

## Firm-Level Climate Change Exposure

ZACHARIAS SAUTNER, LAURENCE VAN LENT, GRIGORY VILKOV,  
and RUISHEN ZHANG\*

### ABSTRACT

We develop a method that identifies the attention paid by earnings call participants to firms' climate change exposures. The method adapts a machine learning keyword discovery algorithm and captures exposures related to opportunity, physical, and regulatory shocks associated with climate change. The measures are available for more than 10,000 firms from 34 countries between 2002 and 2020. We show that the measures are useful in predicting important real outcomes related to the net-zero transition, in particular, job creation in disruptive green technologies and green patenting, and that they contain information that is priced in options and equity markets.

CLIMATE CHANGE WILL PROFOUNDLY AFFECT the way business is conducted. Scientists have developed complex models that estimate the effect of greenhouse gas emissions on the global climate. At the same time, however, little

\*Zacharias Sautner, Laurence van Lent, and Grigory Vilkov are at Frankfurt School of Finance & Management. Ruishen Zhang is at Shanghai University of Finance and Economics. The climate change exposure data are available at <https://doi.org/10.17605/OSF.IO/FD6JQ>. We thank Stefan Nagel, an anonymous Associate Editor, two referees, Artur Hugon, Marcin Kacperczyk, Bryan Kelly, Kelvin Law, Christian Leuz, Tim Loughran, Quentin Moreau, and Ane Tamayo for helpful comments. We are grateful to Aakash Kalyani for preparing the green jobs data set and to Markus Schwedeler for sharing his code. Participants at AFA 2022 Meetings, NBER, NYU Stern (PhD Classes in Empirical Household Finance), University of Zurich, University of Dusseldorf, Stockholm Business School, SUFE, Bocconi University, Columbia University, CEIBS, LSE, Xiamen University, Duke Kunshan University, University of Miami, Tingshua University, Dongbei University of Finance and Economics, Chinese University of Hong Kong, and University of Hong Kong provided helpful feedback. Funding for this project was provided by DFG Project ID 403041268 - TRR 266 (van Lent and Zhang), the Institute for New Economic Thinking (Van Lent), Shanghai Pujiang Program (Zhang), the 111 Project (B18033)(Zhang), and National Natural Science Foundation of China (Project ID 72202128, 72241405, Zhang). Parts of this project were conducted when Sautner was visiting NOVA School of Business and Economics and UCSD Rady School of Management (funded in part by the Norwegian Finance Initiative). We have read *The Journal of Finance* disclosure policy. Sautner is a Regular Research Visitor at the ECB and a member of the Sustainability Council of Lampe Asset Management (since November 2022). He has no other disclosure to make. Van Lent, Vilkov, and Zhang have no conflicts of interest to disclose.

Correspondence: Zacharias Sautner, Frankfurt School of Finance & Management, Adickesallee 32-34, 60322 Frankfurt am Main, Germany; e-mail: [z.sautner@fs.de](mailto:z.sautner@fs.de).

This is an open access article under the terms of the [Creative Commons Attribution](#) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

DOI: 10.1111/jofi.13219

© 2023 The Authors. *The Journal of Finance* published by Wiley Periodicals LLC on behalf of American Finance Association.

evidence exists on the degree to which climate change impacts jobs, innovation, and risk sharing in capital markets. One key challenge in estimating these impacts is that it is difficult to measure how individual firms are affected by climate change (Giglio, Kelly, and Stroebe (2021)), as the effects are multifaceted, originating from multiple sources. For instance, while physical climate changes and regulations implemented to combat global warming can impose costs on some firms, climate change can provide opportunities for other firms, such as those operating in renewable energy, electric cars, or energy storage. It is therefore important to develop disaggregated measures that capture this variation across firms. The measures should also reflect market participants' assessments about how climate change affects individual firms. Such information is important to consider in a finance context given the critical role that market participants play in the resource allocation and price discovery process.

In this paper, we make progress on this front by using transcripts of earnings conference calls to construct time-varying measures of how call participants across the globe view firms' exposures to different facets of climate change. Earnings calls are key corporate events in which financial analysts listen to management and ask questions about current and future developments material to the firm (Hollander, Pronk, and Roelofsen (2010)). We interpret these measures as capturing the attention financial analysts and management devote to climate change topics at a given point in time. A benefit of these measures is that they reflect "soft" information originating from information exchanges between managers and analysts.<sup>1</sup>

To construct the climate change exposure measures, we build on recent work using quarterly earnings calls as a source for identifying firms' various risks and opportunities (Hassan et al. (2019, 2021, 2023a, 2023b), Jamilov, Rey, and Tahoun (2021)). These studies use the proportion of the conversation during an earnings call that relates to a particular topic to capture the firm's exposure to that topic. We follow these papers in defining "exposure" to an issue as the share of the conversation in a transcript devoted to that topic.<sup>2</sup> We depart from these papers, however, along two dimensions. First, our measures capture the market's perception of a firm's exposure to various upside or downside factors related to climate change, namely, physical threats, regulatory interventions, and technological opportunities. Second, to mitigate the

<sup>1</sup> This feature allows us to provide economic insights beyond those derived from existing firm-level exposure measures based on "hard" information (e.g., carbon emissions, extreme local weather events). Note that the exchanges are not limited to soft information but might also discuss specific quantitative data or restate "hard" information in conversational terms. Prior literature provides important insights into the relations between "hard" information and real and financial outcomes at the firm level (e.g., Bolton and Kacperczyk (2021, 2022, 2023), Ilhan, Sautner, and Vilkov (2021) or De Haas and Popov (2023) for carbon emissions, and Kruttli, Roth Tran, and Watugala (2021), Hong, Li, and Xu (2019), Addoum, Ng, and Ortiz-Bobea (2020), or Pankratz and Schiller (2021) for weather events).

<sup>2</sup> This definition of "exposure" is different from how risk exposure is defined in the asset-pricing literature. Our measure is not intended to capture the covariance with aggregate fluctuations. Hassan et al. (2019) discuss the relationship between these two areas of literature.

challenges of identifying “niche languages” that use specific wordings, particularly in the context of climate change, where language use varies among policymakers, journalists, and financial market participants (Webersinke et al. (2021)), we develop a new method that adapts the keyword discovery algorithm proposed in King, Lam, and Roberts (2017) to construct four related sets of climate change bigrams in earnings calls. The first captures broadly defined aspects of climate change. The remaining three measures cover specific climate change “topics:” opportunities, physical shocks (e.g., sea level rise), and regulatory shocks (e.g., carbon taxes, cap and trade markets). We then use these four sets of bigrams to construct firm-level measures reflecting call participants’ topical attention. In particular, the measures count the frequency of specific climate change bigrams in a transcript, scaled by the number of bigrams.<sup>3</sup> The algorithm only requires human input to specify a short list of initial keywords associated with climate change. Our sample covers data from over 10,000 firms in 34 countries between 2002 and 2020.

We conduct several validation exercises to verify our methodology. First, we consider the face validity of the climate change bigrams. Second, we follow Baker, Bloom, and Davis (2016) and perform a structured human audit in which 18 graduate students independently coded over 2,000 transcript text fragments. Both of these exercises suggest that our algorithm reliably captures bigrams identifying climate change discussions. Third, our exposure measures are robust to excluding one keyword at a time from the initial keywords list. Fourth, our keyword search-based measures substantially improve the identification of climate change discussions relative to an alternative approach using the initial keywords only. And fifth, we find plausible industry patterns in the exposure measures. When we aggregate exposure to the industry level, the sector with the highest overall exposure is Electric, Gas, & Sanitary Services (utilities), followed by Construction (top-ranked firms build power generation systems or solar projects) and Transportation Equipment (top-ranked firms build fuel-cell or zero-emission vehicles). Utilities top the exposure ranking for opportunity and regulatory shocks, which indicates that this sector faces both opportunities (e.g., renewable energy) and regulatory risks (e.g., carbon taxes).<sup>4</sup>

Our results reveal sizeable within-industry variation for all measures, which indicates that firms benefit or suffer from climate change to various degrees. A case in point is the comparison between TotalEnergies and ExxonMobil. While

<sup>3</sup> We also construct “sentiment” measures, which count the relative frequency of climate change bigrams that occur in the vicinity of positive and negative tone words (Loughran and McDonald (2011)), and “risk” measures, which count the relative frequency of climate change bigrams mentioned in the same sentence as the words “risk,” “uncertainty,” or their synonyms.

