Predicting CDS Spreads and Stock Returns with Weather Risk: A Study Utilizing NLP/LLM and AI Measures*

Yi Zhou†

April 20, 2024

Abstract

Drawing from a comprehensive and unique dataset encompassing both quantitative and qualitative weather risk measures, the study finds that both numerical and textual representations of weather risk can predict future credit risk, expected stock returns, and firm fundamentals. To explore the textual dimension of weather risk, this paper utilizes advanced natural language processing (NLP) techniques, including Term Frequency-Inverse Document Frequency (TF-IDF), Word2Vec, and leverages Large Language Model (LLM) such as BERT (Bidirectional Encoder Representations from Transformers). To conduct the empirical analysis, this study utilizes Artificial Intelligence (AI) using TensorFlow/Keras, Deep Learning (DL), and Machine Learning (ML).

JEL classification: G02, G11, G12, G13, G14.
Keywords: Weather risk, climate risk, credit risk, credit default swaps (CDS), stock returns, artificial intelligence (AI), deep learning (DL), machine learning (ML), large language model (LLM), Bidirectional Encoder Representations from Transformers (BERT), natural language processing (NLP), Term Frequency-Inverse Document Frequency (TF-IDF), Word to Vector (Word2Vec), and TensorFlow/Keras (TF/Keras).

*I deeply appreciate Kelvin Kay Droegemeier for tremendous encouragement and support. I also thank Mark-it for providing the CDS data. All remaining errors are mine.

†Department of Finance, College of Business, San Francisco State University, 1600 Holloway Avenue, San Francisco, CA 94132, USA, Email: yizhou88@sfsu.edu, Office: 415-338-2661, Fax: 415-338-0596.

Predicting CDS Spreads and Stock Returns with Weather Risk Measures
A Study Utilizing NLP/LLM and AI

Abstract

Drawing from a comprehensive and unique dataset encompassing both quantitative and qualitative weather risk measures, the study finds that both numerical and textual representations of weather risk can predict future credit risk, expected stock returns, and firm fundamentals. To explore the textual dimension of weather risk, this paper utilizes advanced natural language processing (NLP) techniques, including Term Frequency-Inverse Document Frequency (TF-IDF), Word2Vec, and leverages Large Language Model (LLM) such as BERT (Bidirectional Encoder Representations from Transformers). To conduct the empirical analysis, this study utilizes Artificial Intelligence (AI) using TensorFlow/Keras, Deep Learning (DL), and Machine Learning (ML).

JEL classification: G02, G11, G12, G13, G14.

Keywords: Weather risk, climate risk, credit risk, credit default swaps (CDS), stock returns, artificial intelligence (AI), deep learning (DL), machine learning (ML), large language model (LLM), Bidirectional Encoder Representations from Transformers (BERT), natural language processing (NLP), Term Frequency-Inverse Document Frequency (TF-IDF), Word to Vector (Word2Vec), and TensorFlow/Keras (TF/Keras).

1 Introduction

Modern financial literature has given great importance to climate risk. There has been a plethora of papers on how climate risk affects equities, bond, and credit markets[1] While this literature has primarily focused on climate risk, there is a related but distinct area that has received less attention. This paper contributes to the literature in three ways. First, it is the first study to examine how weather risk affects credit and equity markets, offering a new perspective in this emerging area. Although climate risk and weather risk are often conflated, they differ in their scope, magnitude, and predictability[2] This article investigates the consequences of weather risk on corporate credit risk, expected stock returns, and company fundamentals, utilizing a unique dataset comprising both quantitative and qualitative weather risk measures. The findings highlight the significant predictive power of numerical and textual weather risk indicators for future credit risk, expected stock returns, and business fundamentals.

Weather risk refers to short-term, localized weather events and their possible effects on certain industries such as agriculture, transportation, and construction. It includes events like storms, hurricanes, and flooding. In contrast, climate risk refers to the long-term effects of climate change on ecosystems, economies, civilizations, and infrastructure over decadal to centennial timelines. It entails changes in temperature patterns, precipitation levels, and sea levels. Furthermore, weather events are usually ephemeral, lasting hours, days, or weeks, while some calamities, such as extended droughts or floods, can have long-term impacts. Climate change, on the other hand, is a continuing, long-term trend caused by variables such as greenhouse gas emissions, deforestation, and industrial activity, with long-term and probably permanent consequences for ecosystems and weather patterns. To summarize, climate and weather risk differ in their timelines, pre-

[1]Please see, for example, Sautner et al. (2023a), Sautner et al. (2023b), Ilhan et al. (2023), Ilhan et al. (2021), Painter (2020), Pankratz and Schiller (2024), Li et al. (2024), Addoum et al. (2023), Ginglinger and Moreau (2023), Bartrama et al. (2022), and Engle et al. (2020), etc.

[2]Please see, for example, Keller and DeVecchio (2019), Arnell (2017), Muir-Wood (2016), Kunreuther and Michel-Kerjan (2011), and Posner (2005)).

dictability, and repercussions, with climate change causing long-term shifts in environmental conditions and weather patterns, while weather disasters represent short-term hazards requiring fast response and readiness.

The second contribution of this paper to the literature is that this paper utilizes a unique and novel publicly available weather event database. This database encompasses weather risk measures presented in both quantitative and qualitative formats. The weather risk measures are derived from the "Storm Events Database" provided by the National Weather Service (NWS), which is part of the National Oceanic and Atmospheric Administration (NOAA). This database covers weather events from January 1997 to December 2023 across the United States. It contains detailed information about each event, including the start date, location (state and county), event type (e.g., thunderstorm wind, winter storm, tornado), episode and event identification numbers, magnitude or intensity measures, estimated property and crop damages in dollars, casualties (direct and indirect deaths and injuries), and very importantly and textual narratives describing the overall episode and specific event details. This dataset contains both textual narratives and numeric measurements that are scientific in nature, pertaining specifically to weather events. The integration of qualitative textual descriptions alongside quantitative numeric data related to meteorological phenomena provides a unique opportunity to analyze scientific data from multiple perspectives.

The third major contribution of this study is its relevance to the rapidly evolving fields of large language models (LLMs), artificial intelligence (AI), and natural language processing (NLP), which have garnered significant attention recently. To effectively analyze the textual representations of weather risk present in the database, this paper employs a range of cutting-edge NLP techniques. These include traditional methods such as Term Frequency-Inverse Document Frequency (TF-IDF) and more advanced approaches like Word2Vec for generating dense word embeddings. Moreover, the study leverages the power of state-of-the-art LLMs, specifically the Bidirectional Encoder Representations from Transformers (BERT) model, to capture intricate semantic and contextual informa-

tion from the textual narratives.

By combining these textual features extracted through NLP with the numerical measurements from the weather event database, the paper demonstrates the potential for integrating unstructured text and structured data to enhance predictive modeling tasks. Specifically, it explores the ability of this multimodal data to forecast future credit risk and expected stock returns. To streamline the analysis and modeling process, the paper leverages the latest advancements in AI and deep learning (DL) frameworks. It employs TensorFlow/Keras, a widely-used open-source platform for building and deploying DL models, as well as traditional machine learning (ML) algorithms.

The paper has three major findings. The first finding is that both numeric and textual weather risk measures significantly predict future conditions of credit/default risk. This paper shows that weather risk measures like magnitude and property/crop damage significantly and positively predict credit default swap (CDS) levels, changes in CDS levels, percentage changes in CDS levels, and CDS slope across multiple maturity terms. A 1% increase in magnitude corresponds to around a 40% increase in the 5-year CDS average level, while a 1% rise in property damage leads to a 17% increase in the 5-year CDS average level. Similar positive impacts are seen for changes and percentage changes in CDS levels from magnitude, property damage, and crop damage measures. The CDS slope, which reflects the term structure, also increases significantly with higher magnitude and property damage. These evidence show that when weather-related risks surge–whether it's due to more intense events, greater property damage, or increased casualties–the levels and changes in CDS increase, too. This reflects investors' anxieties about the potential impact of weather risks on borrowers' financial stability and their ability to fulfill debt obligations.

Notably, this paper reveals that weather events have a disproportionately greater impact on the creditworthiness of non-investment grade (speculative) firms compared to investment grade firms, with the former experiencing an impact over three times greater. This is likely due to the weaker financial positions and higher debt levels of non-investment

grade firms, making them more vulnerable to weather-related disruptions and increased costs. With limited access to capital markets and higher borrowing costs, these firms may face exacerbated liquidity constraints and increased funding costs, ultimately affecting their creditworthiness to a greater extent than their investment grade counterparts.

The second finding is that both numeric and textual weather risk measures significantly predict expected stock returns. This paper shows that weather risk measures like magnitude and crop damage significantly and negatively impact future stock returns. A 1% increase in magnitude corresponds to a 14% decrease in next month's stock average return, while a 1% rise in crop damage leads to a 30% reduction in the following month's return. Portfolio analyses further reinforce this negative relation between weather risk proxies and expected equity returns. Stocks in the highest magnitude quartile underperform those in the lowest quartile by 0.59% per month on a raw return basis and 0.47-0.57% per month on a risk-adjusted basis across different factor models, with statistically significant alphas. This return spread is even more pronounced among smaller stocks and low-priced firms, with a hedge portfolio going long the lowest magnitude tercile and short the highest magnitude tercile generating significant positive returns of 0.80-0.92% per month in these subsamples. Overall, the analyses provide robust evidence that elevated weather risk negatively impacts future stock performance, especially for smaller, more constrained firms potentially more vulnerable to such risks.

The third finding is that both numeric and textual weather risk measures significantly predict firm fundamentals. This paper shows that higher weather risk, proxied by measures like crop damage and storm magnitude, negatively impacts firm leverage, profitability, and growth opportunities while increasing capital expenditures in the following period. Greater crop damage is associated with higher future leverage ratios and changes in leverage. Higher storm magnitude predicts lower return on assets, sales growth, earnings growth, and Tobin's Q in the next period, suggesting weather events disrupt operations and reduce expected profitability and valuations. However, elevated storm magnitude also forecasts an increase in next period's capital expenditures.

The paper is organized as follows. Section 2 reviews recent literature related to this paper. Section 3 describes the data source and the summary statistics for the key variables. Section 4 presents empirical evidence and robustness checks. Section 5 concludes.

2 Literature Review

This paper is closely intertwined with the extensive body of literature on climate and finance on credit markets, equity returns, and firm fundamentals, while also distinguishing between the distinct yet interconnected concepts of climate and weather. Sautner et al. (2023a) develop a method that identifies firms' climate change exposures from earnings call transcripts using machine learning, capturing opportunity, physical, and regulatory shocks, and show that the measures predict outcomes related to the net-zero transition like green job creation and patenting, and are priced in financial markets. Sautner et al. (2023b) estimate the risk premium for S&P 500 stocks' climate change exposure from 2005 to 2020, finding an overall insignificant unconditional risk premium but noting positive trends pre-financial crisis and post-2014. Ilhan et al. (2023) demonstrate through a survey and empirical evidence that institutional investors highly value and actively seek climate risk disclosures. Ilhan et al. (2021) highlight the necessity for robust regulatory measures to address climate change, revealing how climate policy uncertainty is reflected in option market pricing, with higher costs for downside tail risk protection observed for firms with more carbon-intensive operations. Using a firm-level measure of temperature sensitivity, Cuculiza et al. (2024) find that firms more sensitive to temperature changes have lower future profitability, riskier policies, and lower subsequent returns, suggesting mispricing that nonlocal investors and analysts contribute to, enabling a trading strategy exploiting this mispricing to generate over 4% annual risk-adjusted returns from 1968-2020. Using detailed geographic data on U.S. households' exposure to future sea level rise (SLR), Ilhan (2022) show that households more exposed to long-run SLR risks hold less stock market participation compared to unexposed neighbors. Kölbel et al. (2024) use

BERT to assess regulatory climate risk disclosures' effects on CDS, finding that disclosing transition risks tends to increase CDS spreads post-2015 Paris Climate Agreement, while disclosing physical risks decreases them.

Pankratz and Schiller (2024) investigate how physical climate events like heat and floods at supplier locations negatively impact the operating income of suppliers and their customers, leading customers to terminate supplier relationships when suppliers' realized exposure exceeds expectations, and how customers learn from experience and adapt by replacing suppliers with ones having lower expected and realized climate exposure. Li et al. (2024) develop text-based measures to quantify firms' exposure to physical and transition climate risks from earnings call transcripts, and finds that firms with high transition risk exposure, especially non-responsive ones, have been discounted by investors as climate concerns grow; it also documents how firms respond differently through investments, innovation, and employment policies when facing varying levels of climate risk exposure. Addoum et al. (2023) examine the impact of extreme temperatures on corporate profitability across different industries by combining temperature data with locations of public companies' establishments, finding that extreme temperatures significantly affect earnings in over 40% of industries in a bidirectional manner. Ginglinger and Moreau (2023) use firm-level data on forward-looking physical climate risk to examine its impact on capital structure, finding that greater physical climate risk leads to lower leverage in the post-2015 period after increased disclosure standards, with the reduction in leverage driven by both firms' lower optimal leverage and lenders increasing spreads for riskier firms, consistent with the hypothesis that physical climate risk affects leverage through higher expected distress and operating costs. Bartrama et al. (2022) show that localized climate risk mitigation policies like California's cap-and-trade program can have unintended consequences due to regulatory arbitrage, as financially constrained firms shift emissions and output from California to underutilized plants in other states, leading to an overall increase in total emissions among constrained firms and undermining the policy's effectiveness. Engle et al. (2020) propose and implements a procedure to dynam-

ically hedge climate change risk by extracting innovations from constructed climate news series through textual analysis, using a mimicking portfolio approach with firms' ESG scores to model climate risk exposures to build hedge portfolios that effectively hedge innovations in climate news both in-sample and out-of-sample.

Gounopoulos and Zhang (2024) find companies increase cash reserves in response to rising climate risks driven by heightened environmental enforcement and physical risks, especially for financially constrained firms with low environmental awareness who rely more on equity issuance and cost cuts than debt to bolster cash holdings precautionarily. Lin et al. (2023) show that production inflexibility coupled with product price uncertainty creates price risk, which significantly impacts firms' liquidity management, as evidenced by the finding that higher electricity price volatility leads to increased cash holdings among firms using inflexible production technologies in the deregulated electricity industry, with the effect most pronounced for financially constrained firms and those lacking hedging opportunities, suggesting capital market and balance sheet liquidity are substitutes. Pankratz et al. (2022) link firm performance, analyst forecasts, and earnings announcement returns to firm-specific heat exposure measures , finding that increased exposure to extremely high temperatures reduces firms' revenues and operating income, with analysts and investors failing to fully anticipate the economic repercussions of heat as a physical climate risk..

Additionally, the paper intricately connects with the body of literature concerning CDS, such as Kölbel et al. (2024), Zhang et al. (2009) and Ericsson et al. (2009). Zhang et al. (2009) propose a novel method to explain the factors influencing credit default swap (CDS) premiums. This approach involves analyzing high-frequency equity price data to identify the volatility and jump risks associated with individual firms. The empirical findings demonstrate that volatility risk alone can account for 48% of the variation observed in CDS spread levels. Additionally, jump risk independently forecasts 19% of the variation. Furthermore, after incorporating credit ratings, macroeconomic conditions, and firms' balance sheet information into the analysis, the model can explain 73% of the

total variation in CDS premiums. Ericsson et al. (2009) find that a minimal set of theoretical determinants of default risk, particularly volatility and leverage, have substantial explanatory power in predicting default swap spreads through univariate and multivariate regressions, and a principal component analysis indicates limited evidence for a residual common factor, suggesting that the theoretical variables explain a significant amount of variation in the data.

