NO	NO	NO	NO	YES	YES
N	4,345	4,345	4,345	4,345	4,345	4,345
R^2	75.15%	74.48%	75.17%	69.59%	69.81%	69.45%

Table 5
Nonlinear Dependence

The table presents the results of quantile regression models where the dependent variable is the crash probability spread, or $CPS_{i,t}$. The quantile specification is indicated in the first row. Investor $Sentiment_{i,t}$ is the investor sentiment measure for investor *i* constructed from the survey response text. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the news sentiment measure as well as its first five lags, the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the date level are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Specification:	Q(10%)	Q(25%)	Q(50%)	Q(75%)	Q(90%)
Dependent Variable:	CPS_t	CPS_t	CPS_t	CPS_t	CPS_t
Investor $Sentiment_t$	-0.326***	-0.758***	-2.888***	-5.749***	-10.151***
	(0.057)	(0.108)	(0.231)	(0.381)	(0.643)
Day-of-week FEs	YES	YES	YES	YES	YES
Month FEs	YES	YES	YES	YES	YES
Control Variables	YES	YES	YES	YES	YES
N	12,276	12,276	12,276	12,276	12,276
R^2	0.47%	1.36%	4.32%	4.23%	3.88%

Table 6
Valence-Arousal Model

The table presents the results of regression models where the dependent variable are $\pi^{Survey}_{i,t}$ and $CPS_{i,t}$. Arousal $Score_t$ is the arousal score. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the news sentiment measure as well as its first five lags, the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the date level are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Dependent Variable:	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$	$CPS_{i,t}$
Investor $Sentiment_{i,t}$		-3.869***	-3.848***	-3.830***
		(0.185)	(0.185)	(0.184)
Arousal $Score_{i,t}$	1.713***	1.203***	1.077***	1.071***
	(0.211)	(0.208)	(0.204)	(0.203)
Investor $Sentiment_{i,t}$ × Arousal $Score_{i,t}$			-0.469**	-0.464**
			(0.197)	(0.195)
Day-of-week FEs	YES	YES	YES	YES
Month FEs	YES	YES	YES	YES
Control Variables	YES	YES	YES	YES
N	12,337	12,337	12,337	12,337
R^2	2.50%	6.32%	6.39%	5.40%

Table 7
Specific Emotion Measures

The table presents the results of regression models where the dependent variable is $CPS_{i,t}$. The emotion measured in Emotion $Score_t$ is indicated in the first row. PC is the first principal component of the emotion indicators. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the news sentiment measure as well as its first five lags, the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the date level are displayed in parentheses. Robust standard errors clustered on the date level are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Emotion Specification:	Anxiety	Fear	Sadness	Disgust	Optimism	Excitement	PC
Dependent Variable:	$CPS_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$
Emotion $Score_{i,t}$	3.558***	3.551***	3.186***	1.420***	-3.283***	-2.190***	4.061***
	(0.200)	(0.201)	(0.211)	(0.263)	(0.163)	(0.166)	(0.195)
Day-of-week FEs	YES	YES	YES	YES	YES	YES	YES
Month FEs	YES	YES	YES	YES	YES	YES	YES
Control Variables	YES	YES	YES	YES	YES	YES	YES
N	12,943	12,943	12,943	12,943	12,943	12,943	12,943
R^2	3.90%	3.89%	3.27%	1.17%	3.48%	1.90%	4.91%

Table 8
Proximity and Attention

The dependent variable in these specifications is *SVI,* or the natural log of one plus the internet search volume for the earthquake (Panel A) and lottery (Panel B) terms. The key explanatory variables are the natural log of one plus the closest distance to the event, a dummy variable associated with whether the closest distance to the event is under 30 miles, and a dummy variable associated with whether the closest distance to the event is between 30 and 100 miles. The control variables included where indicated are the one-week lagged dependent variable, the natural log of one plus the historical frequency of earthquakes (Panel A) or lottery winners (Panel B), and the natural log of one plus the population of the designated media area. Date and designated media area fixed effects are included where indicated. Robust standard errors double clustered on the date and designated media area levels are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Panel A: Earthquakes				
	(1)	(2)	(3)	(4)
Dependent Variable:	$SVI_{j,t}^{Earthquake}$	$SVI_{j,t}^{Earthquake}$	$SVI_{j,t}^{Earthquake}$	$SVI_{j,t}^{Earthquake}$
$\ln(1+\text{Distance})_{j,t}$	-0.047***			
	(0.012)			
$\text{Distance}_{j,t,[0mi,30mi)}$		0.165***	0.203***	0.205***
		(0.069)	(0.055)	(0.024)
$\text{Distance}_{j,t,[30mi,100mi)}$		-0.039	0.022	0.103***
		(0.085)	(0.057)	(0.030)
Date FEs	NO	NO	NO	YES
Region FEs	NO	NO	NO	YES
Control Variables	NO	NO	YES	YES
N	116,522	116,522	115,746	115,746
R^2	0.47%	0.10%	11.45%	35.21%

Table 8 (cont.)