<sup>4</sup> That firms with heightened regulatory risks also exhibit climate-related opportunities is consistent with Cohen, Gurun, and Nguyen (2021), who document that several major electricity, oil, and gas firms are not only large CO<sub>2</sub> emitters, but also innovators in green technologies. This finding is consistent with how analysts view sectors with high regulatory risks (e.g., “Morgan Stanley: ‘Second wave of renewables’ to drive 70 GW of coal retirements,” *S&P Global Market Intelligence*, December 20, 2019).

TotalEnergies and ExxonMobil have similar regulatory exposures, TotalEnergies scores more than seven times higher in terms of measured opportunities. This divergence in perceived prospects is consistent with differences in the perceived extent to which these firms embrace renewable energy and the net-zero transition into their business models (Pickl (2019)).

In a final validity check, we find that climate exposure positively correlates with carbon emissions and Engle et al.'s (2020, EGKLS) index of public climate change attention. The association with emissions stems from regulatory and opportunity exposure (since physical exposure is unrelated to emissions).<sup>5</sup> The effect of public attention also arises from positive associations between EGKLS's index and the opportunity and regulatory exposure measures.

We apply our measures to shed light on the nature of climate change exposure among our sample firms. Perhaps surprisingly, as climate change is often seen as an aggregate risk factor associated with global changes in the physical climate, its within-sector impact is far from uniform. A variance analysis that separates the relative contributions of aggregate, sectoral, and firm-level exposure by including the corresponding sets of fixed effects shows that between 70% and 96% of the variation in the exposure measures plays out at the firm level. Only half of this firm-level variation is persistent, suggesting that firms within an industry are exposed to climate change to varying degrees over time. Thus, the effects of climate change are heterogeneous across firms even within an industry. This result is consistent with the idea that many factors that affect a firm's ability to adapt to a greener economy exhibit large firm-level components (e.g., managerial skill, financing constraints).

We interpret the large share of firm-level variance as capturing economically meaningful heterogeneity and argue that a firm's idiosyncratic climate change exposure is the key driver of this heterogeneity. That being said, a plausible alternative is that part of the variation reflects idiosyncratic measurement error. Several tests dispel this alternative for several reasons. First, as discussed below, we report robust associations between our measures and green job creation, green innovation, and risk-related outcomes. Second, following Hassan et al. (2019), we directly quantify the amount of measurement error contained in the firm-level variation. Approximately, 5% to 10% of the variation in measured exposure can be attributed to measurement error. The implied measurement error at the firm level is about 2 percentage points higher than that for the overall variation. Although we interpret these results with due caution, they suggest that measurement error in the firm-level dimension is higher than that in the overall panel, but only modestly so.

Having bolstered confidence that the firm-level variation in measured climate change exposure is meaningful, we apply it to four real and financial market outcomes. In the first two applications, we demonstrate that climate

<sup>5</sup> This result may also reflect the fact that some firms' emissions provides opportunities by supporting the transition to a greener economy (e.g., producers of building materials that make houses more energy-efficient). Such "enabling activities" are also explicitly included in the EU Taxonomy, which identifies activities that help reach the EU's climate targets.

change exposure predicts green-tech hiring and green patents, two key drivers of the low-carbon transition. Using data compiled from Burning Glass (BG) by Bloom et al. (2021), we establish that firms with higher measured climate exposures create more jobs in disruptive green technologies over the subsequent year:<sup>6</sup> a one-standard-deviation increase in climate change exposure is associated with a 109% increase in green jobs in the following year. This overall effect originates from more job creation at firms exhibiting higher measured opportunities and regulatory exposures.

The results for green-tech job creation extend to green patenting. A one-standard-deviation increase in climate change exposure is associated with a 72% increase in the number of green patents in the following year. Once more, this finding stems from firms with higher opportunities and regulatory exposures. High-exposure firms are not simply recruiting more across fields. They are also not more innovative, in general. In fact, firms with higher exposure hire less in nongreen-tech areas and generate fewer nongreen patents.

The remaining two applications relate climate change exposure to financial market outcomes. We first show that measured exposure is related to risks and risk premiums in the options market. Such relationships are plausible, as policy uncertainty surrounding regulation, including climate policy uncertainty, is priced in options (Kelly, Pastor, and Veronesi (2016), Ilhan, Sautner, and Vilkov (2021)). Likewise, there is plenty of uncertainty surrounding green technology or renewable energy investment. Realizing these opportunities leads to significant gains if successful or large losses if unsuccessful. It is therefore plausible that measured exposure relates to investors' propensity to hedge extreme climate risks and/or gamble on climate outcomes. Indeed, for options written on stocks with high overall exposure, the tail regions are relatively more expensive. Effects are similar at firms with high opportunity exposure, for which investors are willing to pay a (variance risk) premium. In comparison, effects are smaller but still statistically significant for firms with high regulatory exposure. This finding corroborates the view that some firms with high regulatory exposure face downside risks and upside potential (due to their innovation activity).

We also document the conditional pricing of a factor that reflects innovations to the aggregate level of climate change exposure. Firms with higher betas to this factor face higher uncertainty related to future developments in climate-related areas and, as a result, earn higher returns.<sup>7</sup> Our estimation applies the approach of Gagliardini, Ossola, and Scaillet (2016), which performs well when—as in our case—the cross section is large relative to the time series. We

<sup>6</sup> Our data do not cover all jobs potentially related to climate change, but they do identify job postings with potential to have a lasting and meaningful real impact, as Bloom et al. (2021) only consider job creation in “disruptive” technologies (e.g., solar or battery technology).

<sup>7</sup> Our primary objective is to show that climate attention in earnings calls is linked to systematic risk, with shocks to such attention being priced in the cross section. We do not want to propose a new factor to be added to the factor zoo, and we do not try to use a conditional model framework to explain asset pricing anomalies (Lewellen and Nagel (2006)).

obtain a positive average conditional risk premium on the factor, and, more importantly, find large time-series variability in the risk premium.<sup>8</sup>

Our keyword discovery approach of extracting climate-related information from text offers an alternative approach to contemporaneous papers that try to accomplish the same task by relying on other advances in natural language processing (NLP). All of this work, including ours, is based on the understanding that standard NLP methods are not well suited for “niche languages,” that is, specialized, highly technical vocabulary that varies substantially across textual sources (Varini et al. (2020), Webersinke et al. (2021)). These frictions are exacerbated when the wordings associated with a topic are complex, ambiguous, and fast moving. A promising approach among these alternatives is to use pretrained language models to learn word patterns in the language. When implementing this pretraining approach in a specific domain of interest (e.g., climate change), rather than using large generic corpora, researchers have found some promising results (Bingler, Kraus, and Leippold (2022), Kölbel et al. (2022)). Work is ongoing on these problems. Which approach works best in the context of climate finance is ultimately an empirical matter.

A valid question is whether our approach delivers meaningful gains above and beyond an alternative, off-the-shelves approach. Our main argument is that keyword discovery is useful when the language of interest is not common. We illustrate this claim by constructing, for comparison purposes, alternative exposure measures using a list of pre-specified keywords from EGKLS. These keywords appear more frequently in earnings calls than the bigrams we identify, probably because EGKLS’s set also contains unigrams and more general terms. However, several of EGKLS’s unigrams are part of our top-100 list of bigrams, and exposure measures based on the pre-specified keywords correlate positively with our measures. Beyond these correlations, a question is why the approaches differ. As mentioned above, our measures have the benefit of capturing context-specific jargon used in specialized economic environments (earnings calls), while an approach using pre-specified keywords better captures broader discussions (e.g., in news media in the case of EGKLS’s keywords). In addition, our approach adjusts the vocabulary over time, while using pre-specified keywords fixes this vocabulary *ex ante*.<sup>9</sup> Finally, especially for the topic-based measures, it is easier to identify initial seed bigrams than to develop keyword lists from authoritative texts.