This paper contributes to the fast growing literature on using BERT family models to process contextual information for predicting financial variables. Chen et al. (2023) employ advanced large language models to extract contextualized representations of news text for predicting returns, surpassing traditional word-based methods like bag-of-words or word vectors. By capturing both syntax and semantics, these representations offer a more comprehensive understanding of text meaning. Lin Tan and Zhang (2024) discuss the use of LLMs to generate article-level representations of Chinese financial news articles. These representations are then employed for two key modeling tasks: sentiment classification to assess the overall positive or negative tone of the articles, and to predict daily cross-sectional stock returns. Kim and Nikolaev (2023) utilize the combination of textual and numeric information in Management Discussion and Analysis (MD&A) to examine the impact of narrative context on financial disclosures, particularly in MD&A sections. Leveraging machine learning techniques like BERT and ANN, it quantifies how narrative communication affects the interpretation of financial data. Results show that contextual information notably enhances predicting future earnings changes and stock returns, indicating its growing significance. Kim and Nikolaev (2024) emphasize the significance of contextual information mandated by the Securities and Exchange Commission (SEC) in interpreting reported numbers, particularly concerning operating profitability. Through the use of a bidirectional recurrent neural network with Long Short-Term Memory (LSTM) cells, the model integrates contextual insights to enhance the understanding of reported profitability.

This paper contributes to the recent increase in research on the application of deep

learning methods in AI in finance. Balachandran et al. (2024) examine how market reactions to the tone of voluntary disclosures in 8-K filings, assessed using a new algorithm focusing on adjectives and adverbs, reveal a positive association with market response, affecting stock liquidity and bid-ask spreads, indicating firms strategically use tone to communicate qualitative information, with the algorithm surpassing traditional methods by considering linguistic nuances. Glasserman et al. (2023) construct a measure of news novelty or unusualness called entropy, using a recurrent neural network applied to a large news corpus. An increase in news entropy predicts negative stock market returns and negative macroeconomic outcomes over the next year. Cao et al. (2023) analyzes corporate executive presentations to assess the impact of visual information on market reactions, finding that short-term abnormal returns correlate positively with forward-looking operational data extracted from presentation slides using deep learning.

3 Data

The weather risk proxies are obtained from the "Storm Events Database"[3] made available by the National Weather Service (NWS) of the National Oceanic and Atmospheric Administration (NOAA). This database contains data from January 1997 to December 2023. The database contains the following variables: *begin_yearmonth* refers to the year and month when the weather event started. *year* denotes the specific year when the weather event commenced. *month_name* indicates the name of the month corresponding to the weather event record, such as January, April, July, or September. *state* identifies the state where the weather event took place. *cz_name* specifies the county name where the weather event occurred. *event_type* specifies the type of weather event, such as Thunderstorm Wind, Winter Storm, or Strong Wind. *episode_id* refers to the identification number assigned by the NWS to denote a storm episode. Episodes encompass significant weather phenomena with the potential for loss of life, injuries, property damage, and disruption

[3]https://www.ncei.noaa.gov/pub/data/swdi/stormevents/csvfiles/

to commerce. *event_id* is the identification number assigned by the NWS to identify a single component contributing to a specific storm episode. A weather event is a distinct, short-term meteorological occurrence, while an episode is a prolonged period of consistent weather conditions that may include multiple events. *Magnitude* measures the scale of the event, primarily used for wind speeds (in knots). For each county and each year-month, the maximum value of *Magnitude* is selected. *Damage_Property* quantifies the estimated property damage in dollars caused by the weather event. *Damage_Crops* quantifies the estimated damage to crops in dollars caused by the weather event. *Deaths_Indirect* specifies the number of deaths indirectly associated with the weather event. *death* is the summation of *Deaths_Direct* and *Deaths_Indirect*. *Injuries_Direct* indicates the number of injuries directly resulting from the weather event. *Injuries_Indirect* represents the number of injuries indirectly related to the weather event. *Injury* is the summation of *Injuries_Direct* and *Injuries_Indirect*. *Deaths_Direct* denotes the number of deaths directly attributed to the weather event. For each county and each year-month, the variables *Damage_Property*, *Damage_Crops*, *Death*, and *Injury* are aggregated by summing their values. *episode_narrative* provides a comprehensive overview of the episode's general nature and overall activity, as reported by the NWS. *event_narrative* offers descriptive details of the individual event, as reported by the NWS. The *event_type* covered in this sample are: thunderstorm wind, high wind, marine thunderstorm wind, strong wind, marine high wind, marine strong wind, tornado, flash flood, dust devil, heavy snow, winter storm, blizzard, seiche, wildfire, heavy rain, waterspout, flood, hurricane (typhoon), dust storm.

To give an example: in March 2023, in Simpson County, Mississippi, there was a thunderstorm wind event with a magnitude of 52[4]. The *episode_id* is 180848. The *event_id* is 1106736. The *damage_property* is 10.00K. The *event_narrative* is *Multiple trees were downed near Pinola near MS-28*. The *episode_narrative* is

[4]Please see https://www.weather.gov/bmx/event_03242023.

In the afternoon and evening of the 24th, clusters of strong to severe storms were stretched from southwest to northeast across southeastern Arkansas as robust instability and very strong wind shear served to produce a volatile severe weather environment. To the east in Mississippi, storm activity proved to be more isolated as deep convection struggled develop across the state. With much of the area not being convectively overturned, an area of thunderstorm activity near the Mississippi River near Vicksburg, MS began to become stronger and took advantage of the open warm sector environment, well east of the competitive storm environment to the west that had thus far impeded tornado development. This area of thunderstorm activity quickly consolidated into an organized supercell and produced a family of long track, strong to violent tornadoes from the Mississippi Delta, across central Mississippi, and finally into portions of northeast Mississippi. This cyclical tornado-producing supercell was responsible for the vast majority of severe weather reports from the day and would claim 22 lives across the state. Dozens more were injured, and hundreds of homes were damaged or destroyed by these tornadoes. Having wrought widespread destruction across the area, this tornado event by most any measure represents a historic degree of devastation for the state of Mississippi and the region at large.

Figure 1 and Figure 2 are graphical depiction of textual information of *episode_narrative* and *event_narrative*, in which the magnitude of each word signifies its frequency or significance within the provided text. In textual layout, less frequent words are typically rendered in smaller font sizes, whereas more frequently occurring words are presented in larger font sizes. Figure 1 highlights the overall weather conditions and main activities during the episode, including terms like ”severe thunderstorm,” ”shower thunderstorm,” and ”heavy rain,” among others. Figure 2 pays direct attention to specific descriptive elements of the individual event, such as ”large tree,” ”tree blow,” and ’tree limb,’ etc.

One of the goals of this paper is to predict credit risk using weather risk measures. The credit risk is proxied by the month-end credit default swap spread provided by Mark-it. A credit default swap (CDS) is a financial transaction in which the seller agrees to reimburse the buyer in the case of debt default or other credit-related catastrophes. In effect, the seller of the CDS insures the customer against the default of a certain asset. The buyer of the CDS pays the seller a periodic fee, known as the CDS ”fee” or ”spread,” in exchange for potential compensation if the asset fails. In the event of a default, the CDS buyer receives compensation, typically the loan’s face value, while the seller gains custody of the

failed loan or its equivalent market value in cash. The CDS spread indicates the market's assessment of the credit risk associated with an individual firm. A bigger gap indicates that investors perceive a higher chance of default, whereas a narrower spread implies a lesser perceived risk. Investors use CDS spreads to assess credit risk and market opinion about a borrower's creditworthiness. Widening spreads may signal deteriorating credit quality or increasing market uncertainty, whilst narrowing spreads may reflect improved credit conditions or a reduction in perceived risk. CDS maturity ranges from one to ten years, with the five-year CDS being the most commonly traded. The combined dataset, merging the "Storm Events Database" with Mark-it's CDS dataset, spans from January 2001 to October 2019. The sample size is 38,927.

The other goal of this paper is predict stock return using using weather risk measures. The Center for Research in Security Prices (CRSP) monthly returns dataset comprises data on monthly stock returns, stock prices, trading volumes, and shares outstanding for a broad spectrum of companies listed on major U.S. stock exchanges, encompassing stocks listed on NYSE, AMEX, and Nasdaq exchanges. Firm-specific control variables utilized in the study are obtained from firm quarterly balance-sheet and annual accounting data provided by Compustat Industrial Quarterly and Annual files (COMPUSTAT) from from Wharton Research Data Services (WRDS).

To examine the impact of weather risk on future stock returns, we merge the "Storm Events Database" with CRSP/COMPUSTAT from WRDS, creating additional variables for empirical analysis. $SIZE$ is defined as the natural logarithm of the market value (MKV), where MKV is the product of the stock price (PRC) and the number of publicly held shares ($SHROUT$), recorded in thousands. TO represents the turnover and is calculated by dividing the sum of trading volumes (VOL), expressed in hundred shares for monthly data, by the product of the shares outstanding ($SHROUT$) and 1,000. MOM refers to the cumulative return over the last six months. *Leverage* is defined as the leverage measure, calculated as the ratio of the book value of total liability to the sum of the book value of total liability and the market value of equity. ROA represents the return on

assets, defined as net income scaled by total assets. $SALE$ is natural logarithm of sales. P/E is the price-earnings ratio. $TobinsQ$ is defined as the market capitalization of common stock plus the liquidation value of preferred shares plus the book value of long-term debt divided by total assets. $CAPEX$ is defined as capital expenditures scaled by sales. The combined dataset, merging the "Storm Events Database" with CRSP/COMPUSTAT, covers the period from January 1997 to December 2023. The sample size is 674,201.

Panels A and B of Table 1 provide the mean and standard deviation of the aforementioned variables, while Panel C displays the correlation matrix between $Log(Magnitude)$, $Log(Damage\_Property)$, $Log(Damage\_Crops)$, $Death$, and $Injury$. Notably, a high correlation is observed between $Log(Magnitude)$ and $Log(Damage\_Property)$, whereas the correlation coefficients for the other weather risk measure variables are low, all below 0.25.

4 Empirical analysis

4.1 Numeric weather risk measure predicts credit risk, expected stock returns and firm fundamentals

In Table 2–Table 11, the standard errors are clustered on the county levels and year-month levels to control for the cross-sectional and time-series autocorrelations, respectively, following Petersen (2009). The climate risk is captured by the county variables. With this empirical setting, the climate risk is effectively controlled. All regression control for the lagged dependent variables to address unobserved firm heterogeneity.

Table 2 demonstrates that both $Magnitude$ and $Damage\_Property$ significantly and positively predict CDS levels across all maturity terms. In Model 1, CDS_{t+1}^{5-Year} regressed on $Log(Magnitude)_t$ yields a coefficient of 0.568, statistically significant at the 5% level. This implies that a 1% increase in $Magnitude_t$ corresponds to a 56.8 basis points (bp) increase in CDS_{t+1}^{5-Year}. Given that the mean of CDS^{5-Year} is 143.73, a 1% increase in $Magnitude_t$ raises CDS_{t+1}^{5-Year} by 40% of its mean. Similar results are observed in Models 3 to 6 for CDS_{t+1}^{3-Year}, CDS_{t+1}^{7-Year}, $CDS_{t+1}^{10-Year}$, and $CDS_{t+1}^{20-Year}$, respectively. In Model

2, regressing CDS_{t+1}^{5-Year} on $Log(Damage_Property)t$ yields a coefficient of 0.245, statistically significant at the 5% level. This implies that a 1% increase in *Damage_Property* leads to a 24.5 basis points (bp) increase in CDS_{t+1}^{5-Year}. Consequently, a 1% increase in *Damage_Property* raises CDS_{t+1}^{5-Year} by 17% of its mean. The empirical results show that as weather risk intensifies, CDS levels escalate, reflecting the heightened concerns of bond investors over the potential impact of climatic uncertainties on the ability of borrowers to meet their obligations.

Table 3 illustrates that *Magnitude*, *Damage_Property*, and *Damage_Crops* significantly and positively predict changes in *CDS* across all maturity terms. In Models 1 and 2, a 1% increase in *Damage_Crops* is associated with an increase of 25.1 bp and 25.5 bp in $Chg\ CDS_{t+1}^{5-Year}$, respectively. Similarly, Models 3 and 4 indicate that a 1% increase in *Magnitude* results in increases of 20.8 bp and 24.4 bp in $Chg\ CDS_{t+1}^{3-Year}$ and $Chg\ CDS_{t+1}^{7-Year}$, respectively. As weather risk escalates, CDS changes exhibit a more pronounced trajectory, mirroring the financial market's amplified concerns over the potential reverberations of weather risk on borrowers' ability to service their debt obligations.

Table 4 demonstrates that *Magnitude*, *Damage_Property*, and *Damage_Crops* significantly and positively predict percentage changes in *CDS* across all maturity terms. The coefficients for *Magnitude*, *Damage_Property*, and *Damage_Crops* are almost all statistically significant at the 1% level. These results further support the findings presented in Table 3 and Table 4. As weather risk intensifies, CDS percentage changes become more pronounced, reflecting the financial market's heightened sensitivity to the potential impacts of weather risk on borrowers' creditworthiness.

Table 5 demonstrates that both *Magnitude* and *Damage_Property* significantly and positively predict *Slope* across all maturity terms. In Model 1, $Slope_{t+1}$ regressed on $Log(Magnitude)_t$ yields a coefficient of 0.432, statistically significant at the 5% level. This implies that a 1% increase in $Magnitude_t$ corresponds to a 43.2 basis points (bp) increase in $Slope_{t+1}$. Given that the mean of *Slope* is 69.56, a 1% increase in $Magnitude_t$ raises $Slope_{t+1}$ by 62% of its mean. In Model 2, regressing $Slope_{t+1}$ on $Log(Damage\_Property)t$

yields a coefficient of 0.187, statistically significant at the 5% level. This implies that a 1% increase in *Damage_Property* leads to a 18.7 basis points (bp) increase in $Slope_{t+1}$. Consequently, a 1% increase in *Damage_Property* raises $Slope_{t+1}$ by 27% of its mean. These results underscore the financial market's perception that heightened weather risk amplifies the credit risk profile of borrowers, thereby increasing the premium demanded by investors for longer-dated credit protection.

The empirical findings unveil a clear pattern: as weather risk escalates, be it through intensifying event magnitudes, rising property damages, increasing deaths or injuries, CDS levels, changes, and percentage changes exhibit a pronounced upward trajectory, mirroring investors' apprehensions over the potential reverberations of weather risk on borrowers' creditworthiness and ability to service debt obligations. This underscores the market's perception that heightened weather risk amplifies the overall credit risk profile of borrowers, thereby commanding a steeper premium for longer-dated credit protection.

Models 1 and 2 of Table 6 reveal intriguing results, indicating that the weather event has a significantly greater impact on the creditworthiness of non-investment grade firms compared to investment grade firms, with the former experiencing an impact more than three times greater than the latter. Non-investment grade firms, also known as speculative firms, typically have weaker financial positions and higher levels of debt compared to investment grade firms. As a result, they may have less financial resilience to withstand the adverse effects of weather events, such as disruptions to operations or increased costs for repairs. Non-investment grade firms may also have limited access to capital markets or face higher borrowing costs compared to investment grade firms. Consequently, the impact of adverse weather events on their financial performance may exacerbate liquidity constraints or increase the cost of raising additional funds, further impacting their creditworthiness.