Panel B: Lottery Winners				
	(1)	(2)	(3)	(4)
Dependent Variable:	$SVI_{j,t}^{Lottery}$	$SVI_{j,t}^{Lottery}$	$SVI_{j,t}^{Lottery}$	$SVI_{j,t}^{Lottery}$
$\ln(1+Distance)_{j,t}$	-0.025***			
	(0.006)			
$Distance_{j,t,[0mi,30mi)}$		0.856***	0.428***	0.347***
		(0.093)	(0.106)	(0.080)
$Distance_{j,t,[30mi,100mi)}$		0.203***	0.126***	0.075***
		(0.075)	(0.059)	(0.037)
Date FEs	NO	NO	NO	YES
Region FEs	NO	NO	NO	YES
Control Variables	NO	NO	YES	YES
N	116,522	116,522	115,746	115,746
R^2	0.40%	0.29%	12.12%	22.02%

Table 9
Time Variation in Attention

The dependent variable in these specifications is *SVI,* or the natural log of one plus the internet search volume for the earthquake (Panel A) and lottery (Panel B) terms. The key explanatory variables are the weekly lead and lags for dummy variables associated with whether there was an earthquake (Panel A) or lottery winner (Panel B) within a 30 mile radius. The leads and lags range from two weeks prior to two weeks after the event. The control variables included where indicated are the one-week lagged dependent variable, the natural log of one plus the historical frequency of earthquakes (Panel A) or lottery winners (Panel B), and the natural log of one plus the population of the designated media area. Date and designated media area fixed effects are included where indicated. Robust standard errors double clustered on the date and designated media area levels are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Panel A: Earthquakes

	(1)	(2)	(3)
Dependent Variable:	$SVI_{j,t}^{Earthquake}$	$SVI_{j,t}^{Earthquake}$	$SVI_{j,t}^{Earthquake}$
Nearby Moderate $Earthquake_{j,t-2}$	-0.036	-0.011	0.000
	(0.025)	(0.024)	(0.012)
Nearby Moderate $Earthquake_{j,t-1}$	-0.013	0.001	0.004
	(0.025)	(0.025)	(0.014)
Nearby Moderate $Earthquake_{j,t}$	0.019	-0.018	0.020
	(0.026)	(0.025)	(0.014)
Nearby Moderate $Earthquake_{j,t+1}$	0.183***	0.204***	0.190***
	(0.030)	(0.030)	(0.022)
Nearby Moderate $Earthquake_{j,t+2}$	-0.008	0.022	0.007
	(0.031)	(0.028)	(0.013)
Date FEs	NO	NO	YES
Region FEs	NO	NO	YES
Control Variables	NO	YES	YES
N	116,522	115,746	115,746
R^2	0.10%	11.45%	35.19%

Table 9 (cont.)

Panel B: Lottery Winners			
	(1)	(2)	(3)
Dependent Variable:	$SVI_{j,t}^{Lottery}$	$SVI_{j,t}^{Lottery}$	$SVI_{j,t}^{Lottery}$
Nearby Lottery $Winner_{j,t-2}$	0.481***	0.026	0.117***
	(0.080)	(0.068)	(0.055)
Nearby Lottery $Winner_{j,t-1}$	0.581***	0.082	0.172***
	(0.076)	(0.063)	(0.054)
Nearby Lottery $Winner_{j,t}$	0.979***	0.505***	0.488***
	(0.087)	(0.097)	(0.076)
Nearby Lottery $Winner_{j,t+1}$	0.820***	0.422***	0.347***
	(0.089)	(0.106)	(0.080)
Nearby Lottery $Winner_{j,t+2}$	0.468***	0.076	0.125***
	(0.076)	(0.055)	(0.050)
Date FEs	NO	NO	YES
Region FEs	NO	NO	YES
Control Variables	NO	YES	YES
N	116,522	115,746	115,746
R^2	0.93%	12.21%	22.12%