Most closely related to our paper is the contemporaneous work by Li et al. (2021, LSTY), who also use earnings calls to identify climate risks. We diverge

<sup>8</sup> A caveat of all four applications is that any evidence of our measures’ ability to predict real and financial outcomes is a success only if the true relationship exists in the data. We therefore face the usual joint-hypothesis problem between the quality of our measures and the true economic model generating the data.

<sup>9</sup> Time-series variation in true (unobservable) climate change exposure, especially over long horizons, is more likely to be picked up by such an “evolutionary” approach. Indeed, the selection of pre-specified keywords may become obsolete over time with changing technologies or climate change concerns.



from their work in terms of our method, focus, and sample. More specifically, LSTY use a pre-specified training library to identify climate risk words, which, we argue, is unlikely to uncover the exact language used in earnings calls to discuss climate change (see also Varini et al. (2020)). In addition, while LSTY focus on physical and regulatory risks among U.S. firms, we provide a more comprehensive analysis based on a global sample and include upside opportunity effects of climate change. Based on a textual analysis of 10K reports, Baz et al. (2022) document that firms with more regulatory climate change exposure experience positive stock return effects after the 2016 Trump election.

Since making our data available, our measures have been related to a series of real and financial outcomes. This “out-of-sample” evidence is reassuring, as it indicates that the measures capture meaningful variation across firms and do not reflect mostly noise. On the real side, as in our paper, von Schickfus (2021) illustrates more green patenting when the overall measure and the opportunity measure are higher, and Li, Lin, and Lin (2022) show that the overall measure predicts depressed overall innovation. Furthermore, our overall measure positively relates to cash holdings (Heo (2021)) and explains how strongly U.S. firms’ emissions declined in response to the EPA’s 2010 Greenhouse Gas Reporting Program (Tomar (2023)). Our physical measure is related to physical risk disclosure in 8K filings (Gostlow (2021)), and the opportunity measure relates to firms’ carbon risk management (Duong et al. (2021)). On the financial side, our physical measure is associated with lower leverage after the Paris Agreement (Ginglinger and Moreau (2022)). Mueller and Sfrappini (2022) show that after regulatory climate risks become salient, bank lending is skewed toward firms with high regulatory exposure in the United States, but away from such firms in the EU. We provide additional evidence in Sautner et al. (2022) that our measures are priced in equity markets, and Kölbel et al. (2022) show that the overall measure is negatively associated with credit default swap (CDS) spreads after the Paris Agreement. Di Giuli et al. (2022) find that investors’ propensity to vote for climate proposals after experiencing hot temperatures is higher at firms with more overall climate change exposure. Heath et al. (2022) find that socially responsible investment (SRI) funds invest less in firms with higher overall climate change exposure. Our keyword dictionary is used by Hail, Kim, and Zhang (2021).

The rest of the paper proceeds as follows. Section I describes the data. Section II presents our method to quantify firm-level climate change exposure. Section III validates the exposure measures. Section IV presents a variance decomposition of the exposure measures and addresses measurement error. Section V presents four applications of the exposure measures. Section VI concludes.

## I. Data

### A. Data on Earnings Conference Calls

We use transcripts of quarterly earnings calls held by publicly listed firms to construct time-varying measures of the attention paid by call participants to firm-level climate change exposure. The measures are constructed using

the entire earnings call, including both the management presentation and the Q&A session with analysts.<sup>10</sup> The transcripts are collected from the Refinitiv Eikon database. We use the complete set of English-language transcripts from 2002 to 2020. Unless indicated otherwise, as most of our other data vary at the year level, we average quarterly transcript-based measures for each firm. We exclude countries with 150 or fewer firm-year observations and drop SIC codes 9900 to 9999 (“Nonclassifiable”) Our final sample includes 86,152 firm-year observations from 10,673 firms headquartered in 34 countries. Variable definitions are presented in Table A.1.<sup>11</sup>

### B. Data on Carbon Emissions

Some tests use data on carbon emissions (*Total Emissions*), calculated as the sum of Scope 1 and Scope 2 emissions, from S&P Global Trucost. These data include emissions reported by firms and emissions estimated by Trucost. We use emission levels, rather than intensities, as emission levels are associated with a risk premium (Bolton and Kacperczyk (2021, 2023)), are the prime target of policy and investor initiatives aiming to achieve net-zero emissions, and are directly linked to carbon budgets (Bolton, Kacperczyk, and Samama (2021)). Furthermore, many firms have witnessed strong investor opposition on reporting emission intensities. To link the emissions data with our sample firms, we apply a series of matching variables based on the following order: (i) GVKEYs, (ii) ISINs, (iii) exact names, (iv) fuzzy names, and (v) tickers plus the first two ISIN digits. We can match 33,789 firm-years with the emissions data (4,999 unique firms from 34 countries between 2004 and 2020).<sup>12</sup>

### C. Data on Public Attention to Climate Change

We borrow an index developed by EGKLS to capture how public climate change attention varies in the time series. The *WSJ CC News Index* is constructed by measuring news about climate change in the *Wall Street Journal* (WSJ). To quantify the intensity of climate news coverage, EGKLS compare the WSJ’s news content to a corpus of authoritative texts on climate change. The resulting measure reflects the fraction of the WSJ dedicated to the topic of

<sup>10</sup> We also provide tests based on the measured exposure constructed from the Q&A session only. The Q&A part is less scripted and may be less subject to strategic disclosure incentives than the presentation part. In some calls, analysts ask no questions (we would calculate a climate change exposure of zero in these cases). However, zero-question calls are a nonrandom event, and treating these calls as if the firm is unexposed to climate change likely introduces bias (Chen, Hollander, and Law (2014)).

<sup>11</sup> Table IA.I in the Internet Appendix provides the distribution of firm-years across countries. The Internet Appendix may be found in the online version of this article.

<sup>12</sup> Table IA.II illustrates that Trucost data coverage is higher for firm-years with higher climate change exposure, larger, more profitable, and less-R&D-intense firms, and non-U.S. firms. The higher climate change exposure scores are expected given that Trucost caters to clients in need of climate risk data (especially risks related to emissions).



climate change each day (we use average annual values). For our sample, *WSJ CC News Index* is available from 2002 to 2017.

#### D. Data on Green-Tech Jobs

Job data related to important green technologies come from Bloom et al. (2021). These authors use textual analysis to identify 29 disruptive technologies over the past decades, of which four are broadly related to climate change (“hybrid vehicle electric car,” “lithium battery,” “solar power,” and “fracking”). Our data from Bloom et al. (2021) contain online job postings by firms related to these four technologies. We refer to the jobs related to these technologies as green-tech jobs.<sup>13</sup> The data do not cover all jobs potentially related to climate change, but do identify those green jobs that, by Bloom et al.’s (2021) construction, have a lasting and meaningful (“disruptive”) real impact. Bloom et al. (2021) obtain these data from BG, which aggregates online job postings using “spider bots” from job boards or employer websites.<sup>14</sup> We match these data by GVKEY and year. Jobs data are available for U.S. firms for 2007 and 2010 to 2020.

The measure *#Green-Tech Jobs* is the number of postings for disruptive green-tech jobs in a firm-year. We assume that no green-tech job was posted if a firm-year does not indicate disruptive green-tech job creation in the BG database. (The results are robust to only considering firm-years within the BG database; many firm-years in BG also show zero green-tech postings). Some tests use *#Nongreen-Tech Jobs*, the number of job postings related to nongreen disruptive technologies in a firm-year. We observe disruptive green job postings in 5.4% of firm-years, and conditional on *#Green-Tech Jobs* being nonzero, the average (median) number of green-tech jobs is 38 (3). The top-5 firms in the cumulative count of new green-tech jobs include Tesla, Sunrun, First Solar, Sunpower Corp, and Viviant Solar.