Table 7 demonstrates that both *Magnitude* and *Damage_Crops* significantly and negatively predict RET. In Model 1, where RET_{t+1} is regressed on $Log(Magnitude)_t$, the coefficient is -0.153, which is statistically significant at the 5% level. This indicates that a

1% increase in $Magnitude_t$ corresponds to a decrease of 0.00153 in RET_{t+1}. Given that the mean of RET is 0.011, a 1% increase in $Magnitude_t$ decreases RET_{t+1} by 14% of its mean. In Models 2 and 3, we observe the even stronger effect of $Magnitude_t$ on RET_{t+1}. In Model 1, where RET_{t+1} is regressed on $Damage\_Crops_t$, the coefficient is -0.332, which is statistically significant at the 1% level. This indicates that a 1% increase in $Damage\_Crops_t$ corresponds to a decrease of 0.00332 in RET_{t+1}. Given that the mean of RET is 0.011, a 1% increase in $Damage\_Crops_t$ decreases RET_{t+1} by 30% of its mean.

Table 8 displays the average raw returns of the quartile portfolios that are equally weighted based on the *Magnitude* variable. It also shows the differences in average raw returns between the portfolios in the bottom and top quartiles. Additionally, the table presents the alphas of the portfolios in relation to three different models: the capital asset pricing model (CAPM) (Fama and French (1992)), the Fama-French three-factor model (which includes market, size, and book-to-market factors) (Fama and French (1993)), and the Carhart four-factor model (which includes market, size, book-to-market factors, and momentum) (Carhart (1997)).

In Table 8, both Panels A and B categorize stocks into four quartiles based on their *Magnitude* from the previous month, and this categorization is done every month. Once equities have been allocated to portfolios, they are retained for a duration of one month. The monthly portfolio return is computed by taking the average returns of all the stocks in the portfolio, using both equal-weighted (Panel A) and value-weighted (Panel B) methods. We create quartile portfolios based on the variable *Magnitude* and adjust them on a monthly basis. Portfolio 1 (Low) consists of equities with the lowest *Magnitude* value in the previous month, whereas Portfolio 4 (High) consists of stocks with the highest *Magnitude* value in the previous month. In our approach, we assign equal weight to stocks in Panel A and value weight to stocks in Panel B for each quartile portfolio. Additionally, we rebalance these portfolios on a monthly basis.

In Table 8, Panel A demonstrates that the average raw return of stocks in the bottom quartile with the lowest *Magnitude* is 0.96% per month, which monotonically falls

to 0.37% per month for stocks in the top quartile. The average difference in raw returns between the bottom and top quartiles is 0.59% per month (7.08% per year), with a significant Newey-West t-statistic of 1.96. The variations in returns between quartiles 1 and 4 are fairly comparable whether we risk-adjust using the CAPM, at 0.57% per month ($t-statistic = 1.85$), the Fama-French three-factor model, at 0.55% per month ($t-statistic = 1.75$), and the Carhart four-factor model, at 0.47% per month ($t-statistic = 1.68$). The three alphas are statistically significant at the 10% level.

Panel B of Table 8 shows that the average raw return of stocks in the bottom quartile with the lowest *Magnitude* is 1.46% per month, and this number decreases to 0.47% per month for stocks in the top quartile. The average difference in raw returns between the bottom and top quartiles is 0.99% per month (11.88% per year), with a very significant Newey-West t-statistic of 2.21. The risk-adjusted returns between quartiles 1 and 4 are quite comparable if we use the CAPM, at 0.87% per month ($t-statistic = 2.01$), the Fama-French three-factor model, at 0.88% per month ($t-statistic = 1.97$), and the Carhart four-factor model, at 0.75% per month ($t-statistic = 1.83$). The first two alphas show statistical significance at the 5% level.

Panel C of Table 8 demonstrates that the negative link between *Magnitude* and expected stock returns is particularly pronounced for small and low-priced firms. Panel C of Table 8 shows the return of an equal-weighted portfolio that is long the lowest tercile of stocks and short the top tercile ranked by *Magnitude*, in subsamples of stocks sorted by proxies of arbitrage–size and stock price level. The entry in the first cell of the first column of Panel C corresponding to the *Magnitude* signal indicates that among firms in lower market capitalization, a strategy that is long on firms in the lowest *Magnitude* tercile and short on firms in the highest *Magnitude* tercile produces a monthly average raw return of 0.92% with the statistical significance at the 1% level. In the subsample with lower market capitalization, the average raw returns of the portfolio strategy that buys low *Magnitude* stocks and shorts high *Magnitude*, as well as the alphas with respect to the CAPM, the Fama-French three-factor model, or the Carhart four-factor model, are all positive and

significant at the 1% level. The fifth cell of the sixth column of Panel C corresponding to the *Magnitude* signal indicates that among firms in the lower price, a strategy that is long on firms in the lowest *Magnitude* tercile and short on firms in the highest *Magnitude* tercile produces a monthly average raw return of 0.80% with the statistical significance at the 5% level. In the subsample with the lower price, the average raw returns of the portfolio strategy that buys low *Magnitude* stocks and shorts high *Magnitude*, as well as the alphas with respect to the CAPM, the Fama-French three-factor model, or the Carhart four-factor model, are all positive and significant at the 5% level.

The above empirical results uncover a significant inverse relationship between weather risk measures and expected equity returns. Weather risk proxies like event magnitude and crop damage exhibit pronounced predictive power, exerting a substantial negative influence on future stock performance. A 1% increase in magnitude foreshadows a considerable 14% decline in average stock returns the following month, while a 1% increase in crop damage forecasts a significant 30% decrease in returns. over the next month. Portfolio-level examinations reinforce this negative association, with stocks in the highest magnitude quartile significantly underperforming those in the lowest quartile by 0.59% per month on a raw basis and 0.47-0.57% per month on a risk-adjusted basis across factor models. This return disparity is amplified among smaller, lower-priced firms, potentially reflecting heightened vulnerability to weather risks. A hedge strategy going long the lowest magnitude tercile and short the highest generates significant positive returns of 0.80-0.92% monthly for these subsamples. Collectively, the empirical analyses unveil robust evidence that elevated weather risk adversely impacts subsequent equity performance, with the effect being particularly acute for smaller, financially constrained firms plausibly less resilient to such weather hazards.

Table 9 demonstrates that *Damage_Crops* significantly and negatively predict *Leverage*. In Model 1, where $Leverage_{t+1}$ is regressed on $Damage\_Crops_t$, the coefficient is 0.001, which is statistically significant at the 5% level. In Model 2, where $Chg\ Leverage_{t+1}$ is regressed on $Damage\_Crops_t$, the coefficient is 0.001, which is statistically significant at

the 5% level. In Model 3, where $PChg\ Leverage_{t+1}$ is regressed on $Damage\_Crops_t$, the coefficient is 0.002, which is statistically significant at the 5% level. Thus, higher values of *Damage_Crops* are associated with higher values of *Leverage*.

Table 10 demonstrates that *Magnitude* significantly and negatively predict future profitability. In Model 1, where ROA_{t+1} is regressed on $Log(Magnitude)_t$, the coefficient is -0.0003, which is statistically significant at the 5% level. This indicates that a 1% increase in $Magnitude_t$ corresponds to a decrease of 0.0003% in ROA_{t+1}. Given that the mean of *ROA* is 0.0006, a 1% increase in $Magnitude_t$ decreases ROA_{t+1} by 0.5% of its mean. In Models 2 to 5, we observe the negative effect of $Magnitude_t$ on $Chg\ ROA_{t+1}$, $Chg\ SALE_{t+1}$, $PChg\ SALE_{t+1}$, and $PChg\ P/E_{t+1}$, which are significant at either the 5% or 1% levels. These results indicate the disruption in operations due to the weather event leads to lower future profitability.

Table 11 demonstrates that *Magnitude* significantly and negatively predict future *TobinsQ*. In Models 1 and 3, where $TobinsQ_{t+1}$ is regressed on $Log(Magnitude)_t$, the coefficient is -0.004, which is statistically significant at the 1% level. Table 11 also demonstrates that *Magnitude* significantly and positively predict future $CAPEX$. In Models 4 and 6, where $CAPEX_{t+1}$ is regressed on $Log(Magnitude)_t$, the coefficient is 0.001, which is statistically significant at the 1% level.

A lower *TobinsQ* generally indicates that the market perceives the growth prospects of the firm to be less favorable or that its assets are overvalued relative to their replacement cost. This could be due to the negative impact of the severe weather event on the firm's operations, revenues, and future cash flows. The results indicate that while higher storm magnitudes may reduce *TobinsQ* by decreasing the market value of assets due to the disruption in operations, they can also lead to higher $CAPEX$ in the next period. Companies may invest more in capital expenditures in the next period to repair or replace damaged assets, infrastructure, or facilities resulting from the severe weather event.

The empirical analyses highlight the strong effects of weather risk on various aspects of corporate financial strategies and performance. Following crop damage, there's a no-

table positive correlation with increased leverage, indicating firms may turn to heightened debt financing post-events. Simultaneously, larger event magnitudes negatively impact profitability metrics such as return on assets, sales growth, and price-to-earnings multiples, signaling disruptive consequences for operational performance. Furthermore, while higher storm intensities reduce Tobin's Q, reflecting diminished growth prospects, they also prompt increased capital expenditures in subsequent periods. This apparent contradiction may signify market concerns about growth limitations alongside strategic investments by firms to address damages, mitigate risks, or capitalize on opportunities post-event. Collectively, these findings accentuate the intricate interplay between weather risk factors and corporate financial decision-making, profitability, valuation, and investment dynamics.

4.2 Textural weather risk measure predicts credit risk, expected stock returns and firm fundamentals

The "Storm Events Database" offers two innovative textual measures for assessing weather risk: *episode_narrative* and *event_narrative*. To retrieve the information contained in *episode_narrative* and *event_narrative*, we employ feature extraction techniques using the following natural language processing (NLP) methods to transform textual data into numerical representations suitable for machine learning algorithms: TF-IDF, Word2Vec, and BERT.

TF-IDF, also known as Term Frequency-Inverse Document Frequency, is a quantitative measure utilized to assess the importance of terms within a document. Term Frequency (TF) quantifies the frequency of a term's occurrence in a document, whereas Inverse Document Frequency (IDF) evaluates the scarcity of a word throughout the complete collection of documents. The TF score is calculated by tallying the number of times a term appears in the document, with greater frequency leading to higher scores. In contrast, the IDF score is calculated by dividing the total number of documents by the number of documents containing the term, and then applying a logarithmic transformation. As a result,

less common terms in the collection of texts have higher IDF values. The TF-IDF score for a phrase in a document is determined by the product of its TF (phrase Frequency) and IDF (Inverse Document Frequency) values. This approach modifies the weight of each term according to its significance and infrequency, thus enabling a precise depiction of the document's content.

Lopez-Lira and Tang (2023) use TF-IDF scores to quantify the qualitative information in ChatGPT's explanations and analyze the words with the highest TF-IDF scores in the positive, negative, and neutral explanations to identify patterns and indicators of financial performance, uncertainty, and other relevant factors. Imerman et al. (2023) use TF-IDF for the bigram frequency of the climate change disclosure in the earnings call transcript.

Word2Vec is a collection of models that are utilized to generate word embeddings. Word embeddings are numerical representations of words that capture both the meaning and structure of those words, derived from their context in a collection of texts. The core concept underlying Word2Vec is that words that occur in comparable settings exhibit comparable meanings. Word2Vec utilizes a neural network to train itself in predicting words by considering the words that surround them. This enables Word2Vec to acquire compact vector representations that capture semantic similarities. Word2Vec consists of two primary model architectures which we use in my analysis–the Continuous Bag-of-Words (CBOW) model which aims to forecast the present word by considering the context words surrounding it and the Skip-gram which attempts to forecast the neighboring context words within a window, based on the current word. Word2Vec embeddings are commonly utilized as inputs for deep learning models in NLP tasks such as text classification.

BERT (Bidirectional Encoder Representations from Transformers) developed by researchers at Google AI in 2018 revolutionized NLP with its bidirectional training of Transformer. BERT is a pre-trained deep learning model that achieves excellent performance on many natural language processing tasks. BERT employs a bidirectional training technique, unlike earlier language models that process text input in a sequential manner

(either left-to-right or right-to-left). This implies that it acquires knowledge about the connections between words in a given context by examining text in both forward and backward directions at the same time. The system utilizes the Transformer encoder architecture, which exploits multi-head self-attention to calculate representations of input sequences. BERT's bidirectional architecture and pre-training on extensive datasets enabled it to acquire sophisticated language comprehension abilities. It initiated a fundamental change in the field of NLP by emphasizing the use of transfer learning from extensive pre-trained language models.

Chen et al. (2023) uses BERT as one of the benchmark models for its empirical analysis. Specifically, the paper adopts BERT as an initial benchmark model for comparison with LLMs. They show that LLMs like BERT achieve better results in sentiment analysis tasks compared to traditional word-based methods. Lin Tan and Zhang (2024) employ BERT to derive contextualized depictions of Chinese news text and forecast stock returns within the Chinese equity market. Kirtac and Germano (2024) use BERT to measure the sentiment of the news of the previous 3 days to forecast daily stock returns. Kim and Nikolaev (2024) employ BERT and artificial neural networks and develops a context-based proxy for future operating profitability, refining traditional measures and emphasizing the importance of qualitative insights in financial analyses. Kölbel et al. (2024) use BERT to classify transition and physical climate risk from 10-K reports.

After performing TF-IDF, Word2Vec, or BERT for feature extraction, we merge the textual features (*episode_narrative* or *event_narrative*) with the numeric features–*magnitude*, *damage_crops*, *damage_property*, *deaths*, and *injuries*. Additionally, we incorporate dummy variables for state, county, and year-month to account for climate-related risks. we aggregate all of these features into a unified feature set for analysis. The length of the *episode_narrative* or *event_narrative* is winsorized at the 10th percentile level.

To conduct the empirical analysis, we employ four supervised machine learning (ML) models: Linear Regression (LR), Decision Tree (DT) Regression, Extreme Gradient Boosting (XGBoost) Regression, and Random Forest (RF) Regression. Additionally, we utilize

three artificial intelligence (AI) and deep learning (DL) models: Deep Neural Networks (DNN), Convolutional Neural Networks (CNN), and Recurrent Neural Networks (RNN).

Linear Regression (LR) is a statistical technique employed to establish a mathematical model that describes the association between a dependent variable, also known as the target variable, and one or more independent variables, referred to as features. This is achieved by fitting a linear equation to the obtained data. Regression projects often employ this method when the target variable is of a continuous nature.

Decision Tree (DT) Regression is a supervised learning technique that works in a non-parametric manner to do regression tasks. The prediction of a target variable is achieved through the acquisition of basic decision rules derived from the properties of the data. Decision trees partition the dataset into subsets by selecting the attribute with the highest significance at each node, hence constructing a hierarchical tree structure.

Extreme Gradient Boosting (XGBoost) Regression is an ensemble learning methodology that constructs a robust predictive model by employing a set of weak learners, specifically decision trees. The process involves the successive addition of trees to the ensemble, wherein each subsequent tree is trained to rectify the faults made by its preceding tree. XGBoost has gained recognition for its notable strengths in terms of efficiency, scalability, and exceptional performance across several machine learning contests.

Random Forest (RF) Regression is also an ensemble learning technique that involves the construction of several decision trees during the training process. The resulting models are then used to perform regression tasks by calculating the average prediction of the individual trees. By aggregating many trees, it enhances the performance of decision trees by mitigating overfitting and improving prediction accuracy.

Deep Neural Networks (DNNs), Convolutional Neural Networks (CNNs), and Recurrent Neural Networks (RNNs) are three types of Artificial Neural Network (ANN) which is a broad family or umbrella term that encompasses various types of neural network architectures, including DNNs, CNNs, and RNNs, among others, such Autoencoders, Generative Adversarial Networks (GANs) and many more.

