

Unmasking Climate Change Impacts: Traversing Storms, Cold, Heat and Fire in Corporate Earnings Calls through a Hybrid Taxonomy and GPT-based Methodology

Michele Cimino^{1,*†}, Annalisa Molino^{1†}, Maria Paola Priola^{2†}, Lorenzo Prospero^{1,*†} and Lea Zicchino^{1,3†}

¹Dept. of Financial Markets and Intermediaries Analysis, Prometeia. Bologna, Italy

²Dept. of Economics, University of Cagliari, Cagliari, Italy

³Johns Hopkins University SAIS Europe. Bologna, Italy

Abstract

Understanding how companies are managing the risks and opportunities of climate change is critical for investors, financial institutions and analysts. Corporate earnings calls are a valuable source of information and fill gaps in climate change data. We use transcripts of these calls in Europe and the US over the past two decades to assess how companies are affected by four climate hazards: storms, cold weather, heat waves and wildfires. Our approach involves several steps. First, we develop a classification system (taxonomy) for each climate hazard by reviewing scientific reports. This taxonomy is then expanded by identifying semantically similar words using a Word2Vec model. We then identify sentences in the transcripts that contain these climate-related keywords. Using generative AI techniques, specifically GPT 3.5, we analyse these sentences to gain insights into how individual companies are exposed to climate change risks. We distinguish between negative impacts (risks) and potential benefits (opportunities) for their business activities. We also identified three main channels through which climate risks affect different companies: 1) disruptions to the company's supply chain, 2) impacts on the company's demand, and 3) direct damage to the company's assets and operations. Our findings show that exposure to physical climate risk varies widely across sectors in terms of the types of events and the channels through which they affect firms. This innovative dataset has the potential to provide investors and analysts with accurate information on past climate risk exposures, thereby enhancing their understanding of how climate change may impact economic activity and corporate decision-making.

Keywords

Climate change, Physical risks, Text Analysis, Pattern Matching, Conference calls, Word Embedding, C58 Financial Econometrics, C63 Computational Techniques, G32 Financial Risk and Risk Management, Q54 Climate, Natural Disasters and Their Management, Global Warming CEUR-WS

1. Introduction

In an era of increasing environmental uncertainty, understanding the exposure of firms to physical climate-related financial risks has become imperative. Indeed, the increasing likelihood of acute hazards such as floods, storms, wildfires, heat waves, and cold waves presents a formidable challenge to corporations worldwide. These risks pose big challenges for companies, not just in terms of how they might affect the value of buildings and infrastructure or generate business interruption, but also because they can disrupt supply chains and affect the demand for products. Capturing the full spectrum of physical climate-related financial risks demands a large

amount of data that needs to be properly analyzed. This task in particular, when applied to the diversified nature of large corporations spanning various geographical regions, increases the complexity of assessing their vulnerability to climate-related hazards.

Our objective is to utilize NLP techniques to extract information from earnings call transcripts, thereby discerning firm-level indicators of physical risk exposure. In particular, the indicators that we extract offer information on firms' exposure in 3 different dimensions. Firstly, the exposure metrics are hazard-specific; the hazards considered in this paper are the following: wildfires, floods, hurricanes, cold waves, heat waves, and droughts. Secondly, we aim to differentiate between the adverse effects of these hazards on business activity, identifying whether they pose risks or present opportunities for firms. Thirdly, we seek to delineate the channels through which these hazards exert their influence, distinguishing between direct impacts, such as the destruction of firms' physical infrastructure, and indirect effects, such as disruptions to supply chains or reductions in consumer demand.

Our paper mainly contributes to the strands of financial literature using Natural Language Processing and

Ital-IA 2024: 4th National Conference on Artificial Intelligence, organized by CINI, May 29-30, 2024, Naples, Italy

* Corresponding author.

† These authors contributed equally.

✉ michele.cimino@prometeia.com (M. Cimino);

annalisa.molino@prometeia.com (A. Molino);

mariap.priola@unica.it (M.P. Priola);

lorenzo.prosperi@prometeia.com (L. Prospero);

lea.zicchino@organization.it (L. Zicchino)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

text analysis to measure firm-level exposure to climate risk. From early applications to more recent advancements, studies utilizing text analysis have provided insights into market perceptions, sector vulnerabilities, and regulatory developments related to climate change. By leveraging textual data from diverse sources including news articles ([1]; [2]; [3]; [4]), regulatory documents ([5], [6]) and firm-specific documents as annual reports, 10-Ks and earnings conference calls ([7], [8], [9], [6]) previous studies offered granular insights into the multifaceted nature of climate risk, possibly enabling targeted risk assessments and mitigation strategies. While some of the previous studies only focus on transition risk, others delved deeper into sub-categories of transition and physical risk. We contribute to the line of literature measuring climate risk exposure at the firm level by developing a new methodology to extract physical risk indicators, separating physical risk into several climate hazards, between risk and business opportunities and between demand and supply-related shocks.