Table 10
Earthquakes

The table presents the results from regression models where the dependent variables are $\pi^{Survey}_{i,t}$ and $CPS_{i,t}$. $\pi^{Survey}_{i,t}$ is the average subjective 6-month crash probabilities based on survey data as of day *t*. The crash probability spread, or $CPS_{i,t}$, is the difference between $\pi^{Survey}_{i,t}$ and the 6-month crash probabilities based on option prices. The key explanatory variable is a dummy based on whether the survey respondent is within 30 miles of the epicenter of an earthquake above 2.5 magnitude but below 5.5 magnitude within the past month. *Institutional* is a dummy based on whether the survey respondent is an institutional investor. The model also includes: a dummy based on whether the survey respondent is within 30 miles of the epicenter of an earthquake at least 5.5 magnitude within the past month; and the natural log of one plus the number of earthquakes per year from 1900 to 2000. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the ZIP code and date levels are displayed in parentheses. Statistical significance is denoted as *, **, and *** for 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)
Investor Type:	Indiv.	Indiv.	Inst.	Inst.	All
Dependent Variable:	$\pi^{Survey}_{i,t}$	$CPS_{i,t}$	$\pi^{Survey}_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$
Nearby Moderate $Earthquake_{i,t}$	0.026**	0.021**	-0.006	-0.003	0.019**
	(0.012)	(0.009)	(0.012)	(0.009)	(0.008)
Nearby Severe $Earthquake_{i,t}$	0.098	0.080	-0.016	-0.004	0.031
	(0.113)	(0.075)	(0.056)	(0.041)	(0.042)
Historical Earthquake $Frequency_{i,t-30}$	-0.003*	-0.002*	0.000	-0.001	-0.001*
	(0.002)	(0.001)	(0.002)	(0.001)	(0.001)
$Institutional_{i,t}$					-0.019***
					(0.003)
Nearby Moderate $Earthquake_{i,t}$ × $Institutional_{i,t}$					-0.019*
					(0.011)

Table 10 (cont.)

Control Variables	NO	YES	NO	YES	YES
Month FEs	YES	YES	YES	YES	YES
Day-of-week FEs	YES	YES	YES	YES	YES
N	6,212	6,212	5,486	5,486	11,698
R^2	0.48%	1.32%	0.42%	1.61%	1.55%

Table 11
Lottery Winners

The table presents the results from regression models where the dependent variables are $\pi^{Survey}{}_{i,t}$ and $CPS_{i,t}$. $\pi^{Survey}{}_{i,t}$ is the average subjective 6-month crash probabilities based on survey data as of day *t*. The crash probability spread, or $CPS_{i,t}$, is the difference between $\pi^{Survey}{}_{i,t}$ and the 6-month crash probabilities based on option prices. The key explanatory variable is a dummy based on whether the survey respondent is within 30 miles of the store that sold a winning MegaMillions or Powerball ticket within the past month. *Institutional* is a dummy based on whether the survey respondent is an institutional investor. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the ZIP code and date levels are displayed in parentheses. Statistical significance is denoted as *, **, and *** for 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)
Investor Type:	Indiv.	Indiv.	Inst.	Inst.	All
Dependent Variable:	$\pi^{Survey}{}_{i,t}$	$CPS_{i,t}$	$\pi^{Survey}{}_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$
Nearby Lottery Winner$_{i,t}$	-0.030**	-0.025***	0.015	0.007	-0.025**
	(0.012)	(0.010)	(0.019)	(0.015)	(0.010)
Lottery Winner Frequency$_{i,t-30}$	-0.010***	-0.007***	-0.009***	-0.007***	-0.007***
	(0.003)	(0.002)	(0.003)	(0.002)	(0.002)
Institutional$_{i,t}$					-0.022***
					(0.003)
Nearby Lottery Winner$_{i,t}$ × Institutional$_{i,t}$					0.031*
					(0.017)
Control Variables	NO	YES	NO	YES	YES
Month FEs	YES	YES	YES	YES	YES
Day-of-week FEs	YES	YES	YES	YES	YES
N	6,212	6,212	5,486	5,486	11,698
R^2	0.69%	1.47%	0.59%	1.76%	1.70%