#### E. Data on Green Patents

To identify green patents, we collect patent data from the Google Patents (GP) database. This database is also used by Kogan et al. (2017) and Kelly et al. (2021). To identify “green” patents, we follow Cohen, Gurun, and Nguyen (2021) and apply an OECD classification that identifies patents with the potential to address environmental problems. A description of how the OECD classifies patents into technology groups is provided by Haščić and Migotto (2015). Green patents include patents on emission abatement technologies, renewable

<sup>13</sup> It is unclear ex ante whether fracking has positive or adverse environmental effects. More specifically, Acemoglu et al. (2019) argue that shale gas has the short-term benefit of lower emissions, when compared to conventional fossil fuels. However, the shale gas boom may lead to less innovation in other emission-reducing technologies in the long run. Furthermore, fracking has negative climate effects due to emission leakage. Our results are robust to excluding fracking jobs.

<sup>14</sup> BG data are also used by Darendeli, Law, and Shen (2022) to measure green hiring. Campello, Gao, and Xu (2021) additionally use BG data, though not in a climate context.

energy, and energy storage. As in Kogan et al. (2017), we use name matching to match patent assignee names to sample firms.<sup>15</sup> Patent data are available for U.S. firms from 2002 to 2019 (GP coverage for 2020 was still limited at the time of writing).

The measure *#Green Patents* is the number of green patents filed in a firm-year. We assume that no green patenting occurred if we are unable to identify a green patent in GP for a firm-year (results are robust to relaxing this assumption). Consistent with Acemoglu et al. (2019), new green patents are relatively rare—we observe green patenting in only 1.4% of firm-years. However, the distribution is highly skewed. If we consider observations within GP, then green patenting is observed in 15% of firm-years. Conditional on green patenting being nonzero, the average (median) number of green patents equals 8.5 (2). The top green patent producer is Caterpillar, with 1,364 green patents over the sample period.<sup>16</sup> We also use the total number of nongreen patents filed (*#Nongreen Patents*).

#### F. Data on Risks and Risk Premiums in the Options Market

Data on option-implied variables are from the Volatility Surface File of Ivy DB OptionMetrics. In these tests, we focus on S&P500 firms, for which data on liquid options are available. We match options data through the historical CUSIP link of OptionMetrics. We construct six measures: implied variance (*IVar*), implied skewness (*ISkew*), implied kurtosis (*IKurt*), implied volatility slopes (*SlopeD* and *SlopeU*), and variance risk premium (*VRP*). The variable construction process is detailed in Section II of the Internet Appendix. The high frequency of the option-implied measures allows us to use quarterly values of *CCExposure*.<sup>17</sup>

#### G. Data on Risk Premiums in the Equity Market

Our tests examining the climate change exposure factor use monthly data on the standard factors from Ken French's data library. Term and default spread data are from the St. Louis Fed's FRED library. The term spread is the difference between the 10-year and three-month Treasury constant maturity data series (variable *T10Y3MM*). The default spread is the difference between the

<sup>15</sup> We track the timing of an invention by matching patents using the priority year, that is, the effective date of a patent filing (De Haas and Popov (2022)). While the "filing date" corresponds to when a patent application is filed at the patent office, the "priority date" is when the novelty of an invention is established.

<sup>16</sup> Caterpillar traditionally manufactured diesel engines and mining equipment, but moved into selling photovoltaic or energy storage technology. The firm also ranks in the top-10 in Cohen, Gurun, and Nguyen's (2021) sample; the slight ranking divergence is due to different sample periods.

<sup>17</sup> To avoid look-ahead bias, we match quarterly exposure values covering earnings calls in quarter  $t$  (typically discusses events of quarter  $t - 1$ ) with option-implied measures from the last day of quarter  $t$ .

Baa and Aaa corporate bond yield (*BAA10YM* and *AAA10YM*). Book-to-market ratio data (defined in log terms as in Fama and French (2008)) come from Compustat North America. Term and default spreads and the book-to-market ratio for each firm are centered and standardized in the time series, and then used as instruments for conditional risk premium estimation. We restrict the risk premium tests to S&P500 firms with more than 28 monthly returns (out of 228) during our sample period.

#### *H. Financial Statement Data*

Data on firm financial variables (e.g., total assets, debt, CAPEX, R&D, or cash holdings) are from Compustat North America and Compustat Global.

## **II. Quantifying Firm-Level Exposure to Climate Change**

### *A. Discovery of Climate Change Bigrams*

To quantify exposure to climate change, we build on Hassan et al. (2019, 2021, 2023a). Extracting climate-related information from text sources is challenging (Webersinke et al. (2021)). Methods using training libraries or pre-specified word lists do not cope well with the niche language used to describe climate change.<sup>18</sup> In addition, discussion in earnings calls considers climate change together with topics such as regulation, tax credits, technological breakthroughs, and performance. This results in substantial ambiguity about when the discussion is genuinely about climate change. Finally, vocabulary used to discuss climate change is fast moving, changing to reflect shifting opinions, regulations, and innovations related to climate change.

To address these challenges, we adapt the keyword discovery algorithm proposed in King, Lam, and Roberts (2017).<sup>19</sup> This algorithm does not require a comprehensive “climate change” training library, but rather only a small set of “initial” bigrams (see Table IA.III). These initial bigrams are chosen because they relate unambiguously to climate change. The algorithm then uses these initial bigrams to search for new bigrams that also likely indicate climate change conversation and searches directly in the transcripts. Because each initial bigram is connected to a specific group of new bigrams discovered through the search algorithm, one can easily decompose the measure of climate change exposure into its constituent parts based on the presence of these bigrams. The initial bigrams allow the algorithm to identify sentences that focus unambiguously on climate change. The algorithm then extracts “features” by relying on supervised learning methods. Features are bigrams beyond the initial set predicting climate change from the identified sentences.

<sup>18</sup> That said, researchers have used the SEC Climate Disclosure Search tool, which looks for pre-specified keywords in SEC filings, to develop a measure of climate risk (Berkman, Jona, and Soderstrom (2019)).

<sup>19</sup> Details, including how we define the set of initial bigrams, are presented in Section I of the Internet Appendix.

Finally, the algorithm constructs a model predicting whether a sentence is related to climate change. We apply this prediction model to sentences that do not include any initial bigrams and then assess whether the predicted sentences are related to climate change. To discover new climate change bigrams, we reverse-engineer the machine-learning (ML) process and trace back the bigrams that best discriminate climate-change-related sentences from other sentences. The resulting set of climate change bigrams  $\mathbb{C}$  includes the initial bigrams and the newly identified bigrams.

That our approach generates meaningful climate change bigrams based on the initial bigrams is helpful for several reasons. First, it extends the rather broadly specified initial bigrams into more specialized word combinations.<sup>20</sup> Second,  $\mathbb{C}$  includes the names of several power stations and wind farms (e.g., “kibby wind” or “coughlin power”), which are of interest to call participants who discuss the climate change exposure of these facilities’ operators. These bigrams illustrate the challenge of using training libraries or pre-specified word lists to identify climate change talk—few researchers have the detailed field knowledge to recognize the relationship between these words and climate change.

Our approach allows us to adapt the bigram-search algorithm to discover three unique sets of bigrams from  $\mathbb{C}$  that capture opportunities as well as regulatory and physical climate shocks. Toward this end, we feed a set of initial bigrams reflecting these three topics to the search algorithm. We then allow the algorithm to discover bigrams related to the topic of interest. Table IA.IV lists the initial bigrams used for the topic search. We construct new initial bigrams for these topics by hand-picking appropriate bigrams from the top-500 bigrams discovered after the first generic, nontopic-specific bigram search. We then rerun the search algorithm to find a broader set of bigrams for each topic. As the topic-based algorithm yields some general climate change bigrams, we drop bigrams appearing in more than one topic to guarantee that we do not have overlapping topic measures. In the final stage, we take the intersection between  $\mathbb{C}$  and each set of topic bigrams to obtain the sets of opportunity, regulatory, and physical climate change bigrams (i.e.,  $\mathbb{C}^{Opp}$ ,  $\mathbb{C}^{Reg}$ , and  $\mathbb{C}^{Phy}$ ), respectively.