Our findings highlight that the effects of climate hazards on firms are not uniform, with variations observed across sectors, time frames, and types of hazards. Moreover, we find a significant portion of firms that exhibit positive exposure to these hazards and a majority of firms indirectly impacted, often through shocks in their end markets or disruptions in their supply chains.

2. Earnings call data

Earnings conference calls serve as a conduit for corporations to relay information and positive guidance to stakeholders and interested parties on the company's financial results. As highlighted by Sautner et al. (2023) [7], earning calls are instrumental for financial analysts and market participants in garnering insights and engaging directly with corporate management. The dynamics of this interaction could significantly mitigate the susceptibility of these communications to corporate greenwashing in contrast to other climate risk disclosure tools [9]. Crucially, the earning calls transcripts typically include not only the presentation of the results by company officials but also a section of questions and answers where officials take direct questions from other participants of the call.

Expanding upon the research of Sautner et al. (2023) [7] we harness publicly available transcripts of corporate earnings calls to extract signals of firms' exposure to physical climate risks.

We collected data from Refinitiv Eikon database for 3'152 publicly-listed firms, spanning two decades from 2002 to 2022, yielding a total of 101'069 transcripts of English-language conference calls. These transcripts were sourced from the comprehensive set of pub-

licly listed companies catalogued in the Refinitiv Eikon database. Our dataset includes companies from 17 sectors, and with headquarters in 34 countries, although we report a concentration of 68% in North America and 23% in European countries.

3. Extracting Physical Risk Metrics

3.1. Taxonomy definition

We develop a comprehensive taxonomy of keywords, categorized into four main hazard groups: heatwaves and droughts, wildfires, storms and floods, and cold waves. We derive a preliminary compilation of text snippets (unigrams, bigrams or trigrams¹) from a list of glossaries, scientific reports, and documents issued by institutional sources (IPCC, NOAA etc.). Furthermore, we introduce synonyms sourced from WordNet and BabelNet, two lexical databases that link semantic word relationships. These resources support the augmentation of our taxonomy, not solely with hazard synonyms but also with additional hazards that may have been initially overlooked.

To construct a measure of physical risk exposure, we initially employ a Boolean model based on keyword matching to pinpoint paragraphs related to physical risk within the transcript text. We define two approaches: the Single Key Matching (SKM) and Neighbourhood Matching (NHM). SKM is designed to identify precise matches for specific keys within the text. On the other hand, NHM employs the proximity search of a word pair to encompass patterns located in close proximity to a specified key, thus facilitating a context-sensitive retrieval of information. NHM is a stricter model that employs a set of "control" keywords that - when located in the proximity of a taxonomy keyword, are used to validate the first match. "control" keywords include all taxonomy keywords and more generic physical risk keywords.

Next, we manually annotate a corpus of paragraphs featuring each text snippet from the taxonomy. We label paragraphs as True Positive (TP), if the occurrence of a keyword is indicative of the associated hazard group's occurrence, and False Positive (FP) otherwise. We compute the Precision, defined as the proportion of actual positives (true conditions) that are correctly identified by the model over the total predicted positives [10]. We proceed with the following scheme: Keywords with a Precision below 50% are omitted from our taxonomy. Keywords with a Precision above 90% are included in the final taxonomy as leading words. Keywords with a Precision between 50% and 90%, labelled as ambiguous words, are only included if they reach a Precision over

¹N-grams are sequences of n consecutive words (or tokens) extracted from a text. Unigrams are a sequence of one word, bigrams of two consecutive words and so on.

	No. Keywords	No. paragraphs	No. transcripts	No. Firms
Initial taxonomy	51	23301	10488	1771
Word2vec expansion*	25	3939	1821	115
Final taxonomy	68	27240	12309	1886

Note: * paragraphs, transcripts and firms not already included in the matched sample

Table 1
Boolean model matching statistics

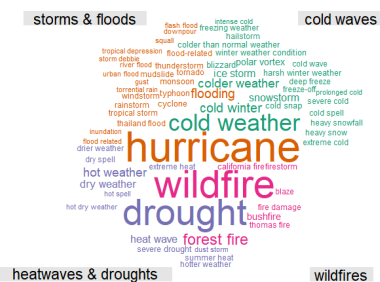


Figure 1: Taxonomy in-sample frequency wordcloud

70% with the NHM Boolean model. This strategy enables the exclusion of overly ambiguous keywords that could introduce excessive noise into our exposure metric. Additionally, we drop keywords that do not occur within our documents. This refinement process ends with a taxonomy of 43 keywords, from our initial sample exceeding 80.