Table 12
COVID-19 Pandemic Responses: Individual Investors

The table presents the results from regression models where the dependent variables are $\pi^{Survey}{}_{i,t}$ and $CPS_{i,t}$ for the individual investor subsample from January 2019 through December 2020. $\pi^{Survey}{}_{i,t}$ is the average subjective 6-month crash probabilities based on survey data as of day *t*. The crash probability spread, or $CPS_{i,t}$, is the difference between $\pi^{Survey}{}_{i,t}$ and the 6-month crash probabilities based on option prices. The key explanatory variables are: the one-week percentage change in county-level COVID cases-per-capita for survey respondent *i* as of day *t*; the natural log of cumulative county-level COVID cases-per-capita for survey respondent *i* as of day *t*; the one-week percentage change in county-level COVID deaths-per-capita for survey respondent *i* as of day *t*; the natural log of cumulative county-level COVID deaths-per-capita for survey respondent *i* as of day *t*; dummy variables associated with whether the residence county of survey respondent *i* as of day *t* has an active count-level order associated with stay-at-home, retail establishment restrictions, gatherings, and school closures. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the ZIP code and date levels are displayed in parentheses. Statistical significance is denoted as *, **, and *** for 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Investor Type:	Indiv.	Indiv.	Indiv.	Indiv.	Indiv.	Indiv.	Indiv.
Dependent Variable:	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	$CPS_{i,t}$
Change in County Cases (p.c.)$_{i,t}$	0.300***					0.288***	0.157**
	(0.072)					(0.075)	(0.062)
County Cases (p.c.)$_{i,t-7}$		-0.052				-0.047	-0.027
		(0.053)				(0.052)	(0.037)
Change in County Deaths (p.c.)$_{i,t}$			-13.168			-47.071	-23.625
			(47.735)			(45.094)	(32.972)
County Deaths (p.c.)$_{i,t-7}$				3.366***		4.273***	2.657**
				(1.269)		(1.499)	(1.107)
County Stay-at-Home Order$_{i,t}$					0.068	0.065	0.035
					(0.045)	(0.043)	(0.032)
County Retail Establishment Order$_{i,t}$					0.010	-0.010	-0.014
					(0.050)	(0.049)	(0.037)

Table 12 (cont.)

County Gatherings Order$_{i,t}$					-0.125	-0.140	-0.727**
					(0.146)	(0.145)	(0.305)
County School Order$_{i,t}$					-0.234***	-0.229***	0.512
					(0.041)	(0.041)	(0.332)
Control Variables	NO	NO	NO	NO	NO	NO	YES
Year-Month FEs	YES	YES	YES	YES	YES	YES	YES
State FEs	YES	YES	YES	YES	YES	YES	YES
Day-of-week FEs	YES	YES	YES	YES	YES	YES	YES
N	653	653	653	653	653	653	653
R^2	15.19%	14.62%	14.46%	15.12%	14.96%	16.71%	16.18%

Table 13
COVID-19 Pandemic Responses: Institutional Investors

The table presents the results from regression models where the dependent variables are $\pi^{Survey}{}_{i,t}$ and $CPS_{i,t}$ for the institutional investor subsample from January 2019 through December 2020. $\pi^{Survey}{}_{i,t}$ is the average subjective 6-month crash probabilities based on survey data as of day *t*. The crash probability spread, or $CPS_{i,t}$, is the difference between $\pi^{Survey}{}_{i,t}$ and the 6-month crash probabilities based on option prices. The key explanatory variables are: the one-week percentage change in county-level COVID cases-per-capita for survey respondent *i* as of day *t*; the natural log of cumulative county-level COVID cases-per-capita for survey respondent *i* as of day *t*; the one-week percentage change in county-level COVID deaths-per-capita for survey respondent *i* as of day *t*; the natural log of cumulative county-level COVID deaths-per-capita for survey respondent *i* as of day *t*; dummy variables associated with whether the residence county of survey respondent *i* as of day *t* has an active count-level order associated with stay-at-home, retail establishment restrictions, gatherings, and school closures. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the ZIP code and date levels are displayed in parentheses. Statistical significance is denoted as *, **, and *** for 10%, 5%, and 1% levels, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Investor Type:	Inst.	Inst.	Inst.	Inst.	Inst.	Inst.	Inst.
Dependent Variable:	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	$CPS_{i,t}$
Change in County Cases (p.c.)$_{i,t}$	0.151					0.155	0.133
	(0.100)					(0.105)	(0.083)
County Cases (p.c.)$_{i,t-7}$		-0.045				-0.039	-0.029
		(0.057)				(0.059)	(0.045)
Change in County Deaths (p.c.)$_{i,t}$			-85.577***			-89.701**	-80.355**
			(29.836)			(39.197)	(35.442)
County Deaths (p.c.)$_{i,t-7}$				-2.452		-1.649	-0.852
				(1.653)		(2.392)	(1.881)
County Stay-at-Home Order$_{i,t}$					0.058	0.078	0.054
					(0.088)	(0.088)	(0.057)
County Retail Establishment Order$_{i,t}$					0.040	0.044	0.007
					(0.080)	(0.080)	(0.066)