### B. Construction of Climate Change Exposure Measures

Using the bigram sets, we construct measures of climate change exposure for each transcript. We interpret these measures as capturing the attention devoted to climate change topics by call participants at a point in time, rather than as measures of fundamental exposure. We use the broad set of climate change bigrams  $\mathbb{C}$  to illustrate how we construct these measures. The topic

<sup>20</sup> For example, “rooftop solar” and “photovoltaic panel” come from the initial bigram “solar energy,” while “nuclear power” and “event fukushima” come from “renewable energy,” and “tesla battery” and “hybrid plug” come from “electric vehicle.”

measures are constructed analogously; we simply replace  $\mathbb{C}$  with the bigrams that relate to the corresponding topic.

We construct an overall exposure measure,  $CCExposure$ , based on how frequently the specified bigrams appear in a transcript. This involves taking the set of climate bigrams  $\mathbb{C}$  to the transcript of firm  $i$  in quarter  $t$  and counting the frequency of these bigrams. To account for the call length, we scale the count by the number of bigrams in the transcript,

$$CCExposure_{i,t} = \frac{1}{B_{i,t}} \sum_b^{B_{i,t}} (1[b \in \mathbb{C}]), \quad (1)$$

where  $b = 0, 1, \dots, B_{i,t}$  are the bigrams in the earnings call transcripts of firm  $i$  in quarter  $t$  and  $1[\cdot]$  is the indicator function. We create an annual measure for each firm by averaging the quarterly measures. We produce exposure measures from  $\mathbb{C}^{Opp}$ ,  $\mathbb{C}^{Reg}$ , and  $\mathbb{C}^{Phy}$ , respectively, by scoring each transcript using the same method. We label the topic-based measures as  $CCExposure^{Opp}$ ,  $CCExposure^{Reg}$ , and  $CCExposure^{Phy}$ .

Some of our tests employ two refinements. In the first refinement, we create two sentiment measures by counting the number of climate change bigrams after conditioning on the presence of the positive or negative tone words in Loughran and McDonald (2011),

$$CCSentiment_{i,t}^{Pos/Neg} = \frac{1}{B_{i,t}} \sum_b^{B_{i,t}} \left\{ (1[b \in \mathbb{C}]) \times \sum_{b \in S} \mathcal{T}^{Pos/Neg}(b) \right\}, \quad (2)$$

where  $S$  represents the sentence containing bigrams  $b = 0, 1, \dots, B_{i,t}$  and  $\mathcal{T}^{Pos/Neg}(b)$  assigns sentiment to each bigram  $b$ :<sup>21</sup>

$$\mathcal{T}^{Pos}(b) = \begin{cases} 1 & \text{if } b \text{ has a positive tone,} \\ 0 & \text{if otherwise,} \end{cases}$$

$$\mathcal{T}^{Neg}(b) = \begin{cases} 1 & \text{if } b \text{ has a negative tone,} \\ 0 & \text{if otherwise.} \end{cases}$$

In the second refinement, we construct a measure of risk by counting the relative frequency of the climate change bigrams mentioned in the same sentence

<sup>21</sup> Though not used in this paper, we also combine both sentiment measures into an overall measure by counting the climate change bigrams after conditioning on the presence of positive and negative tone words,

$$CCSentiment_{i,t} = \frac{1}{B_{i,t}} \sum_b^{B_{i,t}} \left\{ (1[b \in \mathbb{C}]) \times \sum_{b \in S} \mathcal{T}(b) \right\},$$

where  $\mathcal{T}(b) = 1(-1)$  if  $b$  has positive (negative) tone, and zero otherwise.

with the words “risk,” “uncertainty,” or their synonyms,

$$CCRisk_{i,t} = \frac{1}{B_{i,t}} \sum_b^{B_{i,t}} (1[b \in \mathbb{C}] \times 1[b, r \in S]), \quad (3)$$

where  $r$  contains the words “risk,” “uncertainty,” or a synonym.

The exposure measures do not adjust for the differences in the importance or typical frequencies of individual bigrams. For robustness, we account for such differences by constructing measures that weigh each bigram with a score reflecting the bigram’s representativeness for climate discussions. We do this so that common terms that appear in most transcripts receive low scores, as these terms are less informative about a call’s content, as do rare terms in a given transcript, as these terms have low text frequency. This approach follows Hassan et al. (2019), Gentzkow, Kelly, and Taddy (2019), and EGKLS and is commonly referred as “term frequency-inverse document frequency” (TFIDF). Formally,

$$CCExposure_{i,t}^{TFIDF} = \frac{1}{B_{i,t}} \sum_b^{B_{i,t}} \left( 1[b \in \mathbb{C}] \times \log \left( \frac{N_{\mathbb{T}}}{f_{b,\mathbb{T}}} \right) \right), \quad (4)$$

where  $N_{\mathbb{T}}$  refers to the number of transcripts and  $f_{b,\mathbb{T}}$  to the number of transcripts in which bigram  $b$  appears. A bigram appearing in many transcripts therefore has low weight when calculating the TFIDF score, and—in the extreme case—if a given bigram appears in every transcript, it receives zero weight ( $\log(\frac{N_{\mathbb{T}}}{f_{b,\mathbb{T}}}) = 0$ ).

Table I reports summary statistics for the exposure measures (for purposes of exposition, the measures are multiplied by  $10^3$ ).<sup>22</sup> Table IA.V reports the correlations across the exposure measures. A few correlations deserve further comment. The correlation between  $CCExposure^{Reg}$  and  $CCExposure^{Opp}$  is positive at 33%, and  $CCExposure^{Phy}$  is largely unrelated to  $CCExposure^{Reg}$  and  $CCExposure^{Opp}$ . In addition, the correlation between  $CCExposure$  and  $CCExposure^{TFIDF}$  is 99.7%.

Tables IA.VI to IA.VIII report the sample distribution at the earnings-call (transcript) level across countries, years, and industries. We report the distributions for all sampled earnings calls and for those calls with nonzero climate change exposure. The tables show meaningful proportions of calls with nonzero climate change exposure across all three sample cuts; transcripts with  $CCExposure > 0$  are not concentrated in certain countries, years, or industries. Our analysis does not make use of a binary indicator for whether  $CCExposure$  is nonzero, but instead uses a continuous measure.

<sup>22</sup> The magnitudes of  $CCExposure^{TFIDF}$  are larger than those of  $CCExposure$  as the inverse document frequency of the climate change bigrams can be much larger than one (the document frequencies of the climate change bigrams are much smaller than the total number of transcripts).



**Table I**  
**Climate Change Exposure Variables: Summary Statistics**

This table reports summary statistics for different measures of climate change exposure, carbon emissions, and public attention to climate change at the firm-year level. For the climate change exposure measures, we average values of the four earnings calls during the year. The sample includes 10,673 unique firms from 34 countries over the period 2002 to 2020. Table A.1 provides detailed variable definitions.