DNNs refers to a type of ANN composed of multiple layers of interconnected nodes (neurons) arranged in a feedforward manner. Each layer in a DNN typically performs a transformation on the input data, gradually extracting higher-level features as information propagates through the network. DNNs are characterized by their depth, meaning they have multiple hidden layers between the input and output layers. These hidden layers allow DNNs to learn complex patterns and representations from the data. By learning hierarchical representations of features, DNNs can capture intricate relationships and dependencies in the input data, making them powerful models for various machine learning tasks.

In Kim and Nikolaev (2024), textual vectors are utilized as input for training a deep neural network to determine adjustments for profitability context. The model consists of multiple layers, including an input layer with 768 neurons, three hidden layers with decreasing neuron sizes, and an output layer producing structural parameters. The training objective is to minimize the root mean squared error (RMSE) between predicted and actual operating profitability values, with ReLU activation functions applied to all neurons except those in the output layer.

This paper uses a DNN constructed using TensorFlow's Keras API for regression analysis, comprising three hidden layers. These layers apply weights to concatenated text and numeric features, followed by Rectified Linear Unit (ReLU) activation functions, facilitating nonlinear transformations crucial for capturing complex data patterns. With varied neuron sizes and ReLU activation functions, the architecture is designed to progressively capture intricate relationships in the data through multiple layers. The output layer, consisting of a single neuron, utilizes a linear activation function, aligning with the regression objective to estimate continuous outcomes. The model is optimized using the Adam optimizer and mean squared error (MSE) as the loss function, leveraging adaptive learning rates derived from first and second moment estimates of gradients. This optimization approach combines the strengths of AdaGrad and RMSProp, ensuring efficient parameter updates for enhanced model performance.

CNNs is a type of ANNs designed to process structured grid-like data, such as images. CNNs are particularly well-suited for tasks related to computer vision, such as image classification, object detection, and image segmentation. The key feature of CNNs is their ability to automatically learn hierarchical representations of features from raw pixel data. This is achieved through the use of convolutional layers, which apply filters (also known as kernels or feature detectors) to small, overlapping regions of the input image. These filters detect specific patterns or features, such as edges, textures, or shapes, at different spatial locations within the image. CNNs typically consist of multiple layers, including convolutional layers, pooling layers, and fully connected layers. Convolutional layers perform the feature extraction process, pooling layers reduce the spatial dimensions of the feature maps, and fully connected layers perform the final classification or regression task.

Cao et al. (2023) investigates using convolutional neural networks (CNN) to analyze visual data from corporate presentations, categorizing images into Operations Forward, Operations Summary, and Others. Employing transfer learning, CNN models are refined to improve prediction accuracy and minimize training data requirements. Results indicate that visual information, especially forward-looking operational data identified by CNNs, positively affects short-term stock returns, particularly when stocks are predominantly held by AI-equipped institutional investors, highlighting the impact of AI on market responses to visual corporate data.

The CNN neural network architecture in this paper, mirroring the structure of the DNN regression model, integrates three hidden layers with input layers for text and numeric features followed by dense layers for prediction. Text and numeric features are combined and processed through convolutional layers tailored for sequential input, facilitating feature extraction pertinent to tasks like sentiment analysis or text classification. The convolutional layer applies one-dimensional convolutions with 128 filters and a kernel size of 3, utilizing ReLU activation to capture complex patterns. Global max-pooling extracts key features, while dropout regularization prevents overfitting by ran-

domly dropping 10% of neurons. TensorFlow's Keras API is employed for model development, including steps to clear sessions, establish random seeds, and compile the model with MSE loss, followed by training and evaluation utilizing validation data. The CNN model's performance closely resembles that of the DNN.

RNNs is a type of ANNs designed to process sequential data by maintaining internal memory or state. Unlike feedforward neural networks, which process data in a single direction (from input to output), RNNs have connections that form directed cycles, allowing them to capture temporal dependencies in sequential data. The key feature of RNNs is their ability to handle input sequences of varying lengths and to maintain memory of past inputs while processing current inputs. This makes them well-suited for tasks involving sequential data, such as time series prediction. In an RNN, each time step in the input sequence is processed one at a time, with the network's hidden state being updated at each step based on both the current input and the previous hidden state. This recurrent connection enables the network to retain information about past inputs, allowing it to capture long-term dependencies in the data. However, traditional RNNs suffer from the vanishing gradient problem, which makes it difficult for them to learn long-range dependencies in sequences. To address this issue, several variants of RNNs have been developed, including Long Short-Term Memory (LSTM) networks used in this paper, which incorporate mechanisms to better control the flow of information through the network and mitigate the vanishing gradient problem.

Glasserman et al. (2023) use RNN model the probability distributions of words in the articles, allowing calculation of entropy scores that reflect how unusual the text is compared to recent prior months. The process involves training a RNN on articles from the preceding 6 months from $t-6$ to $t-1$ to establish a model, which is then utilized to compute the average entropy score of articles in month t.

The RNN architecture in this paper adopts a bidirectional LSTM design to capture sequential dependencies, with concatenated features fed into dense layers with ReLU activation functions, culminating in a linear activation function in the output layer. Compris-

ing three hidden layers akin to the DNN model, the architecture integrates input layers for both text and numeric features, followed by dense layers for prediction. Each layer applies weights to the input data and ReLU activation, facilitating nonlinear changes to learn complex patterns. Utilizing a bidirectional RNN with LSTM cells, the model incorporates a 64-unit LSTM layer and wraps it with a bidirectional wrapper to process input sequences forward and backward, aiding in capturing information from both past and future contexts. Training the RNN involves employing the Adam optimizer with mean squared error as the loss function. Developed using TensorFlow's Keras API, the RNN regression model undergoes a structured process involving clearing sessions, setting random seeds, and purging log directories, with unique logs generated for each run. The RNN model's performance closely resembles that of the DNN.

Below we will present the empirical results of four ML models (LR, DT, XGBosst, and RF) and three AI/DL models (DNN, CNN, and RNN), after utilizing techniques such as TF-IDF, Word2Vec, or BERT for feature extraction.

We merge the textual features (*episode_narrative* or *event_narrative*) with the numeric features*magnitude*, *damage_crops*, *damage_property*, *deaths*, and *injuries*. Additionally, we incorporate dummy variables for state, county, and year-month to account for climate-related risks. We aggregate all of these features into a unified feature set for analysis. Likewise, Kim and Nikolaev (2024) investigates the crucial role of contextual narrative alongside numerical data in empirical asset pricing models, highlighting the impact of qualitative factors often overlooked in traditional analyses.

For Table 12 to Table 16, Panel A utilizes the textual feature *episode_narrative*, while Panel B employs *event_narrative*. Both panels utilize the following numeric features: *magnitude*, *damage_crops*, *damage_property*, *deaths*, and *injuries*. Additionally, dummy variables for state, county, and year-month are included to account for climate-related risks, along with the value of the target variable of the current month to control for all the firm-specific characteristics. The target variables are defined as follows: for Table 12, it is the 5-year CDS spread of the next month; for Table 13, it is the CDS slope, defined as the

difference between the 5-year CDS spread and the 1-year CDS spread of the next month; for Table 14, it is the next month-end stock price; for Table 15, it is the next month leverage as defined in Table 9; and for Table 16, it is the *TobinsQ* of the next month as defined in Table 11. The target variables are all winsorized at 5% level.

For Table 12 to Table 16, Models 1 to 4 and Models 5 to 8 represent four supervised machine learning (ML): Linear Regression (LR), Decision Tree Regression (DT), Random Forest Regression (RF), and Extreme Gradient Boosting (XGB). Models 1 to 4 utilize TF-IDF for textual feature extraction, while Models 5 to 8 utilize Word2Vec. Similarly, Models 9 to 11 and Models 12 to 14 represent three artificial intelligence (AI) and deep learning (DL) models: Deep Neural Networks (DNN), Convolutional Neural Networks (CNN), and Recurrent Neural Networks (RNN). Models 9 to 11 utilize TF-IDF for textual feature extraction, whereas Models 10 to 14 utilize BERT.

Using the input features, we first train the machine learning (ML) models and the artificial intelligence (AI) and deep learning (DL) models. We fit each model to the training data and then use them to predict the target variable on the test data. We evaluate each model using the statsmodels library, fitting a linear regression model to the actual and predicted target values. Finally, we obtain the model summaries, providing detailed information about each model's performance, including coefficients, standard errors, p-values, and goodness-of-fit statistics for the predicted target values.

ML and AI/DL models approach data analysis with a predictive perspective. They aim to uncover patterns and relationships within data to facilitate predictions or decisions without heavily relying on specific assumptions about the underlying data generation process. ML and DL models exhibit greater flexibility in terms of assumptions, enabling them to manage nonlinear relationships, high-dimensional data, and intricate patterns without strict adherence to particular data distribution assumptions. These models learn patterns directly from data, allowing them to handle both structured and unstructured datasets. ML and AI/DL models are frequently referred to as 'black boxes' because of their complex architectures and vast parameter space. Unlike traditional panel regres-

sion analysis, these models do not provide coefficients, t-stats for the regressors, or other interpretable measures of relationship between variables. Instead, their evaluation primarily relies on predictive accuracy[5]. For Table 12 to Table 16, we report the regression results of the target values on the predicted values of the test set, including the coefficients, standard errors, p-values, and the adjusted R^2s.

Table 12 illustrates that the ensemble combination of the textual weather risk measures, *episode_narrative* and *event_narrative*, and the numeric weather risk measures, *magnitude*, *damage_crops*, *damage_property*, *deaths*, and *injuries*, significantly predict future 5-year CDS levels, which serve as proxies for a firm's credit risk. The statistical significance of the predicted target values is at the 1% level. The average adjusted R^2 values are 93% for Panel A and 92% for Panel B. Consequently, when controlling for climate risk and lagged credit risk measures, these weather risk measures effectively forecast the future creditworthiness of the firm.

Table 13 confirms the results observed in Table 12, demonstrating a similar predictive relationship with alternative credit risk measures, specifically CDS slope. The textual weather risk measures, *episode_narrative* and *event_narrative* and the numeric weather risk measures, *magnitude*, *damage_crops*, *damage_property*, *deaths*, and *injuries*, significantly predict future CDS slopes. The statistical significance of the predicted target values is at the 1% level. The average adjusted R^2 values are 89% for both Panels A and B.

In Table 14, the label variable represents the stock price at time $t+1$. The feature set comprises textual weather risk measures, including *episode_narrative* and *event_narrative*, numeric weather risk measures such as *magnitude*, *damage_crops*, *damage_property*, *deaths*, and *injuries*, and lagged stock prices at times t, $t-1$, and $t-2$. This is on the monthly frequency, as compared to Chen et al. (2023), Lopez-Lira and Tang (2023) and Kim and Nikolaev (2024) which all predict stock returns on the daily frequency. The average adjusted R^2 values are 96% for both Panels A and B. The statistical significance of the pre-

[5]We thank Gregory Ryslik at COMPASS Pathways for insightful discussion about how to evaluate ML and AI/DL models.

dicted target values is at the 1% level. This indicates that when controlling for climate risk and lagged stock prices, these weather risk measures effectively forecast stock prices.

In Table 15 and Table 16, weather risk measures significantly forecast future firm fundamentals, including leverage (Table 15) and *TobinsQ* (Table 16), respectively. This forecasting remains significant after controlling for climate risk measures and lagged firm fundamentals. The predicted target values demonstrate statistical significance at the 1% level. The average adjusted R^2 values exceed 93% for both Panels A and B across both tables.

The empirical analyses leveraging machine learning and AI techniques and textual data underscore the significant predictive power of weather risk measures for an array of corporate outcomes. Tables 12 and 13 demonstrate that an ensemble of textual narratives and numerical indicators pertaining to weather events exhibits remarkable forecasting ability for firms' future credit risk proxies, including 5-year CDS levels and CDS slopes, even after accounting for lagged credit risk factors. This predictive prowess extends to equity markets, as evidenced by Table 14, where weather risk measures effectively anticipate subsequent stock price movements when combined with historical pricing data. Furthermore, Tables 15 and 16 reveal the capacity of these weather risk metrics to presage future firm fundamentals, accurately forecasting leverage ratios and Tobin's Q, a gauge of growth prospects. Notably, these predictive relationships persist after controlling for climate risk factors and past firm characteristics. The exceptionally high adjusted R-squared values, exceeding 90% across all analyses, coupled with the robust statistical significance of the predicted targets, accentuate the remarkable explanatory power of weather risk indicators. These findings underscore the pressing need for market participants and corporate decision-makers to integrate weather risk assessments into their analytical frameworks and strategic planning processes.

5 Conclusion

This research makes a substantial contribution to the body of knowledge regarding the effects of weather risk, as opposed to long-term climate risk, on corporate fundamentals and financial markets. Through the utilization of an extensive NWS database containing qualitative narratives and quantitative measurements of weather phenomena such as floods and cyclones, the analyses combine state-of-the-art natural language processing methods with conventional structured data.

Elevated weather risk is predictive of increased credit risk and wider spreads, particularly for speculative-grade firms with limited financial flexibility, according to the findings. Furthermore, this augurs diminished anticipated stock returns due to the rational pricing by equity investors of these transient operational disruptions and cash flow disturbances. In response to the negative impact of weather hazards on short-term profitability, growth prospects, and valuations, companies increase their capital expenditures the following period, most likely for investments in long-term positioning, resilience, and recovery.

This research demonstrates the potential of modern AI/machine learning frameworks, such as deep learning and large language models, to enable multimodal analyses that integrate numerical and textual data in order to extract novel predictive signals from scientific datasets. Qualitative textual indicators and quantitative weather risk measures both accurately forecast future conditions pertaining to corporate fundamentals, credit risk, and equity returns.

References

Addoum, Jawad M., David T. Ng, and Ariel Ortiz-Bobea (2023), "Temperature Shocks and Industry Earnings News," *Journal of Financial Economics*, vol. 150, 1–45.

Arnell, Nigel (2017), *A Short Guide to Climate Change Risk (Short Guides to Business Risk)*, Routledge.

Balachandran, Sudhakar, Jalaj Pathak, Konduru Sivaramakrishnan, and Rustam Zufarov (2024), "Is Tone in Voluntary Disclosure Credible?" Working Paper.

Bartrama, Söhnke M., Kewei Hou, and Sehoon Kim (2022), "Real Effects of Climate Policy: Financial Constraints and Spillovers," *Journal of Financial Economics*, vol. 143, 668–696.

Cao, Sean, Yichen Cheng, Meng Wang, Yusen Xia, and Baozhong Yang (2023), "Visual Information in the Age of AI: Evidence from Corporate Executive Presentations," Working Paper.

Carhart, Mark (1997), "On Persistence in Mutual Fund Performance," *The Journal of Finance*, vol. 52, 57–82.

Chen, Yifei, Bryan Kelly, and Dacheng Xiu (2023), "Expected Returns and Large Language Models," Working Paper.

Cuculiza, Carina, Alok Kumar, Wei Xin, and Chendi Zhang (2024), "Temperature Sensitivity, Mispricing, and Predictable Returns," Working Paper.

Engle, Robert F., Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel (2020), "Hedging Climate Change News," *Review of Financial Studies*, vol. 33, 1184–1216.

Ericsson, Jan, Kris Jacobs, and Rodolfo Oviedo (2009), "The Determinants of Credit Default Swap Premia," *Journal of Financial and Quantitative Analysis*, vol. 44, 109–132.

Fama, Eugene F. and Kenneth R. French (1992), "The Cross-Section of Expected Stock Returns," *The Journal of Finance*, vol. 47, 427–465.

Fama, Eugene F. and Kenneth R. French (1993), "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, vol. 33, 3–56.

Ginglinger, Edith and Quentin Moreau (2023), "Climate Risk and Capital Structure," *Management Science*, vol. 69, 7492–7516.