Table 13 (cont.)

County Gatherings $Order_{i,t}$					0.110	0.091	-0.055
					(0.127)	(0.127)	(0.094)
County School $Order_{i,t}$					-0.408***	-0.431***	-0.466*
					(0.113)	(0.114)	(0.251)
Control Variables	NO	NO	NO	NO	NO	NO	YES
Year-Month FEs	YES	YES	YES	YES	YES	YES	YES
State FEs	YES	YES	YES	YES	YES	YES	YES
Day-of-week FEs	YES	YES	YES	YES	YES	YES	YES
N	497	497	497	497	497	497	497
R^2	20.19%	19.86%	20.10%	19.92%	20.22%	21.47%	20.73%

Appendix Table A1
Correlation Matrix

The table displays correlation matrix of the variables used in the analysis. π^{Survey} is the crash probability responses from the survey. π^{Option} is the natural probability of a stock market crash based on Martin (2017). CPS is the difference between π^{Survey} and π^{Option}. R^M_t is the market return. VIX is the CBOE VIX index. Investor Sentiment is the seven-day moving average of the daily average investor sentiment measure constructed from the survey response text.

		(1)	(2)	(3)	(4)	(5)
(1)	π^{Survey}_t					
(2)	π^{Option}_t	31.55%				
(3)	CPS_t	92.00%	-7.61%			
(4)	R^M_t	-0.83%	-5.91%	1.49%		
(5)	VIX_t	30.86%	87.37%	-3.39%	2.25%	
(6)	Average Investor $Sentiment_t$	-19.54%	-2.44%	-19.41%	-1.80%	3.54%

Appendix Table A2
Sample Splits of Option-Implied Crash Probabilities

The table displays average π^{Option}_t values across various sample splits: sample period splits, splits across VIX quintiles, daily returns quintiles and past month return quintiles. Robust standard errors are displayed in the parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Sample Period Splits:	2000 - 2020	2000 - 2004	2005 - 2009	2010 - 2014	2015 - 2020
	5.845***	6.265***	6.006***	6.043***	5.195***
	(0.130)	(0.192)	(0.397)	(0.205)	(0.197)
Volatility Splits:	Low VIX_t	2	Middle VIX_t	4	High VIX_t
	2.998***	4.082***	5.477***	6.960***	9.707***
	(0.025)	(0.027)	(0.038)	(0.046)	(0.070)
Daily Returns Splits:	Low R^M_t	2	Middle R^M_t	4	High R^M_t
	7.157***	5.292***	4.795***	5.241***	6.739***
	(0.084)	(0.067)	(0.062)	(0.071)	(0.085)
Past Month Returns Splits:	Low $R^M_{t-30,t-1}$	2	Middle $R^M_{t-30,t-1}$	4	High $R^M_{t-30,t-1}$
	7.930***	5.601***	4.778***	4.698***	6.219***
	(0.082)	(0.068)	(0.064)	(0.063)	(0.077)

Appendix Table A3
Sample Splits on the Survey-based Crash Probabilities

The table displays average π^{Survey}_t values across various sample splits: sample period splits, splits across VIX quintiles, daily returns quintiles and past month return quintiles. Robust standard errors are displayed in the parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

Sample Period Splits:	2000 - 2020	2000 - 2004	2005 - 2009	2010 - 2014	2015 - 2020
	19.052***	15.872***	17.652***	21.122***	21.142***
	(0.292)	(0.431)	(0.557)	(0.534)	(0.596)
Volatility Splits:	Low VIX_t	2	Middle VIX_t	4	High VIX_t
	17.779***	18.264***	17.497***	18.863***	22.863***
	(0.195)	(0.182)	(0.180)	(0.178)	(0.225)
Daily Returns Splits:	Low R^M_t	2	Middle R^M_t	4	High R^M_t
	19.828***	18.484***	18.420***	18.483***	20.042***
	(0.210)	(0.194)	(0.185)	(0.193)	(0.213)
Past Month Returns Splits:	Low $R^M_{t-30,t-1}$	2	Middle $R^M_{t-30,t-1}$	4	High $R^M_{t-30,t-1}$
	21.440***	18.982***	18.579***	17.879***	18.382***
	(0.208)	(0.193)	(0.199)	(0.186)	(0.200)