	Mean	STD	25%	Median	75%	N
CC Measures ( $\times 10^3$ )						
$CCExposure_{i,t}$	1.01	2.53	0.10	0.30	0.78	86,152
$CCExposure_{i,t}^{Opp}$	0.31	1.23	0.00	0.00	0.15	86,152
$CCExposure_{i,t}^{Reg}$	0.04	0.23	0.00	0.00	0.00	86,152
$CCExposure_{i,t}^{Phy}$	0.01	0.11	0.00	0.00	0.00	86,152
CC Measures (TFIDF Version) ( $\times 10^3$ )						
$CCExposure_{i,t}$	7.99	19.69	0.77	2.44	6.26	86,152
$CCExposure_{i,t}^{Opp}$	2.35	9.08	0.00	0.00	1.18	86,152
$CCExposure_{i,t}^{Reg}$	0.32	1.68	0.00	0.00	0.00	86,152
$CCExposure_{i,t}^{Phy}$	0.10	0.81	0.00	0.00	0.00	86,152
CC Q&A Measure ( $\times 10^3$ )						
$CCExposure_{i,t}^{Q\&A}$	0.67	1.95	0.00	0.12	0.54	86,152
CC Sentiment and Risk Measures ( $\times 10^3$ )						
$CCSentiment_{i,t}^{Pos}$	0.38	1.10	0.00	0.07	0.32	86,152
$CCSentiment_{i,t}^{Neg}$	0.19	0.55	0.00	0.00	0.16	86,152
$CCRisk_{i,t}$	0.04	0.17	0.00	0.00	0.00	86,152
Carbon Emissions and Climate Change Attention						
$TotalEmissions_{i,t}$	2,961,549	13,608,989	27,472	133,847	751,772	33,789
$WSJCCNewsIndex_t$	0.007	0.001	0.006	0.006	0.008	68,794

### III. Validation

#### A. Validation at the Bigram Level

##### A.1. Face Validity of Climate Change Bigrams

We validate our exposure measures using a multipronged approach. First, we consider the bigrams' face validity. Table II lists the 100 highest-frequency bigrams in C. The top bigrams associated with  $CCExposure$  capture aspects of the opportunities and risks associated with climate change. The top bigrams include both opportunity-related word pairs (e.g., "battery power," "new energy") and risk-related terms (e.g., "environmental concern," "extreme weather").

Table IA.IX considers the three topic-based measures. When we construct  $CCExposure^{Opp}$  using initial bigrams such as "wind power" or "solar energy," we find several new bigrams that refer to new (green) technologies (e.g., "solar farm," "carbon free") (Panel A). Several word combinations are linked to developments in "electric vehicles," including "charge infrastructure" and "battery electric." With respect to  $CCExposure^{Reg}$  (Panel B), when we use initial bigrams

**Table II**  
**Top-100 Bigrams Captured by Climate Change Exposure**  
**(CCE<sub>exposure</sub>)**

This table reports the top-100 bigrams associated with *CCE<sub>exposure</sub>*, which measures the relative frequency with which bigrams related to climate change occur in earnings call transcripts. Table A.1 defines all variables in detail.

Bigram	Frequency	Bigram	Frequency	Bigram	Frequency
renewable energy	15,605	onshore wind	878	carbon intensity	641
electric vehicle	9,508	electric motor	869	energy application	615
clean energy	6,430	provide energy	851	produce electricity	604
new energy	4,544	efficient solution	839	help state	604
climate change	4,374	global warm	837	environmental	593
				standard	
wind power	4,253	power generator	828	power agreement	586
wind energy	4,035	solar pv	827	supply energy	585
energy efficient	3,899	scale solar	827	electric hybrid	585
greenhouse gas	3,416	need clean	821	source power	575
solar energy	2,511	coastal area	816	sustainability goal	572
air quality	2,409	energy star	793	energy reform	571
clean air	2,301	environmental	792	plant power	564
		footprint			
carbon emission	2,088	design use	777	compare con-	560
				ventional	
gas emission	1,910	area energy	777	gas vehicle	560
extreme weather	1,773	charge station	762	effort energy	560
carbon dioxide	1,583	clean water	759	pass house	559
water resource	1,423	major design	747	carbon free	558
autonomous vehicle	1,394	vehicle manu-	740	driver assistance	545
		facturer			
energy environment	1,279	future energy	737	electrical energy	543
wind resource	1,245	motor control	726	solar installation	541
government india	1,201	combine heat	718	snow ice	538
battery power	1,147	electric bus	709	renewable natural	536
air pollution	1,127	distribute power	703	promote use	536
battery electric	1,121	environmental	695	farm project	531
		benefit			
integrate resource	1,052	eco friendly	695	laser diode	528
clean power	1,008	electrical vehicle	695	deliver energy	526
carbon price	999	carbon neutral	690	protect environ-	525
				ment	
world population	977	fast charge	675	sustainable energy	523
solar farm	971	cell power	657	manage energy	522
energy regulatory	967	energy team	650	invest energy	521
obama administration	957	cycle gas	646	electric energy	519
heat power	941	coal gasification	643	forest land	512
carbon tax	928	environmental	643	capacity energy	512
		concern			
unite nation	925				


“carbon tax,” “air pollution,” or “air quality,” that is, terms related to regulatory interventions, we discover bigrams that explicitly include the word “regulation” or its synonyms (e.g., “control regulation,” “environmental legislation”). Turning to the top bigrams for *CCExposure<sup>Phy</sup>* (Panel C), we use initial bigrams such as “natural hazard” or “sea level” to identify word pairs intuitively linked to physical climate change (e.g., “area florida,” “ice control,” and “wind speed”).

For the 10 highest-scoring firms on *CCExposure*, Table [IA.X](#) provides “snippets.” These snippets are text fragments taken from the point in the transcript that the algorithm identifies as the moment when the participants discuss climate issues. Consider Ocean Power Technologies, a U.S. firm that turns ocean wave power into electricity for offshore applications. In its 2008Q4 call, bigrams such as “energy requirement,” “powerbuoy wave,” “wave condition,” and “wave power” were heavily featured. In the top snippet, participants discuss the increased demand for the firm’s trademark technology (the Power-Buoy®) due to heightened attention to renewable energy. Not surprisingly, high-scoring firms are involved in energy production or the broader energy infrastructure. Indeed, when ECOTALITY call participants use climate change bigrams, they discuss how charging infrastructures are central to advancing zero-emissions transportation.

### A.2. Audit Study Based on Human Reading

We developed a two-stage snippet-based audit to evaluate the scoring of our algorithm (Baker, Bloom, and Davis (2016), Hassan et al. (2019)). While our algorithm should be judged in the context of the entire transcript, a snippet-based audit improves our ability to sample across a large number of transcripts. In the first stage, we define a snippet as the 10 sentences around the climate change bigram with the highest text frequency in a transcript. For transcripts with *CCExposure* = 0, we randomly choose a snippet of 10 consecutive sentences for the audit. In our pilot study, each of the authors independently coded 250 identical and randomly selected snippets using a binary coding scheme. The coding used the variable *CCAudit*, which equals one if the rater classifies the text as providing evidence of climate change exposure, and zero otherwise. In addition, for each snippet we record *Coding Confidence*, which ranges from three (the rater is highly confident that their coding is correct) to one (“hard calls”). We identified some slight coding differences between the authors and resolved discrepancies. Based on this iterative procedure, we developed a detailed guide with definitions of what text should be coded as climate change exposure and which snippets should not qualify as such. The audit guide describes examples of snippets and offers interpretations and suggested coding to help the raters solve complex cases in the audit process. We then instructed two graduate students based on the audit guide and asked them to audit the same 250 snippets that the author team coded to assess any remaining inconsistencies.

In the second stage, we recruited 19 graduate students to each independently code 250 new snippets from the audit universe. Together they



**Figure 1. Probability of correctly identified positives by decile.** This figure plots on the vertical axis the predicted probability of having a correctly identified positive (i.e., the audit study of the snippet confirms climate change-related text) against deciles of the *CCEXposure* distribution. The median score of *CCEXposure* in a given decile is shown on the axis. Predicted probabilities are computed by estimating a logit model on the sample of 2,090 audited snippets. (Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com))

assessed 2,090 unique snippets.<sup>23</sup> Auditors received training based on the audit guide. The snippets were partially overlapping to allow us to conduct some inter-rater correspondence tests. Our goal is to verify the information content of *CCEXposure* at various points of its distribution. Following Hassan et al. (2019), we create portfolios with the same number of transcripts based on their percentile of the *CCEXposure* distribution. We then count the number of transcripts at that percentile that the auditors rated as *CCAudit* = 1 (i.e., the snippet is classified as containing a clear discussion of a firm's climate change exposure). We count 310 true positives out of 339 snippets (91% correct positives) in the top-decile portfolio (transcripts with the highest value of *CCEXposure*). The rate of correct positives declines almost linearly as we move to the median and bottom portfolios. This is displayed in Figure 1, which plots the relationship between (the predicted probability of) true positives (as judged by the human reading) at each decile and the median percentile score of *CCEXposure* at that