Glasserman, Paul, Harry Mamaysky, and Jimmy Qin (2023), "New News is Bad News," Working Paper.

Gounopoulos, Dimitrios and Yu Zhang (2024), "Temperature Trend and Corporate Cash Holdings," *Finaical Management*, pages 1–24.

Ilhan, Emirhan (2022), "Sea Level Rise and Portfolio Choice," Working Paper.

Ilhan, Emirhan, Philipp Krueger, Zacharias Sautner, and Laura T. Starks (2023), "Climate Risk Disclosure and Institutional Investors," *Review of Financial Studies*, vol. 36, 2617–2650.

Ilhan, Emirhan, Zacharias Sautner, and Grigory Vilkov (2021), "Carbon Tail Risk," *Review of Financial Studies*, vol. 34, 15401571.

Imerman, Michael B., Xiaoxia Ye, and Ran Zhao (2023), "Voluntary Disclosures and Climate Change Uncertainty: Evidence from CDS Premiums," Working Paper.

Keller, Edward A. and Duane E. DeVecchio (2019), *Natural Hazards: Earth's Processes as Hazards, Disasters and Catastrophes*, Routledge.

Kim, Alex G. and Valeri Nikolaev (2023), "Context-Based Interpretation of Financial Information," Working Paper.

Kim, Alex G. and Valeri Nikolaev (2024), "Profitability Context and the Cross-Section of Stock Returns," Working Paper.

Kirtac, Kemal and Guido Germano (2024), "Sentiment Trading with Large Language Models," *Finance Research Letters*, vol. 62.

Kölbel, Julian F., Markus Leippold, Jordy Rillaerts, and Qian Wang (2024), "Ask BERT: How Regulatory Disclosure of Transition and Physical Climate Risks Affects the CDS Term Structure," *Journal of Financial Econometrics*, vol. 22, 30–69.

Kunreuther, Howard C. and Erwann O. Michel-Kerjan (2011), *At War With the Weather: Managing Large-scale Risks in a New Era of Catastrophes*, MIT Press.

Li, Qing, Hongyu Shan, Yuehua Tang, and Vincent Yao (2024), "Corporate Climate Risk: Measurements and Responses," vol. 00, 1–53.

Lin, Chen, Thomas Schmid, and Michael S. Weisbach (2023), "Price Risk, Production Flexibility, and Liquidity Management: Evidence from Electricity Generating Firms," Working Paper.

Lin Tan, Huihang Wu and Xiaoyan Zhang (2024), "Large Language Models and Return Prediction in China," Working Paper.

Lopez-Lira, Alejandro and Yuehua Tang (2023), "Can ChatGPT Forecast Stock Price Movements? Return Predictability and Large Language Models," Working Paper.

Muir-Wood, Robert (2016), *The Cure for Catastrophe: How We Can Stop Manufacturing Natural Disasters*, Basic Books.

Painter, Marcus (2020), "An Inconvenient Cost: The Effects of Climate Change on Municipal Bonds," *Journal of Financial Economics*, vol. 135, 468–482.

Pankratz, Nora, Rob Bauer, and Jeroen Derwall (2022), "Climate Change, Firm Performance, and Investor Surprises," Working Paper.

Pankratz, Nora M. C. and Christoph M. Schiller (2024), "Climate Change and Adaptation in Global Supply-Chain Networks," *Review of Financial Studies*, vol. 00, 149.

Petersen, Mitchell A. (2009), "Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches," *The Review of Financial Studies*, vol. 22, 435–480.

Posner, Richard A. (2005), *Catastrophe: Risk and Response*, Oxford University Press.

Sautner, Zacharias, Laurence Van Lent, Grigory Vilkov, and Ruishen Zhang (2023a), "Firm-Level Climate Change Exposure," *Journal of Finance*, vol. 78, 1449–1498.

Sautner, Zacharias, Laurence Van Lent, Grigory Vilkov, and Ruishen Zhang (2023b), "Pricing Climate Change Exposure," *Management Science*, vol. 69, 75407561.

Zhang, Benjamin Yibin, Hao Zhou, and Haibin Zhu (2009), "Explaining Credit Default Swap Spreads with the Equity Volatility and Jump Risks of Individual Firms," *The Review of Financial Studies*, vol. 22, 5099–5131.

Figure 1: The Word Cloud Graph for Episode Narrative

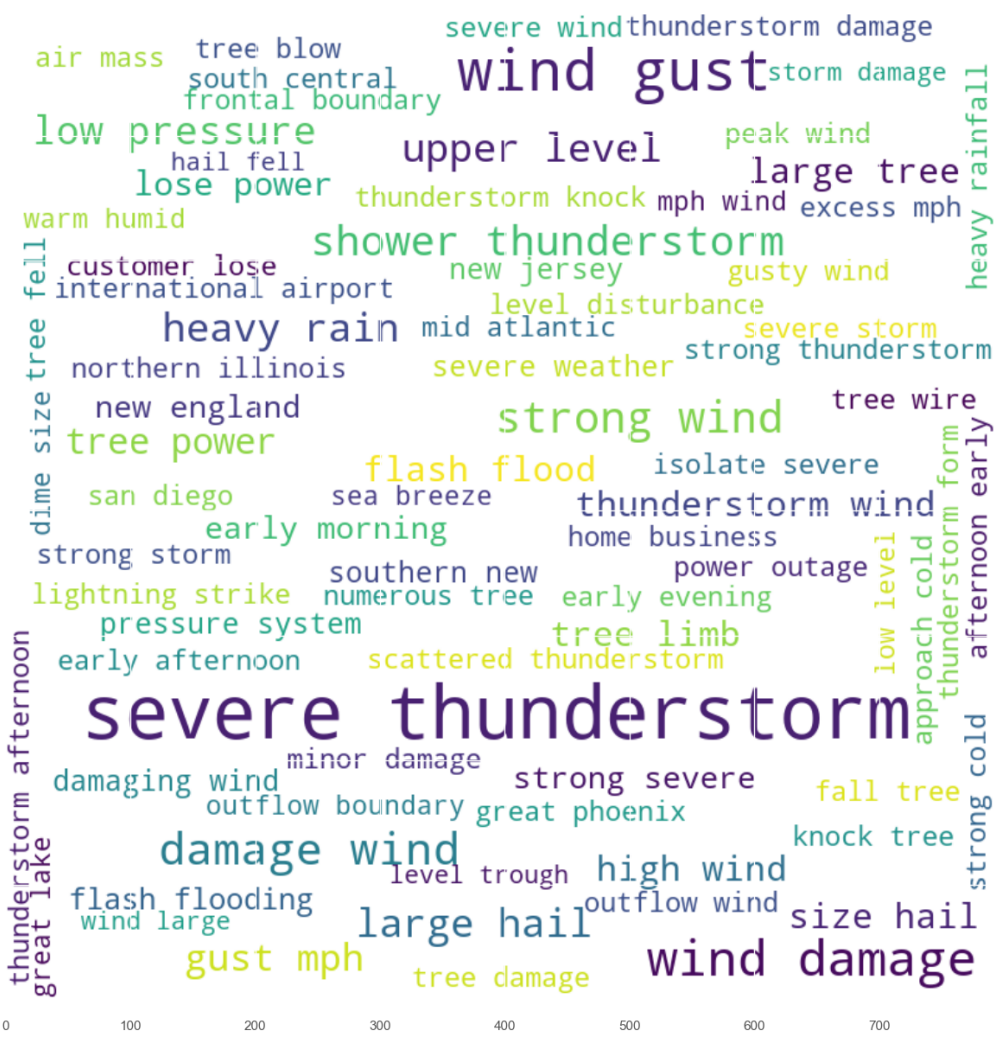

This figure is a graphical representation of the textual information from the *episode_narrative*, depicted in a word cloud format and employing a Bag of Words (BoW) methodology. The visibility of each word in the visualization is determined by its frequency within the specified corpus of text.

Figure 2: The Word Cloud Graph for Event Narrative

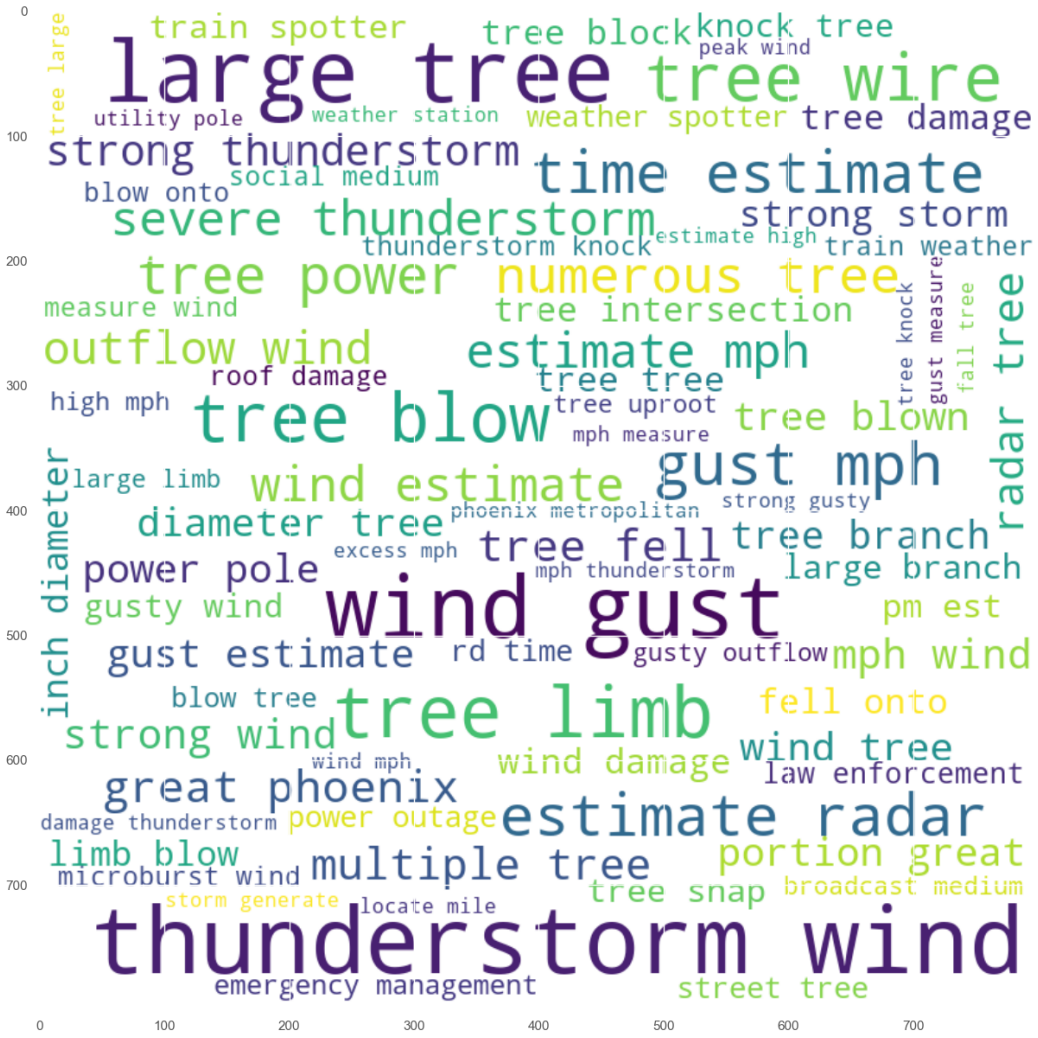

This figure is a graphical representation of the textual information from the *event_narrative*, depicted in a word cloud format and employing a Bag of Words (BoW) methodology. The visibility of each word in the visualization is determined by its frequency within the specified corpus of text.

Table 1: **Summary Statistics**

This table presents the summary statistics, including the mean and standard deviation of the variables. In Panel A, *Magnitude* is the weather event's scale for wind speeds (measured in knots). *Damage_Property* is the estimated property damage caused by the weather event. *Damage_Crops* is the estimated crop damage. Additionally, the logarithm of these variables is used for analysis. *Death* comes from the total of *Deaths_Direct* and *Deaths_Indirect*. *Deaths_Direct* and *Deaths_Indirect* show the number of deaths directly and indirectly caused by the weather event, respectively. *Injury* is the sum of *Injuries_Direct* and *Injuries_Indirect*. *Injuries_Direct* and *Injuries_Indirect* represent the number of injuries directly and indirectly caused by the weather event, respectively. The sample size is 674,201. CDS^{X-Year} represent the CDS level of the maturity term of X years in basis points, where X is set to 5, 3, 7, 10, and 20, respectively. *Slope* is the difference between 5-year CDS and 1-year CDS levels. The CDS sample size is 38,927. In Panel B, RET represent the stock return for the month. $SIZE$ is defined as the natural logarithm of the market value (MKV), where MKV is the product of the stock price (PRC) and the number of publicly held shares ($SHROUT$), recorded in thousands. TO represents the turnover and is calculated by dividing the sum of trading volumes (VOL), expressed in hundred shares for monthly data, by the product of the shares outstanding ($SHROUT$) and 1,000. MOM refers to the cumulative return over the last six months. *Leverage* is defined as the leverage measure, calculated as the ratio of the book value of total liability to the sum of the book value of total liability and the market value of equity. ROA represents the return on assets, defined as net income scaled by total assets. $SALE$ is the natural logarithm of sales. P/E is the price-earnings ratio. *TobinsQ* is defined as the market capitalization of common stock plus the liquidation value of preferred shares plus the book value of long-term debt divided by total assets. $CAPEX$ is defined as capital expenditures scaled by sales. The sample size is 674,201. Panel C is the correlation matrix between $Log(Magnitude)$, $Log(Damage\_Property)$, $Log(Damage\_Crops)$, $Death$ and $Injury$.

Panel A

Variable	Mean	SD
$Log(Magnitude)$	0.36	1.15
$Log(Damage\_Property)$	0.56	2.35
$Log(Damage\_Crops)$	0.02	0.40
$Death$	0.005	0.10
$Injury$	0.03	0.65
$CDS^{5\text{-}Year}$	143.73	254.60
$CDS^{3\text{-}Year}$	109.37	260.32
$CDS^{7\text{-}Year}$	158.62	241.74
$CDS^{10\text{-}Year}$	168.70	228.01
$CDS^{20\text{-}Year}$	177.70	226.52
$Slope$	69.56	128.37

Panel B

Variable	Mean	SD
RET	0.011	0.20
$SIZE$	13.03	2.17
TO	0.22	3.80
MOM	0.07	0.48
$Leverage$	0.21	0.17
ROA	0.0006	0.80
$SALE$	9.08	1.10
P/E	24.92	67.73
$TobinsQ$	1.82	1.65
$CAPEX$	0.03	0.04

Panel C

Variable	$Log(Magnitude)$	$Log(Damage\_Crops)$	$Log(Damage\_Property)$	$Death$	$Injury$
$Log(Magnitude)$	1.00				
$Log(Damage\_Crops)$	0.14	1.00			
$Log(Damage\_Property)$	0.76	0.16	1.00		
$Death$	0.15	0.07	0.13	1.00	
$Injury$	0.15	0.03	0.18	0.25	1.00

Table 2: **The levels of CDS of multiple terms and numeric weather risk measures**

This table presents the results of the panel regression with two-way clustering of standard errors to control for potential correlations within time periods (years and months) and cross-sectional units (counties). In Models 1 to 6, the dependent variables CDS_{t+1}^{X-Year} represent the CDS level for the next period of the maturity term of X years, where X is set to 5 for Models 1 and 2, and to 3, 7, 10, and 20 for Models 3 to 6 respectively. In all models, the independent variables include several metrics, in addition to the CDS level for the current period CDS_t^{X-Year}. *Magnitude* represents the event's scale, primarily utilized for wind speeds (measured in knots). *Damage_Property* quantifies the estimated property damage incurred by the weather event, while *Damage_Crops* quantifies the estimated damage to crops. Additionally, the logarithm of these variables is taken for analysis. *Death* is derived from the summation of *Deaths_Direct* and *Deaths_Indirect*. *Deaths_Direct* and *Deaths_Indirect* represent the count of deaths directly and indirectly attributed to the weather event, respectively. *Injury* is derived from the summation of *Injuries_Direct* and *Injuries_Indirect*. *Injuries_Direct* and *Injuries_Indirect* signify the count of injuries directly and indirectly resulting from the weather event, respectively. For each county and each year-month, the maximum value of *Magnitude* is selected, and the variables *Damage_Property*, *Damage_Crops*, *Death*, and *Injury* are aggregated by summing their values. The sample period is from January 2001 to October 2019. t-statistics are shown in parentheses. $*p < .1; **p < 0.05; ***p < 0.01$..