Appendix Table A4
Internal Consistency

The table presents the results of regression models where the dependent variable is $\pi^{Survey}_{i,t}$. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the date level are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Dependent Variable:	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$
Valuations Too High	0.080***							0.051***
	(0.003)							(0.003)
Valuations Too Low		-0.057***						-0.018***
		(0.004)						(0.004)
Buy Recommendation			-0.055***					-0.027***
			(0.003)					(0.002)
Sell Recommendation				0.110***				0.072***
				(0.006)				(0.006)
Speculative Environment					0.028***			0.013***
					(0.003)			(0.002)
Investors Too Optimistic						0.040***		0.021***
						(0.003)		(0.003)
Investors Too Pessimistic							-0.010***	-0.006**
							(0.003)	(0.003)

Appendix Table A4 (cont.)

Day-of-week FEs	YES	YES	YES	YES	YES	YES	YES	YES
Month FEs	YES	YES	YES	YES	YES	YES	YES	YES
Control Variables	YES	YES	YES	YES	YES	YES	YES	YES
N	15,695	15,695	15,695	15,695	15,695	15,695	15,695	15,695
R^2	8.2%	3.4%	4.9%	6.4%	2.7%	3.5%	1.9%	12.3%

Appendix Table A5
Baseline Specification Using Survey Data

The table presents the results of regression models where the dependent variable are $\pi^{Survey}_{i,t}$ and $CPS_{i,t}$. Investor $Sentiment_{i,t}$ is the investor sentiment measure for respondent *i* and date *t*. News $Sentiment_t$ is the daily average sentiment measure constructed from newspaper articles. Sentiment $Spread_{i,t}$ is the difference between the investor and news sentiment measures. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the news sentiment measure as well as its first five lags, the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the date level are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable:	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$
Investor $Sentiment_{i,t}$	-4.170***		-4.162***	-3.857***	-3.963***	
	(0.184)		(0.184)	(0.183)	(0.182)	
News $Sentiment_t$		-0.509***	-0.426**	-0.053	-0.132	
		(0.189)	(0.182)	(0.183)	(0.181)	
Sentiment $Spread_{i,t}$						-1.970***
						(0.140)
Day-of-week FEs	YES	YES	YES	YES	YES	YES
Month FEs	YES	YES	YES	YES	YES	YES
Control Variables	NO	NO	NO	NO	YES	YES
N	12,808	12,808	12,808	12,808	12,808	12,808
R^2	4.82%	0.36%	4.86%	4.19%	4.79%	2.64%

Appendix Table A6
Quantile Regressions on Option-implied Crash Probabilities

The table presents the results of quantile regression models where the dependent variable is π^{Option}_t. The quantile specification is indicated in the first row. Average Investor Sentiment$_{i,t}$ is the investor sentiment measure for investor *i* constructed from the survey response text. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the news sentiment measure as well as its first five lags, the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the date level are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)	(5)
Specification:	Q(10%)	Q(25%)	Q(50%)	Q(75%)	Q(90%)
Dependent Variable:	π^{Option}_t	π^{Option}_t	π^{Option}_t	π^{Option}_t	π^{Option}_t
Average Investor Sentiment$_t$	-0.034**	-0.008	-0.010**	-0.010	-0.028*
	(0.013)	(0.008)	(0.005)	(0.009)	(0.016)
Lagged CPS Terms	YES	YES	YES	YES	YES
Day-of-week FEs	YES	YES	YES	YES	YES
Month FEs	YES	YES	YES	YES	YES
Control Variables	YES	YES	YES	YES	YES
N	6,570	6,570	6,570	6,570	6,570
R^2	92.25%	92.14%	92.00%	92.27%	92.29%

Appendix Table A7
Self versus Other Investor Sentiment

The table presents weights on each variable in the principal component analysis for the first three components. The bottom displays the eigenvalues associated with each component.