Finally, we enrich our taxonomy through the implementation of a Skip-Gram Word2Vec model [11], a neural network-based approach for learning vector representations of words, known as word embedding. We train the Word2Vec model on a trigram corpus generated from the matched transcripts. Once trained, the model is used to identify semantically related words for each leading keyword in the taxonomy with a similarity ratio threshold of 80%. By applying the Word2Vec model, we expand our taxonomy with an additional set of 25 text snippets. A snapshot of the full taxonomy is shown in Figure 1. Descriptive statistics of the matched paragraphs are shown in Table 1.

3.2. NLP signal extraction

We employ Generative AI models, specifically, GPT-3.5 deployed through Azure Azure Machine Learning studio [12], to analyze paragraphs identified by the Boolean

model (refer to Section 3.1) and extract structured indicators of corporate physical hazard exposure. Our goal is to complete three primary tasks:

1. **Validation of Climate Hazard Exposure.** This entailed confirming the precise occurrence of the hazard (True Positive) within the paragraphs and eliminating False Positives erroneously flagged by the Boolean model.
2. **Risk versus Opportunities.** In instances where the text references are physical hazards, we measure whether the event has an impact on the firm and its directionality. The Generative AI model is prompted to categorize a mention of a climate hazard as “risk” if the paragraphs provided evidence of potential harm to the company’s economic operations. Conversely, the mention could be classified as “opportunity” if the text clearly indicated that the hazard presented a business prospect for the company.
3. **Channels of Corporate Exposure.** we identify the channels through which the hazard impacted the corporation’s business. The model was prompted to differentiate between “direct impacts”, such as the destruction of corporate physical infrastructure or immediate disruptions to operations, and “indirect impacts”, where the hazard affected either the company’s supply chains, influencing the supply of intermediate goods for production, or the company’s end markets, affecting the demand of company’s services and goods.

To extract the aforementioned indicators, we craft unique prompts for each task. For our study, we devise three zero-shot prompts, each with explicit instructions detailing: a) the tasks to be executed, including descriptions of the indicators to be generated, b) the format of the desired output, which is a JSON document with labelled variables, c) the analytical context, encompassing a brief definition of an earnings conference call, the paragraphs subject to analysis, the keyword identifying the hazard type, accompanied by a succinct hazard group description and the NACE 1 prevalent economic sector of the company.

		Precision	Recall	F1	Balanced Accuracy
Task 1	Class: Exposure	0.91	0.86	0.88	0.84
Task 2	Class: 0	0.54	0.55	0.54	0.71
	Class: Risk	0.86	0.78	0.82	0.82
	Class: Opportunity	0.71	0.85	0.77	0.87
Task 3	Class: 0	0.72	0.79	0.75	0.74
	Class: Direct	0.85	0.66	0.74	0.81
	Class: Indirect	0.62	0.63	0.62	0.73

Table 2
Evaluation results for GPT tasks

To evaluate the model’s precision, we manually annotate a subset of paragraphs for each task and compare human and machine classification. Given that the sample distribution in task 1 predominantly favours True Positives (the Boolean model is applied solely to keywords with 90% and 80% Precision, as delineated in section 3.1), we augment our validated database with out-of-sample Negative paragraphs (without any mention of physical hazards) to balance our dataset, sourcing from the ambiguous text snippets dropped by the Boolean models.

The model’s performance results are in Table 2 for each task. Balanced accuracy approximates 83% for the initial task and 80-76% for the subsequent tasks. F1 scores for the first task equals 77%, 71% for the second task and 70% for task three.

Subsequently, we apply the model to the entire corpus of identified paragraphs, with each paragraph being categorized in terms of exposure, risk versus opportunity, and exposure channel. Only paragraphs verified by GPT in Task 1 were deemed valid from the initial set pinpointed by the Boolean model. These paragraph indicators were then aggregated at both the transcript and firm levels.