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
	CDS_{t+1}^{5-Year}	CDS_{t+1}^{5-Year}	CDS_{t+1}^{3-Year}	CDS_{t+1}^{7-Year}	$CDS_{t+1}^{10-Year}$	$CDS_{t+1}^{20-Year}$
$Log(Magnitude)_t$	0.568**		0.466*	0.553**	0.550**	0.618**
	(2.08)		(1.72)	(2.16)	(2.14)	(2.06)
$Log(Damage\_Property)_t$		0.245**				
		(2.03)				
$Log(Damage\_Crops)_t$	−0.248	−0.220	−0.349	−0.228	−0.255	−0.320
	(−0.88)	(−0.82)	(−1.61)	(−0.82)	(−0.87)	(−0.87)
$Death_t$	2.532	2.528	0.693	1.772	2.061	2.203
	(0.78)	(0.78)	(0.30)	(0.58)	(0.64)	(0.81)
$Injury_t$	0.05	0.03	0.054	0.069	−0.011	0.032
	(0.18)	(0.11)	(0.23)	(0.23)	(−0.03)	(0.11)
CDS_t^{5-Year}	0.947***	0.947***				
	(87.50)	(87.50)				
CDS_t^{3-Year}			0.927***			
			(65.43)			
CDS_t^{7-Year}				0.956***		
				(97.51)		
$CDS_t^{10-Year}$					0.958***	
					(109.04)	
$CDS_t^{20-Year}$						0.953***
						(92.85)
$Constant$	5.927***	5.985***	5.593***	5.745***	6.030***	7.306***
	(6.47)	(6.51)	(6.25)	(6.30)	(6.99)	(6.13)
N	37,733	37,733	35,620	35,703	35,395	26,930
R^2	94.1%	94.1%	92.9%	94.8%	95.0%	94.3%

Table 3: **The changes of CDS of multiple terms and numeric weather risk measures**

This table presents the results of the panel regression with two-way clustering of standard errors to control for potential correlations within time periods (years and months) and cross-sectional units (counties). In Models 1 to 6, the dependent variables $Chg\ CDS_{t+1}^{X-Year}$ represent the CDS change for in next period of the maturity term of X years, where X is set to 5 for Models 1 and 2, to 3 for Model 3, and to 7 for Model 4, respectively. In all models, the independent variables include several metrics, in addition to the CDS change for the current period $Chg\ CDS_{t}^{X-Year}$. *Magnitude* represents the event's scale, primarily utilized for wind speeds (measured in knots). *Damage_Property* quantifies the estimated property damage incurred by the weather event, while *Damage_Crops* quantifies the estimated damage to crops. Additionally, the logarithm of these variables is taken for analysis. *Death* is derived from the summation of *Deaths_Direct* and *Deaths_Indirect*. *Deaths_Direct* and *Deaths_Indirect* represent the count of deaths directly and indirectly attributed to the weather event, respectively. *Injury* is derived from the summation of *Injuries_Direct* and *Injuries_Indirect*. *Injuries_Direct* and *Injuries_Indirect* signify the count of injuries directly and indirectly resulting from the weather event, respectively. For each county and each year-month, the maximum value of *Magnitude* is selected, and the variables *Damage_Property*, *Damage_Crops*, *Death*, and *Injury* are aggregated by summing their values. The sample period is from January 2001 to October 2019. t-statistics are shown in parentheses. $*p < .1; **p < 0.05; ***p < 0.01$.

	Model 1	Model 2	Model 3	Model 4
	$Chg\ CDS_{t+1}^{5-Year}$		$Chg\ CDS_{t+1}^{3-Year}$	$Chg\ CDS_{t+1}^{7-Year}$
$Log(Magnitude)_t$	0.079*		0.208**	0.244**
	(1.71)		(2.16)	(2.03)
$Log(Damage\_Property)_t$		0.035*		
		(1.77)		
$Log(Damage\_Crops)_t$	0.251***	0.255***	0.027	0.042
	(2.63)	(2.73)	(0.20)	(0.24)
$Death_t$	−0.38	−0.385	−0.771	−0.607
	(−0.59)	(−0.59)	(−1.07)	(−0.54)
$Injury_t$	0.031	0.029	0.019	0.102
	(0.32)	(0.28)	(0.11)	(0.48)
$Chg\ CDS_{t}^{5-Year}$	0.009**	0.009**		
	(2.03)	(2.03)		
$Chg\ CDS_{t}^{3-Year}$			0.036***	
			(3.61)	
$Chg\ CDS_{t}^{7-Year}$				0.037***
				(3.45)
$Constant$	−0.517***	−0.510***	−0.524***	−0.145
	(−8.96)	(−9.11)	(−5.80)	(−1.50)
N	34,303	34,303	35,053	35,143
R^2	0.1%	0.1%	0.6%	0.4%

Table 4: **The percentage changes of CDS of multiple terms and numeric weather risk measures**

This table presents the results of the panel regression with two-way clustering of standard errors to control for potential correlations within time periods (years and months) and cross-sectional units (counties). In Models 1 to 6, the dependent variables $PChg\ CDS_{t+1}^{X-Year}$ represent the CDS percentage change for the next period of the maturity term of X years, where X is set to 5 for Models 1 and 2, to 3 for Models 3 and 4, and to 7 for Models 5 and 6, respectively. In all models, the independent variables include several metrics, in addition to the CDS percentage change for the current period $PChg\ CDS_{t}^{X-Year}$. *Magnitude* represents the event's scale, primarily utilized for wind speeds (measured in knots). *Damage_Property* quantifies the estimated property damage incurred by the weather event, while *Damage_Crops* quantifies the estimated damage to crops. Additionally, the logarithm of these variables is taken for analysis. *Death* is derived from the summation of *Deaths_Direct* and *Deaths_Indirect*. *Deaths_Direct* and *Deaths_Indirect* represent the count of deaths directly and indirectly attributed to the weather event, respectively. *Injury* is derived from the summation of *Injuries_Direct* and *Injuries_Indirect*. *Injuries_Direct* and *Injuries_Indirect* signify the count of injuries directly and indirectly resulting from the weather event, respectively. For each county and each year-month, the maximum value of *Magnitude* is selected, and the variables *Damage_Property*, *Damage_Crops*, *Death*, and *Injury* are aggregated by summing their values. The sample period is from January 2001 to October 2019. t-statistics are shown in parentheses. $*p < .1; **p < 0.05; ***p < 0.01$.

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
	$PChg\ CDS_{t+1}^{5-Year}$		$PChg\ CDS_{t+1}^{3-Year}$		$PChg\ CDS_{t+1}^{10-Year}$	
$Log(Magnitude)_t$	0.162**		0.222**		0.144**	
	(2.29)		(2.46)		(2.28)	
$Log(Damage\_Property)_t$		0.110***		0.081*		0.052*
		(3.06)		(1.71)		(1.66)
$Log(Damage\_Crops)_t$	0.487***	0.675***	0.522***	0.543***	0.323***	0.338***
	(3.68)	(2.73)	(3.73)	(4.16)	(2.63)	(2.92)
$Death_t$	−0.285	−0.036	−0.517	−0.472	−0.427	−0.392
	(−0.32)	(−0.05)	(−0.59)	(−0.52)	(−0.47)	(−0.42)
$Injury_t$	−0.023	−0.034	0.058	0.058	−0.019	−0.019
	(−0.20)	(−0.20)	(0.45)	(0.44)	(−0.17)	(−0.17)
$PChg\ CDS_t^{5-Year}$	0.069***	0.070***				
	(6.15)	(6.76)				
$PChg\ CDS_t^{3-Year}$			0.032***	0.032***		
			(3.66)	(3.67)		
$PChg\ CDS_t^{7-Year}$					0.028***	0.028***
					(2.97)	(2.99)
$Constant$	0.589***	0.607***	0.818***	0.851***	0.699***	0.721***
	(10.85)	(11.53)	(11.08)	(11.55)	(13.05)	(13.70)
N	36,639	35,746	35,059	35,059	35,153	35,153
R^2	0.5%	0.6%	0.2%	0.2%	0.2%	0.1%

Table 5: **CDS Slope and numeric weather risk measures**

This table presents the results of the panel regression with two-way clustering of standard errors to control for potential correlations within time periods (years and months) and cross-sectional units (counties). The dependent variables $Slope_{t+1}$ represent the next period CDS slope which is defined as the difference of 5-Year CDS level and 1-Year CDS level. The independent variables include several metrics, in addition to the current period CDS slope $Slope_t$. *Magnitude* represents the event's scale, primarily utilized for wind speeds (measured in knots). *Damage_Property* quantifies the estimated property damage incurred by the weather event, while *Damage_Crops* quantifies the estimated damage to crops. Additionally, the logarithm of these variables is taken for analysis. *Death* is derived from the summation of *Deaths_Direct* and *Deaths_Indirect*. *Deaths_Direct* and *Deaths_Indirect* represent the count of deaths directly and indirectly attributed to the weather event, respectively. *Injury* is derived from the summation of *Injuries_Direct* and *Injuries_Indirect*. *Injuries_direct* and *Injuries_Indirect* signify the count of injuries directly and indirectly resulting from the weather event, respectively. For each county and each year-month, the maximum value of *Magnitude* is selected, and the variables *Damage_Property*, *Damage_Crops*, *Death*, and *Injury* are aggregated by summing their values. The sample period is from January 2001 to October 2019. t-statistics are shown in parentheses. $*p < .1; **p < 0.05; ***p < 0.01$.

	Model 1	Model 2	Model 3	Model 4
$Log(Magnitude)_t$	0.432**		2.844***	
	(2.45)		(3.10)	
$Log(Damage\_Property)_t$		0.187**		1.070**
		(2.42)		(2.21)
$Log(Damage\_Crops)_t$	0.033	0.053	−0.374	−0.120
	(0.15)	(0.26)	(−0.50)	(−0.15)
$Death_t$	0.467	0.444	2.234	2.683
	(0.44)	(0.43)	(0.52)	(0.62)
$Injury_t$	0.018	0.006	0.002	−0.006
	(0.12)	(0.04)	0.00	(−0.01)
$Slope_t$	0.871***	0.871***		
	(21.50)	(21.53)		
Constant	8.792***	8.828***	63.305***	63.707***
	(3.32)	(3.32)	(17.23)	(17.44)
N	32,907	32,907	33,606	33,606
R^2	84.5%	84.5%	0.3%	0.2%

Table 6: **Investment Grade vs Speculative Grade**

This table presents the results of the panel regression with two-way clustering of standard errors to control for potential correlations within time periods (years and months) and cross-sectional units (counties). The dependent variables CDS_{t+1}^{5-Year} represent the CDS level for the next period of the maturity term of 5 years. The independent variables include several metrics, in addition to the CDS level for the current period CDS_t^{X-Year}. *Magnitude* represents the event's scale, primarily utilized for wind speeds (measured in knots). *Damage_Crops* quantifies the estimated damage to crops. Additionally, the logarithm of these variables is taken for analysis. *Death* is derived from the summation of *Deaths_Direct* and *Deaths_Indirect*. *Deaths_Direct* and *Deaths_Indirect* represent the count of deaths directly and indirectly attributed to the weather event, respectively. *Injury* is derived from the summation of *Injuries_Direct* and *Injuries_Indirect*. *Injuries_direct* and *Injuries_Indirect* signify the count of injuries directly and indirectly resulting from the weather event, respectively. For each county and each year-month, the maximum value of *Magnitude* is selected, and the variables *Damage_Crops*, *Death*, and *Injury* are aggregated by summing their values. The sample period is from January 2001 to October 2019. *t*-statistics are shown in parentheses. $*p < .1; **p < 0.05; ***p < 0.01$.

	Model 1	Model 2
	Investment	Junk
$Log(Magnitude)_t$	0.260*	0.909**
	(1.96)	(2.23)
$Log(Damage\_Crops)_t$	0.018	−1.211
	(0.13)	(1.59)
$Death_t$	0.427	12.667
	(0.37)	(1.13)
$Injury_t$	0.085	−0.068
	(0.63)	(−0.09)
CDS_t^{5-Year}	0.921***	0.927***
	(86.75)	(109.23)
Constant	5.791***	16.838***
	(6.85)	(7.74)
N	27,212	10,521
R^2	90.1%	91.5%

Table 7: **Expected stock return and numeric weather risk measures**

This table presents the results of the panel regression with two-way clustering of standard errors to control for potential correlations within time periods (years and months) and cross-sectional units (counties). In all models, the dependent variables RET_{t+1} represent the stock return for the next month. In all models, the independent variables include several metrics, in addition to the stock return for the current month RET_t. *Magnitude* represents the event's scale, primarily utilized for wind speeds (measured in knots). *Damage_Crops* quantifies the estimated damage to crops. Additionally, the logarithm of these variables is taken for analysis. *Death* is derived from the summation of *Deaths_Direct* and *Deaths_Indirect*. *Deaths_Direct* and *Deaths_Indirect* represent the count of deaths directly and indirectly attributed to the weather event, respectively. *Injury* is derived from the summation of *Injuries_Direct* and *Injuries_Indirect*. *Injuries_direct* and *Injuries_Indirect* signify the count of injuries directly and indirectly resulting from the weather event, respectively. For each county and each year-month, the maximum value of *Magnitude* is selected, and the variables *Damage_Crops*, *Death*, and *Injury* are aggregated by summing their values. $SIZE$ is defined as the natural logarithm of the market value (MKV), where MKV is the product of the stock price (PRC) and the number of publicly held shares ($SHROUT$), recorded in thousands. TO represents the turnover and is calculated by dividing the sum of trading volumes (VOL), expressed in hundred shares for monthly data, by the product of the shares outstanding ($SHROUT$) and 1,000. MOM refers to the cumulative return over the last six months. The sample period is from January 1997 to December 2023. t-statistics are shown in parentheses. $*p < .1; **p < 0.05; ***p < 0.01$.