<sup>23</sup> We first sorted all transcripts with nonzero *CCEXposure* into deciles. We then randomly selected 10 snippets from each decile and another 10 from *CCEXposure* = 0 transcripts for each sample year.

percentile. The association is positive and nearly linear, as would be expected if our algorithm reliably identifies climate change discussions.<sup>24</sup>

### A.3. Comparison with Approach Using Pre-Specified Keywords

We construct alternative exposure measures from a list of pre-specified climate change keywords to compare these measures with those produced by our algorithm. To obtain such a list, we use the set of unique stemmed unigrams and bigrams  $\mathbb{C}^{EGKLS}$  used by EGKLS to build their time-varying, news-based index of climate change attention. These keywords originate from 74 authoritative texts. To create  $CCExposure^{EGKLS}$ , we replace  $\mathbb{C}$  with  $\mathbb{C}^{EGKLS}$  and recompute the relative frequency with which the alternative terms appear in the transcripts. We construct a frequency-unweighted and a TFIDF version, denoted by  $CCExposure^{EGKLS-EW}$  or  $CCExposure^{EGKLS-TFIDF}$ , respectively.

Table IA.XI illustrates that the unigrams and bigrams in  $\mathbb{C}^{EGKLS}$  appear more frequently in earnings calls than the bigrams in  $\mathbb{C}$ . This finding is unsurprising as  $\mathbb{C}^{EGKLS}$  includes more unigrams and more general terms (the top-3 bigrams are “market,” “increase,” and “time”). Using unigrams rather than bigrams trades off the higher likelihood of a given term occurring in a text against the higher probability of a false positive, that is, wrongly classifying a fragment as climate change text (van Zaanen and Kanters (2010)). Several of the unigrams in  $\mathbb{C}^{EGKLS}$  are part of the top-100 bigrams in  $\mathbb{C}$  (e.g., “carbon,” “energy,” or “water”). As would be expected from Table IA.XI, the mean values of both alternative exposure measures (Table IA.XII, Panel A) are larger than those of  $CCExposure$ . Thus, larger parts of the earnings calls are classified as discussing climate topics if we use  $\mathbb{C}^{EGKLS}$  instead of  $\mathbb{C}$ . At the same time, Table IA.XII, Panel B, indicates that the measures correlate positively with  $CCExposure$ . The correlation table illustrates that our main measure and the alternative measures yield more similar assessments when the public pays close attention to climate change (*WSJ CC News Index* is in the top quartile). One possible explanation is that at times when the *WSJ* devotes a lot of space to climate topics, terms from a more general climate library (on which the index and pre-specified keyword measures build) become more commonly used in earnings calls. Intuitively, media attention might homogenize the language used to talk about climate change. When the media pivots to other events, the vocabulary likely used to discuss climate change in earnings calls becomes more idiosyncratic again. Such instances are plausibly better reflected in our keyword-search-based approach.

A question that remains is how our measure and a measure using pre-defined keywords differ economically. Our measure is well suited to capture

<sup>24</sup> These findings suggest that our algorithm correctly identifies climate change text, even at relatively low  $CCExposure$  scores. A benchmark is provided in Hassan et al. (2019), where the number of correct positives reduces to below five out of 20 at the 90<sup>th</sup> percentile of their text-based political risk score. Weighting observations by *Coding Confidence* does not materially change our findings.

context-specific jargon used in specialized environments with experts and allows us to construct topic-based measures. The pre-specified keyword approach better captures broader discussions by the public, as reflected in articles published in the WSJ, while identifying specific or emerging topics with a pre-specified keyword approach is more challenging. A further difference is that our approach is “evolutionary,” that is, it will reflect changes in the vocabulary used in transcripts over time, while an approach using pre-specified keywords fixes this vocabulary *ex ante*. Time-series variation in true (unobservable) climate change exposure, especially over long horizons, is more likely to be picked up by such an “evolutionary” approach. Any selection of pre-specified keywords is due to become obsolete the further out one moves in time.

#### A.4. Perturbation Tests for Individual Initial Bigrams

We evaluate the extent to which our overall exposure measure depends on *individual* bigrams in the initial bigram list (Table IA.III) by performing a perturbation test. We successively exclude one initial bigram at a time, recomputing the modified set of bigrams  $\mathbb{C}^{Pert}$  as well as the modified  $CCExposure^{Pert}$ . Given that our initial short list contains 50 bigrams, we construct 50 new versions of  $CCExposure^{Pert}$ . After aggregating the measure to the firm-year level, we calculate the correlation of each of these exposure measures with  $CCExposure$ . These correlations are above 85%, which means that  $CCExposure$  does not depend much on specific initial seed bigrams.

#### A.5. Comparison with Approach Using Initial Bigrams Only

Table II shows that the initial keywords dominate the top-100 bigrams used in the construction of  $CCExposure$ . This raises the question of how big the performance gain of the keyword discovery approach is relative to the alternative that only uses the initial seed bigrams. To address this question, we construct the new exposure measure  $CCExposure^{Initial}$  from the initial bigrams only. Figure 2, Panel A, shows how frequently the new measure signals zero exposure, while  $CCExposure$  instead reveals that climate topics are discussed. Results are reported by  $CCExposure$  decile. In the top decile,  $CCExposure^{Initial}$  indicates no exposure in 27% of transcripts. Hence, even among the most exposed firms, there is a performance gain when applying our approach. This gain increases when we consider other deciles—already in the second decile,  $CCExposure^{Initial}$  deviates from  $CCExposure$ , indicating the absence of exposure in more than 62% of transcripts. The effects increase monotonically as we move to lower exposure deciles.

Panel B reports results of the topic-based exposure measures, with the alternative measures using only the topic-based initial bigrams (Table IA.IV). For all three measures and deciles, significant fractions of the transcripts are incorrectly classified as having zero exposure. Even in the three respective top deciles, the alternative approach misses positive exposure in 10% to 40% of the



transcripts. Across all deciles, the gain from the keyword discovery approach is largest for  $CCExposure^{Opp}$  (especially in the lower deciles).

Beyond these statistics, identifying exposure using bigrams beyond the initial seed words is economically important. Below we show that, among the set of firms for which  $CCExposure^{Initial} = 0$ , our exposure measures keep predicting green outcomes. These effects are identified purely from the bigrams obtained through the keyword search algorithm.

## B. Validation at the Climate Change Exposure Level

### B.1. Climate Change Exposure: Industry Variation

We now move away from the bigram level to examine the properties of the exposure measures. This involves several steps. In the first step, we compute averages by industry sector (two-digit SIC code level) and present a ranking of these means in Table III. In Panel A, using  $CCExposure$ , the sectors with the highest overall exposure include Electric, Gas, & Sanitary (SIC49). Top-ranked firms within this sector include China Longyuan Power Group, China's largest producer of wind power, and the U.S. utility Allete. This sector is followed by Heavy Construction (SIC16) and Construction (SIC17). High-ranking firms in these sectors include A-Power Energy Generation Systems, a Chinese firm providing on-site power generation systems, ReneSola, a U.S. firm developing and operating solar projects, and Quanta Services, a U.S. infrastructure solutions provider for firms in the energy and pipeline business. Top-ranked firms in the Transportation Equipment sector (SIC37), ranked next, include alternative fuel and zero-emission vehicle firms.

A few sectors are worth commenting on in Panels B to D, which report the topic-based measures. Utilities top the list for  $CCExposure^{Opp}$  (Panel B) and  $CCExposure^{Reg}$  (Panel C). While the latter ranking position is expected, given the sector's exposure to carbon taxes or related regulations, the earlier position is more surprising. Notwithstanding, it is consistent with Cohen, Gurun, and Nguyen (2021), who find that this sector is a key innovator in the energy transition space. Coal Mining (SIC12) displays high exposure to regulatory and physical shocks (Panels C and D). While high regulatory exposure is expected given the large emissions associated with burning coal, high physical exposure is less obvious. This may reflect mining firms' exposure to heavy precipitation, or heat, which pose physical challenges to their operations. Stone, Clay & Glass Products (SIC32), in the top-5 for  $CCExposure^{Reg}$ , includes mostly cement producers among its top-ranked firms (they belong to the largest CO<sub>2</sub> emitters). A sector in the top-10 of  $CCExposure^{Phy}$  (Panel D) is the insurance industry, which, unsurprisingly, is highly exposed to the costs of storms or flooding.