4. Analysis

In this section, we aim to provide an analysis of indicators derived at the firm level. Our initial focus is on evaluating whether the total number of exposed firms in our sample to a specific hazard tends to be higher in years coinciding with significant events. To accomplish this, we conduct a preliminary assessment as follows. After determining the individual exposure of each firm, we track the number of exposed firms across different hazards over time. Subsequently, we compile a list of major global events. For hurricanes, we reference NOAA data², selecting Katrina and Sandy as significant events with a category exceeding 5. Additionally, we include hurricanes Harvey, Irma, and Maria, which occurred in the same year (2017). For

²<https://www.nhc.noaa.gov/news/UpdatedCostliest.pdf>




Figure 2: Number of firms exposed by Risk Driver

other hazards, we source events from Wikipedia. The outcomes of this analysis are depicted in Figure 2. The data reveals an increase in the number of exposed firms to hurricanes following the years in which the considered hurricanes struck the US. A similar trend is observed for the other risks. These findings provide preliminary evidence of a consistent correlation between exposed firms and the incidence of major physical events.

An analysis of the distribution of firms exposed to physical hazards reveals distinct patterns. Initially, an assessment of sector-specific vulnerability (refer to Figure 3) indicates a pronounced disparity in hazard exposure among sectors. Flood-related events constitute the primary risk across all sectors, followed by cold waves as the secondary hazard. Heatwaves and wildfires present a more varied pattern of exposure. Specifically, the agricultural sector (NACE A) and water-related industries (NACE E) are comparatively more susceptible to heatwaves and droughts. In opposition, the construction industry faces a heightened risk from wildfires.

Furthermore, it is significant to note that a considerable fraction of firms are subject to multiple hazards (illustrated in Figure 6), with nearly 30% of all firms categorized under this multi-hazard exposure.




Figure 3: Sectoral distribution of exposures

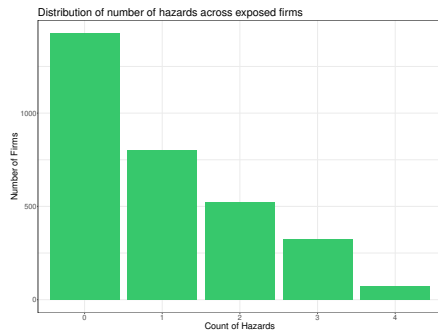


Figure 4: Number of firms with multiple exposures

The second analysis we conduct aims to examine the risk and opportunity patterns associated with each hazard across various sectors. As depicted in Table 3, a significant majority of firms present negative exposure to hazards. This phenomenon is especially marked in the case of floods and wildfires. Nonetheless, there exists a substantial proportion of firms that demonstrate positive exposure, especially after cold waves and heat waves. Looking at the diversification by economic sector could provide some intuitions in explaining this tendency in our data. Figure 5 portrays the distribution of risk and opportunity across sectors for different hazards. It becomes apparent that certain sectors face relatively lower exposure to specific hazards, whereas others might even reap advantages. For example, firms engaged in electric power generation (NACE D) are intuitively positioned to gain from cold waves, as demand for utilities tends to rapidly increase during such events.

Next, we investigated which sectors are impacted directly through damage to assets and operations, and which indirectly, disruptions to the supply chain or to the demand. It is clear from Table 3 that firms are more likely to be indirectly rather than directly exposed to physical hazards. While there is more than half of the firms which is indirectly exposed to the hazard, there

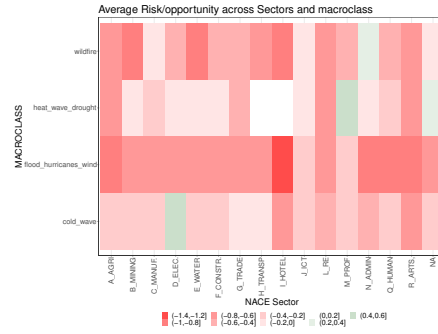


Figure 5: Average Risk/Opportunity indicator across sectors and hazard classes

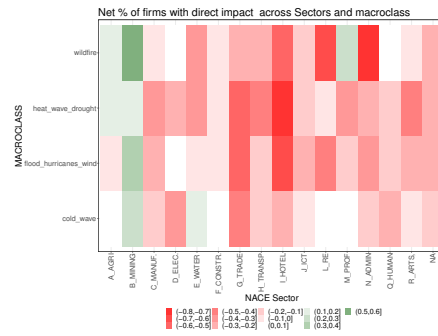


Figure 6: Average channel indicator across sectors and hazard classes

is also a proportion of firms that are indirectly affected, especially through cold waves and floods. The results of the analysis at the sectoral level are shown in Figure 6. The mining sector stands out because it is directly exposed to all risks, a result somehow expected. The agricultural sector tends more often to be exposed directly rather than indirectly, while among the most indirectly exposed sectors, we find the trade and accommodation services.