Principal Component:	1	2	3
Factor Loadings:			
Anxiety	0.497	0.309	-0.162
Fear	0.496	0.300	-0.165
Sadness	0.410	0.302	-0.069
Disgust	0.113	0.248	0.960
Optimism	-0.435	0.498	-0.093
Excitement	-0.370	0.644	-0.111
Eigenvalue:	3.203	1.124	0.964

Appendix Table A8
Own versus Other Investor Sentiment

The table presents the results of regression models where the dependent variable is the crash probability spread, or $CPS_{i,t}$. Investor $Sentiment_{i,t}$ is the investor sentiment measure for investor *i* constructed from the survey response text. Other Investor $Sentiment_{-i,t}$ is the investor sentiment measure of investors that are not investor *i* on date *t* constructed from the survey response text. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the news sentiment measure as well as its first five lags, the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the date level are displayed in parentheses. Robust standard errors clustered on the date level are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Dependent Variable:	$CPS_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$
Investor $Sentiment_{i,t}$		-3.966***		-3.950***
		(0.182)		(0.182)
Other Investor $Sentiment_{-i,t}$	-0.456**	-0.475***		
	(0.178)	(0.174)		
Other Investor $Sentiment_{-i,t-1}$			-0.469***	-0.401***
			(0.134)	(0.129)
Day-of-week FEs	YES	YES	YES	YES
Month FEs	YES	YES	YES	YES
Control Variables	YES	YES	YES	YES
N	12,808	12,808	12,808	12,808
R^2	0.73%	4.84%	0.78%	4.86%

Appendix Table A9
Alternative Sentiment Measures

The table presents the results of regression models where the dependent variable are $\pi^{Survey}_{i,t}$ and $CPS_{i,t}$. Investor $Sentiment^{LM}_{i,t}$ is the investor sentiment measure based on Loughran and McDonald (2011) for investor *i* constructed from the survey response text. Investor $Sentiment^{GI}_{i,t}$ is the investor sentiment measure based on the General Inquirer for investor *i* constructed from the survey response text. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the news sentiment measure as well as its first five lags, the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the date level are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Dependent Variable:	$\pi^{Survey}_{i,t}$	$\pi^{Survey}_{i,t}$	$CPS_{i,t}$	$CPS_{i,t}$
Investor $Sentiment^{LM}_{i,t}$	-1.968***		-1.958***	
	(0.154)		(0.153)	
Investor $Sentiment^{GI}_{i,t}$		-2.141***		-2.129***
		(0.153)		(0.152)
Day-of-week FEs	YES	YES	YES	YES
Month FEs	YES	YES	YES	YES
Control Variables	YES	YES	YES	YES
N	12,877	12,877	12,877	12,877
R^2	2.76%	1.66%	2.94%	1.85%

Appendix Table A10
Investor Sophistication and Coherency

The table presents the results of regression models where the dependent variable is the crash probability spread, or $CPS_{i,t}$. Investor $Sentiment_{i,t}$ is the investor sentiment measure for investor *i* constructed from the survey response text. Investor $Sophistication_{i,t}$ is the score for investor sophistication based on the survey response text. Writing $Coherency_{i,t}$ is the score for the coherency in writing style of the survey response text. Two-way fixed effects on the month and day-of-week levels are included where indicated, but not reported. Control variables included but displayed are the news sentiment measure as well as its first five lags, the market return at date *t* as well as its first five lags, VIX at date *t* as well as its first five lags, and π^{Option} at date *t* as well as its first five lags. Robust standard errors clustered on the date level are displayed in parentheses. Robust standard errors clustered on the date level are displayed in parentheses. Statistical significance at the 10%, 5%, and 1% levels are denoted by *, **, and ***, respectively.

	(1)	(2)	(3)	(4)
Dependent Variable:	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	$\pi^{Survey}{}_{i,t}$	CPS_t
Investor $Sentiment_{i,t}$			-3.506***	-3.489***
			(0.197)	(0.196)
Investor $Sophistication_{i,t}$	-1.033***		-0.680***	-0.672***
	(0.185)		(0.177)	(0.176)
Writing $Coherency_{i,t}$		-2.572***	-1.192***	-1.184***
		(0.210)	(0.218)	(0.217)
Day-of-week FEs	YES	YES	YES	YES
Month FEs	YES	YES	YES	YES
Control Variables	YES	YES	YES	YES
N	12,600	12,600	12,600	12,600
R^2	1.98%	3.42%	6.22%	5.23%