	Model 1 RET_{t+1}	Model 2 RET_{t+2}	Model 3 RET_{t+3}
$Log(Magnitude)_t$	−0.153**	−0.239**	−0.200*
	(−2.24)	(−2.41)	(−1.78)
$Log(Damage\_Crops)_t$	−0.332***	−0.084	−0.157
	(−2.74)	(−0.66)	(−0.78)
$Death_t$	−0.267	−0.127	−1.021
	(−0.38)	(−0.40)	(−1.18)
$Injury_t$	0.097	−0.038	−0.015
	(1.33)	(−0.43)	(−0.22)
RET_t	−1.190	0.783	−0.188
	(−1.35)	(0.59)	(−0.17)
$SIZE_t$	−0.054	−0.209***	−0.217***
	(−1.05)	(−2.63)	(−2.60)
TO_t	−0.191***	−0.265**	−0.022
	(−2.68)	(−2.27)	(−1.40)
MOM_t	0.133	0.538	0.108
	(0.35)	(1.02)	(0.20)
$Constant$	1.991**	3.830***	3.976***
	(2.50)	(3.03)	(2.98)
N	600,256	589,596	604,931
R^2	0.1%	0.1%	0.1%

Table 8: **Expected portfolio return and numeric weather risk measures**

This table reports the average monthly returns of portfolios. Panel A reports the equal-weighted quartile portfolio returns sorted by *Magnitude*. Panel B reports the value-weighted quartile portfolio returns sorted by *Magnitude*. Panel C reports the return of an equal-weighted portfolio that is long the bottom tercile of stocks and short the top tercile ranked by *Magnitude*, in the subsamples of stocks sorted by proxies of limits to arbitrage–size and stock price level. The table reports the average raw returns *Ret*, the capital asset pricing model (CAPM) alphas α_{CAPM}, Fama-French 3-factor alphas (FF3) alphas α_{FF3}, and Carhart four-factor (Carhart-4) alphas α_{C4}. All returns are in percent. The sample period is from January 1997 to December 2023. The numbers enclosed in brackets represent *t*-statistics adjusted by the Newey-West method. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Equal-Weighted Quartile Portfolio Returns and Alphas for *Magnitude*

Magnitude	1 (Low)	2	3	4 (High)	Low−High
Ret	0.96	0.70	0.70	0.37	0.59**
	(2.59)	(1.82)	(1.91)	(0.94)	(1.96)
α_{CAPM}	0.88	0.69	0.58	0.31	0.57*
	(2.22)	(1.69)	(1.47)	(0.71)	(1.85)
α_{FF3}	0.91	0.69	0.57	0.36	0.55*
	(2.31)	(1.68)	(1.45)	(0.81)	(1.75)
α_{C4}	0.95	0.75	0.62	0.48	0.47*
	(2.39)	(1.79)	(1.53)	(1.09)	(1.68)

Panel B: Value-Weighted Quartile Portfolio Returns and Alphas for *Magnitude*

Magnitude	1 (Low)	2	3	4 (High)	Low−High
Ret	1.46	0.93	1.01	0.47	0.99**
	(3.51)	(2.29)	(2.56)	(1.28)	(2.21)
α_{CAPM}	1.33	0.90	0.99	0.46	0.87**
	(3.08)	(2.06)	(2.25)	(1.17)	(2.01)
α_{FF3}	1.34	0.88	1.00	0.46	0.88**
	(3.11)	(2.03)	(2.28)	(1.18)	(1.97)
α_{C4}	1.32	0.94	1.04	0.57	0.75*
	(3.18)	(2.11)	(2.33)	(1.45)	(1.83)

Panel C: Double Sort Portfolio Returns and Alphas for *Magnitude*

	Size				*Stock Price*			
	Ret	α_{CAPM}	α_{FF3}	α_{C4}	*Ret*	α_{CAPM}	α_{FF3}	α_{C4}
1 (low)	0.92***	1.01***	0.98***	0.90**	0.80***	0.81**	0.77**	0.68*
	(2.51)	(2.72)	(2.66)	(2.46)	(2.10)	(2.11)	(2.01)	(1.77)
2 (high)	−0.18	−0.19	−0.18	−0.23	−0.04	−0.02	−0.01	−0.04
	(−0.66)	(−0.68)	(−0.67)	(−0.84)	(−0.18)	(−0.08)	(−0.02)	(−0.18)

Table 9: **Leverage and numeric weather risk measures**

This table presents the results of the panel regression with two-way clustering of standard errors to control for potential correlations within time periods (years and months) and cross-sectional units (counties). In all models, the dependent variables $Leverage_{t+1}$ represent *Leverage* for the next month. *Leverage* is defined as the leverage measure, calculated as the ratio of the book value of total liability to the sum of the book value of total liability and the market value of equity. In all models, the independent variables include several metrics. *Magnitude* represents the event's scale, primarily utilized for wind speeds (measured in knots). *Damage_Crops* quantifies the estimated damage to crops. Additionally, the logarithm of these variables is taken for analysis. *Death* is derived from the summation of *Deaths_Direct* and *Deaths_Indirect*. *Deaths_Direct* and *Deaths_Indirect* represent the count of deaths directly and indirectly attributed to the weather event, respectively. *Injury* is derived from the summation of *Injuries_Direct* and *Injuries_Indirect*. *Injuries_direct* and *Injuries_Indirect* signify the count of injuries directly and indirectly resulting from the weather event, respectively. For each county and each year-month, the maximum value of *Magnitude* is selected, and the variables *Damage_Crops*, *Death*, and *Injury* are aggregated by summing their values. The sample period is from January 1997 to December 2023. *t*-statistics are shown in parentheses. $*p < .1; **p < 0.05; ***p < 0.01$.

	Model 1	Model 2	Model 3
	$Leverage_{t+1}$	$Chg\ Leverage_{t+1}$	$PChg\ Leverage_{t+1}$
$Log(Magnitude)_t$	0.000	0.000	0.001
	(1.27)	(1.07)	(0.79)
$Log(Damage\_Crops)_t$	0.001**	0.001***	0.002**
	(2.51)	(2.64)	(2.33)
$Death_t$	0.004	0.004	0.012
	(1.35)	(1.38)	(1.20)
$Injury_t$	0.000	0.000	−0.001
	(0.89)	(0.91)	(−0.94)
$Leverage_t$	0.991***		
	(487.51)		
$Constant$	0.004***	0.000	0.002
	(4.33)	(0.37)	(0.81)
N	38,693	38,693	38,693
R^2	98.4%	0.1%	0.1%

Table 10: **Profitability and numeric weather risk measures**

This table presents the results of the panel regression with two-way clustering of standard errors to control for potential correlations within time periods (years and months) and cross-sectional units (counties). In Model 1, the dependent variable ROA_{t+1} represents the return on assets, defined as net income scaled by total assets in the next period. In Model 2, the dependent variable is the change in ROA in the current period. In Models 3 and 4, the dependent variables are the change and percentage change in $SALE$ (natural logarithm of sales) in the next period. In Model 5, the dependent variable is the percentage change in P/E (price-earnings ratio) in the next period. In all models, the independent variables include several metrics. *Magnitude* represents the event's scale, primarily utilized for wind speeds (measured in knots). *Damage_Crops* quantifies the estimated damage to crops. Additionally, the logarithm of these variables is taken for analysis. *Death* is derived from the summation of *Deaths_Direct* and *Deaths_Indirect*. *Deaths_Direct* and *Deaths_Indirect* represent the count of deaths directly and indirectly attributed to the weather event, respectively. *Injury* is derived from the summation of *Injuries_Direct* and *Injuries_Indirect*. *Injuries_direct* and *Injuries_Indirect* signify the count of injuries directly and indirectly resulting from the weather event, respectively. For each county and each year-month, the maximum value of *Magnitude* is selected, and the variables *Damage_Crops*, *Death*, and *Injury* are aggregated by summing their values. The sample period is from January 1997 to December 2023. t-statistics are shown in parentheses. $*p < .1; **p < 0.05; ***p < 0.01$.

	Model 1	Model 2	Model 3	Model 4	Model 5
	ROA_{t+1}	$Chg\ ROA_{t+1}$	$Chg\ SALE_{t+1}$	$PChg\ SALE_{t+1}$	$PChg\ P/E_{t+1}$
$Log(Magnitude)_t$	−0.0003**	−0.0002**	−0.001**	−0.0001**	−0.008***
	(−2.56)	(−2.40)	(−2.34)	(−2.48)	(−3.35)
$Log(Damage\_Crops)_t$	0.000	0.000	0.000	0.000	−0.003**
	(0.01)	(0.27)	(0.26)	(0.21)	(−2.45)
$Death_t$	0.003	0.003	0.001	0.000	−0.014
	(0.93)	(0.88)	(0.81)	(0.88)	(−1.29)
$Injury_t$	0.000	0.000	0.000	0.000	0.000
	(0.76)	(0.73)	(0.98)	(1.03)	(0.07)
ROA_t	0.956***				
	(67.10)				
$Constant$	0.002***	0.000	0.003*	0.000**	0.040***
	(2.74)	(0.43)	(1.89)	(1.98)	(4.06)
N	38,693	38,693	38,693	38,693	33,352
R^2	91.5%	0.1%	0.1%	0.1%	0.1%

Table 11: **Corporate investment and numeric weather risk measures**

This table presents the results of the panel regression with two-way clustering of standard errors to control for potential correlations within time periods (years and months) and cross-sectional units (counties). In Models 1 to 3, the dependent variables $TobinsQ_{t+1}$ represent $TobinsQ$ for the next month, which is defined as the market capitalization of common stock plus the liquidation value of preferred shares plus the book value of long-term debt divided by total assets. In Models 4 to 6, the dependent variables $CAPEX_{t+1}$ represent capital expenditures ($CAPEX$) for the next month, which is defined as capital expenditures scaled by sales. In all models, the independent variables include several metrics. *Magnitude* represents the event's scale, primarily utilized for wind speeds (measured in knots). *Damage_Crops* quantifies the estimated damage to crops. Additionally, the logarithm of these variables is taken for analysis. *Death* is derived from the summation of *Deaths_Direct* and *Deaths_Indirect*. *Deaths_Direct* and *Deaths_Indirect* represent the count of deaths directly and indirectly attributed to the weather event, respectively. *Injury* is derived from the summation of *Injuries_Direct* and *Injuries_Indirect*. *Injuries_direct* and *Injuries_Indirect* signify the count of injuries directly and indirectly resulting from the weather event, respectively. For each county and each year-month, the maximum value of *Magnitude* is selected, and the variables *Damage_Crops*, *Death*, and *Injury* are aggregated by summing their values. The sample period is from January 1997 to December 2023. t-statistics are shown in parentheses. $*p < .1; **p < 0.05; ***p < 0.01$.

	Model 1	Model 2	Model 3	Model 4	Model 5	Model 6
	$TobinsQ_{t+1}$	$TobinsQ_{t+1}$	$TobinsQ_{t+1}$	$CAPEX_{t+1}$	$CAPEX_{t+1}$	$CAPEX_{t+1}$
$Log(Magnitude)_t$	−0.004***		−0.004***	0.001***		0.001***
	(−2.80)		(−2.65)	(4.53)		(4.48)
$Log(Damage\_Crops)_t$		−0.004*	−0.002		0.0004**	0.0001
		(−1.65)	(−0.86)		(2.02)	(0.63)
$Death_t$			−0.006			0.001
			(−0.89)			(0.93)
$Injury_t$			0.001			0.0005
			(0.61)			(0.53)
$TobinsQ_t$	0.963***	0.963***	0.963***			
	(78.74)	(78.77)	(78.73)			
$CAPEX_t$				0.855***	0.856***	0.855***
				(32.15)	(32.28)	(32.15)
$Constant$	0.070***	0.068***	0.070***	0.004***	0.004***	0.004***
	(3.37)	(3.31)	(3.37)	(10.99)	(11.18)	(10.99)
N	162,719	162,719	162,719	162,465	162,465	162,465
R^2	92.2%	92.2%	92.2%	73.2%	73.1%	73.2%

Table 12: **5-year CDS level and textural and numeric weather risk measures**

This table presents the results of both Machine Learning (ML) and Artificial Intelligence (AI) with Deep Learning (DL) models. Models 1 to 4 and Models 5 to 8 represent four ML models: Linear Regression (LR), Decision Tree Regression (DT), Random Forest Regression (RF), and Extreme Gradient Boosting (XGB). Models 1 to 4 use TF-IDF for textual feature extraction, while Models 5 to 8 utilize Word2Vec. Similarly, Models 9 to 11 and Models 12 to 14 represent three AI and DL models: Deep Neural Networks (DNN), Convolutional Neural Networks (CNN), and Recurrent Neural Networks (RNN). Models 9 to 11 utilize TF-IDF, whereas Models 10 to 14 use BERT for textual feature extraction. Panel A incorporates the textual feature *episode_narrative*, while Panel B employs *event_narrative*. Both panels utilize numeric features: *magnitude*, *damage_crops*, *damage_property*, *deaths*, and *injuries*. Additionally, dummy variables for state, county, and year-month are included to account for climate-related risks, along with the current value of the target variable, which is the 5-year CDS level. The target variable represents the value of the next month. The regression results of the target values on the predicted values of the test set, including coefficients, standard errors, p-values, and adjusted R^2 values, are reported. The sample size is 3,130 firm-month observations from January 2001 to October 2019. N represents the size of the test sample which is 20% of the sample size. Numbers enclosed in brackets represent standard errors. Statistical significance levels are denoted by ***, **, and *, representing significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Episode

	TF-IDF				Word2Vec (DL)				TF-IDF & AI/DL			BERT/LLM & AI/DL		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
	LR	DT	RF	XGB	LR	DT	RF	XGB	DNN	CNN	RNN	DNN	CNN	RNN
Pred.	1.043***	1.031***	1.034***	1.021***	1.054***	0.983***	1.030***	1.030***	0.989***	1.094***	1.089***	1.379***	1.051***	1.050***
	(0.011)	(0.012)	(0.010)	(0.010)	(0.010)	(0.016)	(0.011)	(0.013)	(0.009)	(0.010)	(0.011)	(0.013)	(0.010)	(0.010)
Const.	−5.088***	−0.555	−3.214*	−1.879	−5.698***	3.319	−3.930**	−2.466	−5.179***	−7.013***	−8.326***	−7.488***	−4.210**	−6.919***
	(−1.901)	(−2.055)	(−1.727)	(−1.714)	(−1.681)	(2.764)	(−1.927)	(−2.222)	(−1.660)	(−1.682)	(−1.885)	(−1.703)	(−1.693)	(−1.723)
N	626	626	626	626	626	626	626	626	626	626	626	626	626	626
Adj. R^2	93.56%	92.25%	94.57%	94.59%	94.95%	86.07%	93.33%	91.15%	95.05%	95.00%	93.83%	94.90%	94.82%	94.75%

Panel B: Event

	TF-IDF				Word2Vec (DL)				TF-IDF & AI/DL			BERT/LLM & AI/DL		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
	LR	DT	RF	XGB	LR	DT	RF	XGB	DNN	CNN	RNN	DNN	CNN	RNN
Pred.	0.957***	0.947***	0.980***	0.967***	0.978***	0.947***	0.995***	0.995***	0.956***	1.167***	1.045***	1.243***	1.021***	1.022***
	(0.013)	(0.015)	(0.012)	(0.011)	(0.012)	(0.015)	(0.011)	(0.012)	(0.012)	(0.019)	(0.016)	(0.013)	(0.010)	(0.010)
Const.	6.159***	9.420***	4.022*	4.579**	3.231	7.996***	1.324	2.381	4.573**	−0.754	1.511	−5.471***	−0.422	−0.084
	(2.353)	(2.691)	(2.137)	(2.058)	(2.147)	(2.766)	(2.074)	(2.214)	(2.155)	(−2.840)	(2.744)	(−1.857)	(−1.840)	(−1.834)
N	548	548	548	548	548	548	548	548	548	548	548	514	514	514
Adj. R^2	91.01%	88.13%	92.64%	93.13%	92.63%	87.64%	93.21%	92.23%	92.50%	87.90%	88.45%	95.03%	94.91%	94.93%