5. Conclusion

Our study reveals a multifaceted landscape of climate risk exposure among firms. The application of NLP techniques to earnings call transcripts allows the identification of firm-specific, hazard-related exposure metrics that vary significantly across different dimensions. Our findings indicate that while physical climate hazards generally pose risks to business activities, there are instances where they can present opportunities, particularly for firms that benefit from the aftermath of such events, and most notably, the highest share of physical risk exposures for firms come from an indirect source, being either a shock in demand caused by hazard impacts in end mar-

Hazard class	Opportunity	Risk	No exposure	Direct	Indirect	No Channel
cold_wave	24.31	49.94	25.74	22.44	40.59	36.96
flood_hurricanes_wind	13.29	71.26	15.45	20.94	45.49	33.57
heat_wave_drought	27.27	45.11	27.62	11.84	48.20	39.97
wildfire	18.60	57.44	23.97	15.29	28.10	56.61
Total	18.39	66.42	15.20	21.52	51.39	27.09

Table 3
Distribution of Exposure across Hazards, in percentages

kets, or disruption of value and supply chains in the firms' network.

The taxonomy developed has proven effective in discerning the nuanced impacts of climate hazards. This has facilitated the deployment of generative language models, and the consequent granular understanding of the direct and indirect channels through which these hazards affect firms, highlighting the complex interdependencies within supply chains and consumer markets. Our research underscores the importance of considering these multifaceted effects when assessing climate-related financial risks.

Looking forward, the insights garnered from this study could serve as a valuable resource for navigating the evolving landscape of climate risks. Investors and financial institutions could leverage this research to enhance their understanding of how climate change may impact economic activity and corporate decision-making. Ultimately, our work contributes to the broader discourse on sustainable finance, emphasizing the need for innovative approaches to understand and mitigate the financial repercussions of climate change. Additional work is still required to leverage our newly created database and investigate more deeply the intricate impacts of physical climate risk on corporations.

Acknowledgments

We would like to thank Jakob Jones and Michele Tursi for their support in the construction of the Earnings call database. We would also thank Michele Filannino for his guidance on GPT deployments and perceptive feedback, and Professor Giovanna Bua and Professor Daniel Kapp for their availability to share data and provide valuable insights. We also thank participants at the 2023 International Fintech Research Conference in Naples.

References

- [1] R. F. Engle, S. Giglio, H. Lee, B. T. Kelly, J. Stroebel, Hedging climate change news, Yale ICF Working Paper 2019 (2019).
- [2] K. B. K. I. David Ardia, Keven Bluteau, Climate change concerns and the performance of green vs. brown stocks, *Management Science* 69 (2023) 7607–7632.
- [3] M. R. . S. G. Faccini, R., Dissecting climate risks: Are they reflected in stock prices? (2023).
- [4] G. Bua, D. Kapp, F. Ramella, L. Rognone, Transition versus physical climate risk pricing in european financial markets: A text-based approach (2022).
- [5] H. Berkman, J. Jona, N. S. Soderstrom, Firm-specific climate risk and market valuation (2024).
- [6] J. Koelbel, M. Leippold, J. Rillaerts, Q. Wang, Ask bert: How regulatory disclosure of transition and physical climate risks affects the cds term structure, Swiss Finance Institute Research Paper, *Journal of Financial Econometrics*, forthcoming 21 (2020).
- [7] Z. Sautner, L. Van Lent, G. Vilkov, R. Zhang, Firm-level climate change exposure, *The Journal of Finance* 78 (2023) 1449–1498.
- [8] Q. Li, H. Shan, Y. Tang, V. Yao, Corporate climate risk: Measurements and responses, *Review of Financial Studies*, Forthcoming (2020).
- [9] J. A. Bingler, M. Kraus, M. Leippold, N. Webersinke, Cheap talk and cherry-picking: What climatebert has to say on corporate climate risk disclosures, *Finance Research Letters* 47 (2022) 102776.
- [10] S. Raschka, An overview of general performance metrics of binary classifier systems, arXiv preprint arXiv:1410.5330 (2014).
- [11] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, *Advances in neural information processing systems* 26 (2013).
- [12] Azure-AI-Services, Openai (gpt-35-turbo - version 2023-05-15) [large language model] (2023). URL: <https://learn.microsoft.com/en-us/azure/ai-services/openai>.