Table 13: **CDS slope and textural and numeric weather risk measures**

This table presents the results of both Machine Learning (ML) and Artificial Intelligence (AI) with Deep Learning (DL) models. Models 1 to 4 and Models 5 to 8 represent four ML models: Linear Regression (LR), Decision Tree Regression (DT), Random Forest Regression (RF), and Extreme Gradient Boosting (XGB). Models 1 to 4 use TF-IDF for textual feature extraction, while Models 5 to 8 utilize Word2Vec. Similarly, Models 9 to 11 and Models 12 to 14 represent three AI and DL models: Deep Neural Networks (DNN), Convolutional Neural Networks (CNN), and Recurrent Neural Networks (RNN). Models 9 to 11 utilize TF-IDF, whereas Models 10 to 14 use BERT for textual feature extraction. Panel A incorporates the textual feature *episode_narrative*, while Panel B employs *event_narrative*. Both panels utilize numeric features: *magnitude*, *damage_crops*, *damage_property*, *deaths*, and *injuries*. Additionally, dummy variables for state, county, and year-month are included to account for climate-related risks, along with the current value of the target variable, which is the CDS slope. The target variable represents the value of the next month. The regression results of the target values on the predicted values of the test set, including coefficients, standard errors, p-values, and adjusted R^2 values, are reported. The sample size is 3,130 firm-month observations from January 2001 to October 2019. N represents the size of the test sample which is 20% of the sample size. Numbers enclosed in brackets represent standard errors. Statistical significance levels are denoted by ***, **, and *, representing significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Episode

	TF-IDF				Word2Vec (DL)				TF-IDF & AI/DL			BERT/LLM & AI/DL		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
	LR	DT	RF	XGB	LR	DT	RF	XGB	DNN	CNN	RNN	DNN	CNN	RNN
Pred.	0.983***	0.920***	0.971***	0.943***	0.999***	0.914***	0.978***	0.970***	0.998***	1.030***	5.992***	1.175***	0.981***	0.984***
	(0.015)	(0.015)	(0.013)	(0.013)	(0.015)	(0.015)	(0.013)	(0.015)	(0.014)	(0.022)	(0.394)	(0.015)	(0.012)	(0.013)
Const.	1.185	5.546***	2.017	2.814*	0.201	5.438***	0.726	1.632	2.161	−0.156	−0.643	5.019***	0.049	1.067
	(1.561)	(1.645)	(1.372)	(1.444)	(1.571)	(1.609)	(1.397)	(1.570)	(1.450)	(−2.112)	(−3.259)	(1.305)	(1.319)	(1.316)
N	571	571	571	571	571	571	571	571	571	571	571	571	571	571
Adj. R^2	88.25%	86.22%	90.65%	89.59%	88.26%	86.80%	90.50%	88.04%	89.60%	80.11%	88.88%	91.12%	91.56%	91.48%

Panel B: Event

	TF-IDF				Word2Vec (DL)				TF-IDF & AI/DL			BERT/LLM & AI/DL		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
	LR	DT	RF	XGB	LR	DT	RF	XGB	DNN	CNN	RNN	DNN	CNN	RNN
Pred.	0.923***	0.912***	0.982***	0.989***	0.920***	0.903***	0.998***	0.998***	1.022***	0.981***	1.041***	1.181***	1.021***	1.020***
	(0.019)	(0.017)	(0.014)	(0.013)	(0.019)	(0.016)	(0.014)	(0.014)	(0.014)	(0.021)	(0.016)	(0.018)	(0.014)	(0.013)
Const.	10.237***	8.694***	2.603*	2.715*	10.077***	8.424***	1.2	1.981	1.937	−5.067**	4.544***	2.36	1.065	−1.639
	(2.021)	(1.863)	(1.479)	(1.416)	(2.037)	(1.795)	(1.526)	(1.514)	(1.491)	(−2.366)	(1.603)	(1.687)	(1.492)	(−1.503)
N	504	504	504	504	504	504	504	504	504	504	504	473	473	473
Adj. R^2	82.82%	85.49%	91.31%	91.99%	82.63%	86.50%	90.95%	90.99%	91.25%	81.43%	89.61%	90.17%	92.34%	92.50%

Table 14: **Expected stock return and textural and numeric weather risk measures**

This table presents the results of both Machine Learning (ML) and Artificial Intelligence (AI) with Deep Learning (DL) models. Models 1 to 4 and Models 5 to 8 represent four ML models: Linear Regression (LR), Decision Tree Regression (DT), Random Forest Regression (RF), and Extreme Gradient Boosting (XGB). Models 1 to 4 use TF-IDF for textual feature extraction, while Models 5 to 8 utilize Word2Vec. Similarly, Models 9 to 11 and Models 12 to 14 represent three AI and DL models: Deep Neural Networks (DNN), Convolutional Neural Networks (CNN), and Recurrent Neural Networks (RNN). Models 9 to 11 utilize TF-IDF, whereas Models 10 to 14 use BERT for textual feature extraction. Panel A incorporates the textual feature *episode_narrative*, while Panel B employs *event_narrative*. Both panels utilize numeric features: *magnitude*, *damage_crops*, *damage_property*, *deaths*, and *injuries*. Additionally, dummy variables for state, county, and year-month are included to account for climate-related risks, along with the current value and the lagged one-period and two-period values of the target variable, which is the stock price at month-end. The target variable represents the value of the next month. The regression results of the target values on the predicted values of the test set, including coefficients, standard errors, p-values, and adjusted R^2 values, are reported. The sample size is 18,865 firm-month observations from January 1997 to December 2023. N represents the size of the test sample which is 20% of the sample size. Numbers enclosed in brackets represent standard errors. Statistical significance levels are denoted by ***, **, and *, representing significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Episode

	TF-IDF				Word2Vec (DL)				TF-IDF & AI/DL			BERT/LLM & AI/DL		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
	LR	DT	RF	XGB	LR	DT	RF	XGB	DNN	CNN	RNN	DNN	CNN	RNN
Pred.	1.005***	0.965***	1.002***	1.002***	1.007***	0.971***	1.002***	0.999***	1.013***	0.997***	1.003***	1.995***	0.996***	1.002***
	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)	(0.003)	(0.004)	(0.004)	(0.003)	(0.006)	(0.003)	(0.003)
Const.	−0.086	0.810***	−0.022	−0.025	−0.114	0.767***	−0.025	0.056	0.030***	−0.054***	−0.005	−1.530***	−0.628***	−0.454***
	(−0.081)	(0.110)	(−0.083)	(−0.081)	(−0.081)	(0.112)	(−0.083)	(0.083)	(0.003)	(−0.004)	(−0.003)	(−0.085)	(−0.081)	(−0.080)
N	3,773	3,773	3,773	3,773	3,773	3,773	3,773	3,773	3,773	3,773	3,773	3,773	3,773	3,773
Adj. R^2	96.44%	93.29%	96.27%	96.45%	96.50%	92.99%	96.33%	96.29%	95.56%	94.55%	95.80%	96.50%	96.60%	96.61%

Panel B: Event

	TF-IDF				Word2Vec (DL)				TF-IDF & AI/DL			BERT/LLM & AI/DL		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
	LR	DT	RF	XGB	LR	DT	RF	XGB	DNN	CNN	RNN	DNN	CNN	RNN
Pred.	0.994***	0.975***	0.994***	0.995***	0.996***	0.960***	0.993***	0.988***	0.993***	1.012***	0.933***	1.874***	1.000***	0.980***
	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	(0.01)	(0.00)	(0.01)	(0.00)	(0.00)
Const.	0.128	0.559***	0.101	0.088	0.085	0.989***	0.119	0.212**	0.018***	0.007	−0.001	−0.461***	−0.225**	0.355***
	(0.10)	(0.13)	(0.10)	(0.10)	(0.10)	(0.14)	(0.10)	(0.10)	(0.00)	(0.00)	(−0.00)	(−0.10)	(−0.10)	(0.10)
N	2,796	2,796	2,796	2,796	2,796	2,796	2,796	2,796	2,796	2,796	2,796	2,796	2,796	2,796
Adj. R^2	96.65%	93.92%	96.58%	96.72%	96.71%	93.22%	96.55%	96.38%	95.79%	94.71%	95.89%	96.56%	96.75%	96.77%

Table 15: **Leverage and textural and numeric weather risk measures**

This table presents the results of both Machine Learning (ML) and Artificial Intelligence (AI) with Deep Learning (DL) models. Models 1 to 4 and Models 5 to 8 represent four ML models: Linear Regression (LR), Decision Tree Regression (DT), Random Forest Regression (RF), and Extreme Gradient Boosting (XGB). Models 1 to 4 use TF-IDF for textual feature extraction, while Models 5 to 8 utilize Word2Vec. Similarly, Models 9 to 11 and Models 12 to 14 represent three AI and DL models: Deep Neural Networks (DNN), Convolutional Neural Networks (CNN), and Recurrent Neural Networks (RNN). Models 9 to 11 utilize TF-IDF, whereas Models 10 to 14 use BERT for textual feature extraction. Panel A incorporates the textual feature *episode_narrative*, while Panel B employs *event_narrative*. Both panels utilize numeric features: *magnitude*, *damage_crops*, *damage_property*, *deaths*, and *injuries*. Additionally, dummy variables for state, county, and year-month are included to account for climate-related risks, along with the current value of the target variable, which is the leverage ratio as defined in Table 9. The target variable represents the value of the next month. The regression results of the target values on the predicted values of the test set, including coefficients, standard errors, p-values, and adjusted R^2 values, are reported. The sample size is 13,230 firm-month observations from January 1997 to December 2023. N represents the size of the test sample which is 20% of the sample size. Numbers enclosed in brackets represent standard errors. Statistical significance levels are denoted by ***, **, and *, representing significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Episode

	TF-IDF				Word2Vec (DL)				TF-IDF & AI/DL			BERT/LLM & AI/DL		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
	LR	DT	RF	XGB	LR	DT	RF	XGB	DNN	CNN	RNN	DNN	CNN	RNN
Pred.	0.991***	0.965***	0.987***	0.989***	0.993***	0.967***	0.991***	0.991***	1.007***	0.988***	1.001***	1.070***	0.971***	0.984***
	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)	(0.004)	(0.003)	(0.004)	(0.004)	(0.004)	(0.003)	(0.005)	(0.005)	(0.004)
Const.	0.002**	0.007***	0.003***	0.003***	0.002*	0.008***	0.002**	0.002**	−0.001	−0.001	−0.006***	0.000	0.006***	0.016***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(−0.001)	(−0.001)	(−0.001)	(0.001)	(0.001)	(0.001)
N	2,646	2,646	2,646	2,646	2,646	2,646	2,646	2,646	2,646	2,646	2,646	2,296	2,296	2,296
Adj. R^2	97.18%	95.29%	97.13%	97.01%	97.28%	94.59%	97.09%	96.71%	96.16%	94.83%	97.00%	96.03%	95.03%	95.50%

Panel B: Event

	TF-IDF				Word2Vec (DL)				TF-IDF & AI/DL			BERT/LLM & AI/DL		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
	LR	DT	RF	XGB	LR	DT	RF	XGB	DNN	CNN	RNN	DNN	CNN	RNN
Pred.	0.998***	0.982***	0.995***	0.996***	1.000***	0.979***	1.001***	0.999***	0.992***	0.957***	1.019***	1.021***	0.976***	0.978***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)	(0.004)	(0.004)	(0.005)	(0.004)	(0.005)	(0.005)	(0.005)
Const.	0.001	0.005***	0.001	0.001	0	0.005***	0	0.001	0.003***	−0.005***	−0.005***	0.034***	0.035***	0.021***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(−0.001)	(−0.001)	(0.001)	(0.001)	(0.001)
N	2,327	2,327	2,327	2,327	2,327	2,327	2,327	2,327	2,327	2,327	2,327	2,017	2,017	2,017
Adj. R^2	96.95%	95.70%	96.94%	96.83%	97.15%	94.08%	96.71%	96.77%	96.25%	94.20%	96.36%	95.50%	94.58%	94.77%

Table 16: **Corporate investment and textural and numeric weather risk measures**

This table presents the results of both Machine Learning (ML) and Artificial Intelligence (AI) with Deep Learning (DL) models. Models 1 to 4 and Models 5 to 8 represent four ML models: Linear Regression (LR), Decision Tree Regression (DT), Random Forest Regression (RF), and Extreme Gradient Boosting (XGB). Models 1 to 4 use TF-IDF for textual feature extraction, while Models 5 to 8 utilize Word2Vec. Similarly, Models 9 to 11 and Models 12 to 14 represent three AI and DL models: Deep Neural Networks (DNN), Convolutional Neural Networks (CNN), and Recurrent Neural Networks (RNN). Models 9 to 11 utilize TF-IDF, whereas Models 10 to 14 use BERT for textual feature extraction. Panel A incorporates the textual feature *episode_narrative*, while Panel B employs *event_narrative*. Both panels utilize numeric features: *magnitude*, *damage_crops*, *damage_property*, *deaths*, and *injuries*. Additionally, dummy variables for state, county, and year-month are included to account for climate-related risks, along with the current value of the target variable, which is *TobinsQ* as defined in Table 11. The target variable represents the value of the next month. The regression results of the target values on the predicted values of the test set, including coefficients, standard errors, p-values, and adjusted R^2 values, are reported. *N* represents the size of the test sample. The sample size is 12,825 firm-month observations from January 1997 to December 2023. *N* represents the size of the test sample which is 20% of the sample size. Numbers enclosed in brackets represent standard errors. Statistical significance levels are denoted by ***, **, and *, representing significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Episode

	TF-IDF				Word2Vec (DL)				TF-IDF & AI/DL			BERT/LLM & AI/DL		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
	LR	DT	RF	XGB	LR	DT	RF	XGB	DNN	CNN	RNN	DNN	CNN	RNN
Pred.	1.009***	0.949***	0.989***	0.997***	1.016***	0.937***	0.996***	0.994***	0.983***	0.999***	1.035***	2.230***	1.003***	0.992***
	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.015)	(0.007)	(0.006)
Const.	−0.009	0.082***	0.017*	0.008	−0.019**	0.101***	0.008	0.015	0.019**	−0.002	−0.031***	−0.775***	0.135***	0.045***
	(−0.009)	(0.011)	(0.009)	(0.009)	(−0.009)	(0.012)	(0.009)	(0.009)	(0.010)	(−0.011)	(−0.009)	(−0.016)	(0.009)	(0.009)
N	2,565	2,565	2,565	2,565	2,565	2,565	2,565	2,565	2,565	2,565	2,565	2,049	2,049	2,049
Adj. R^2	93.64%	89.96%	94.07%	94.04%	93.68%	89.09%	94.04%	93.63%	93.13%	90.53%	94.13%	91.26%	92.06%	92.48%

Panel B: Event

	TF-IDF				Word2Vec (DL)				TF-IDF & AI/DL			BERT/LLM & AI/DL		
	1	2	3	4	5	6	7	8	9	10	11	12	13	14
	LR	DT	RF	XGB	LR	DT	RF	XGB	DNN	CNN	RNN	DNN	CNN	RNN
Pred.	1.001***	0.956***	0.999***	0.997***	1.009***	0.943***	1.003***	0.986***	0.972***	1.005***	0.994***	0.972***	1.005***	0.994***
	(0.005)	(0.006)	(0.005)	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.006)	(0.006)	(0.005)	(0.006)	(0.006)	(0.005)
Const.	0	0.065***	0.002	0.004	−-0.014	0.080***	−0.004	0.021**	0.015	0.012	0.015	0.015	0.012	0.015
	(0.009)	(0.011)	(0.008)	(0.009)	(−0.009)	(0.011)	(−0.008)	(0.009)	(0.010)	(0.011)	(0.009)	(0.010)	(0.011)	(0.009)
N	2,252	2,252	2,252	2,252	2,252	2,252	2,252	2,252	2,252	2,252	2,252	2,252	2,252	2,252
Adj. R^2	94.04%	91.04%	95.20%	94.91%	94.58%	91.34%	95.10%	94.69%	92.87%	92.22%	94.35%	92.87%	92.22%	94.35%