The large variation in exposure *between* sectors masks important heterogeneity *within* each sector (apparent from the large within-sector standard deviations). To illustrate this heterogeneity, we compare TotalEnergies and ExxonMobil. Both firms operate in Petroleum Refining (SIC29), a sector ranking among the top 10 for  $CCExposure^{Opp}$  and  $CCExposure^{Reg}$ . In terms of the

Table III  
Industry Distribution of Climate Change Exposure Measures

This table reports firms' climate change exposure measures for the top-10 industries. Statistics are reported at the firm-year level across different SIC2 industries. We rank sectors by the average values of the climate change exposure measures.  $CCExposure$  measures the relative frequency with which climate change bigrams occur in earnings calls.  $CCExposure^{Opp}$  measures the relative frequency with which bigrams that capture opportunities related to climate change occur in earnings calls.  $CCExposure^{Reg}$  measures the relative frequency with which bigrams that capture regulatory shocks related to climate change occur in earnings calls.  $CCExposure^{Phy}$  measures the relative frequency with which bigrams that capture physical shocks related to climate change occur in earnings calls. For all measures, we average values of the four earnings calls during the year. We report results only those industries for which we have more than 20 firm-year observations. Table A.1 defines all variables in detail.

Panel A:  $CCExposure$  ( $\times 10^3$ )

Industry (SIC2)	Mean	Std.Dev.	Median	N
49 Electric, Gas, & Sanitary Services	6.95	6.23	5.34	3,259
16 Heavy Construction, Except Building	3.04	4.35	1.53	537
17 Construction	2.26	2.95	1.16	131
37 Transportation Equipment	2.12	3.17	1.07	2,021
36 Electronic & Other Electric Equipment	2.07	4.20	0.57	5,812
12 Coal Mining	2.05	1.48	1.70	253
29 Petroleum Refining	1.72	2.14	1.06	730
41 Local & Suburban Transit	1.69	2.06	0.84	94
55 Automotive Dealers & Service Stations	1.63	3.90	0.69	484
33 Primary Metal	1.56	1.54	1.14	1,149

Panel B:  $CCExposure^{Opp}$  ( $\times 10^3$ )

Industry (SIC2)	Mean	Std.Dev.	Median	N
49 Electric, Gas, & Sanitary Services	2.50	3.30	1.26	3,259
16 Heavy Construction, Except Building	1.37	2.78	0.30	537
17 Construction	0.91	1.71	0.34	131
36 Electronic & Other Electric Equipment	0.90	2.38	0.09	5,812
37 Transportation Equipment	0.81	1.70	0.23	2,021
55 Automotive Dealers & Service Stations	0.54	1.34	0.16	484
29 Petroleum Refining	0.47	0.93	0.16	730
35 Industrial Machinery & Equipment	0.46	1.85	0.07	4,056
75 Auto Repair, Services, & Parking	0.42	1.04	0.11	171
87 Engineering & Accounting & Research	0.38	0.94	0.00	1,443

Panel C:  $CCExposure^{Reg}$  ( $\times 10^3$ )

Industry (SIC2)	Mean	Std.Dev.	Median	N
49 Electric Gas & Sanitary Services	0.34	0.61	0.10	3,259
12 Coal Mining	0.14	0.24	0.00	253
29 Petroleum Refining	0.14	0.32	0.00	730
32 Stone Clay Glass Products	0.12	0.35	0.00	622
10 Metal Mining	0.08	0.32	0.00	1,465
33 Primary Metal	0.08	0.22	0.00	1,149
37 Transportation Equipment	0.08	0.27	0.00	2,021
35 Industrial Machinery & Equipment	0.08	0.47	0.00	4,056

(Continued)

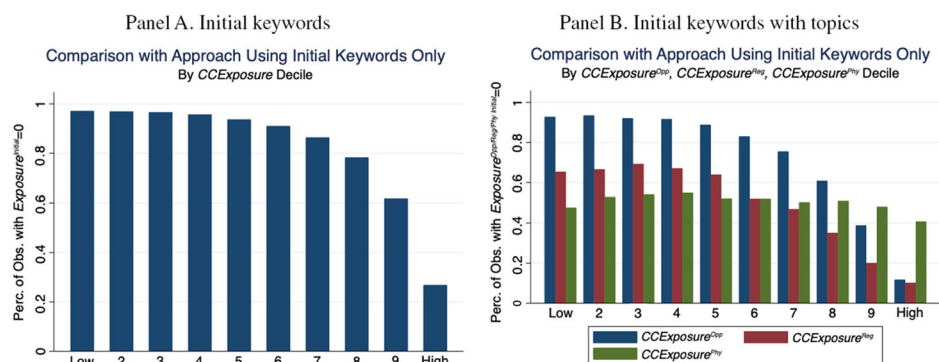
Table III—Continued

Panel C: $CCExposure^{Reg} (\times 10^3)$				
24 Lumber & Wood	0.07	0.43	0.00	471
16 Heavy Construction	0.07	0.21	0.00	537
Panel D: $CCExposure^{Phy} (\times 10^3)$				
Industry (SIC2)	Mean	Std.Dev.	Median	N
41 Local and Suburban Transit	0.17	0.47	0.00	94
26 Paper & Allied Products	0.08	0.35	0.00	852
24 Lumber & Wood	0.07	0.26	0.00	471
49 Electric, Gas, & Sanitary Services	0.06	0.24	0.00	3,259
14 Mining & Quarrying	0.05	0.14	0.00	208
12 Coal Mining	0.04	0.19	0.00	253
64 Insurance Agents, Brokers, & Service	0.03	0.15	0.00	297
10 Metal Mining	0.03	0.12	0.00	1,465
15 Building Construction	0.03	0.09	0.00	600
35 Industrial Machinery & Equipment	0.03	0.25	0.00	4,056

average regulatory exposure since 2010, TotalEnergies’ score is similar to that of ExxonMobil ( $CCExposure^{Reg}_{TotalEnergies} = 0.21$  vs.  $CCExposure^{Reg}_{ExxonMobil} = 0.18$ ), but the French oil major exhibits much higher average opportunity exposure ( $CCExposure^{Opp}_{TotalEnergies} = 1.13$  vs.  $CCExposure^{Opp}_{ExxonMobil} = 0.15$ ). This divergence reflects a broader perception in the market about the extent to which these firms embrace renewable energy and the net-zero transition in their business models (see Pickl (2019)). More generally, the large within-industry variation indicates that sectors have “winners” and “losers.” Investors may therefore be able to address climate risks and opportunities by maintaining a broad industry diversification (rather than banning some industries) and then performing negative screening of climate change “losers.” This observation echoes arguments by both academics (Andersson, Bolton, and Samama (2016)) and providers of low-carbon index solutions.

B.2. Climate Change Exposure: Time-Series Variation

In Figure 3, Panels A to D, we compute the cross-sectional means for  $CCExposure$  and the topic-based measures and plot them over time (for each measure, we focus on top-10 sectors). This figure also highlights key moments in public awareness of climate change, covering climate policy events relevant to regulatory and opportunity shocks (Panels B and C), select physical shocks (Panel D), or both (Panel A). In Panel A,  $CCExposure$  generally increases over the sample period, especially since the mid-2000s. The increase in the early years indicates that earnings calls discussed climate issues earlier than we might have expected. A plateau is reached around 2009 (the year of the unsuccessful Copenhagen Climate Summit). We then observe a slight decline in the years leading up to the 2012 Doha Climate Summit. We note a renewed



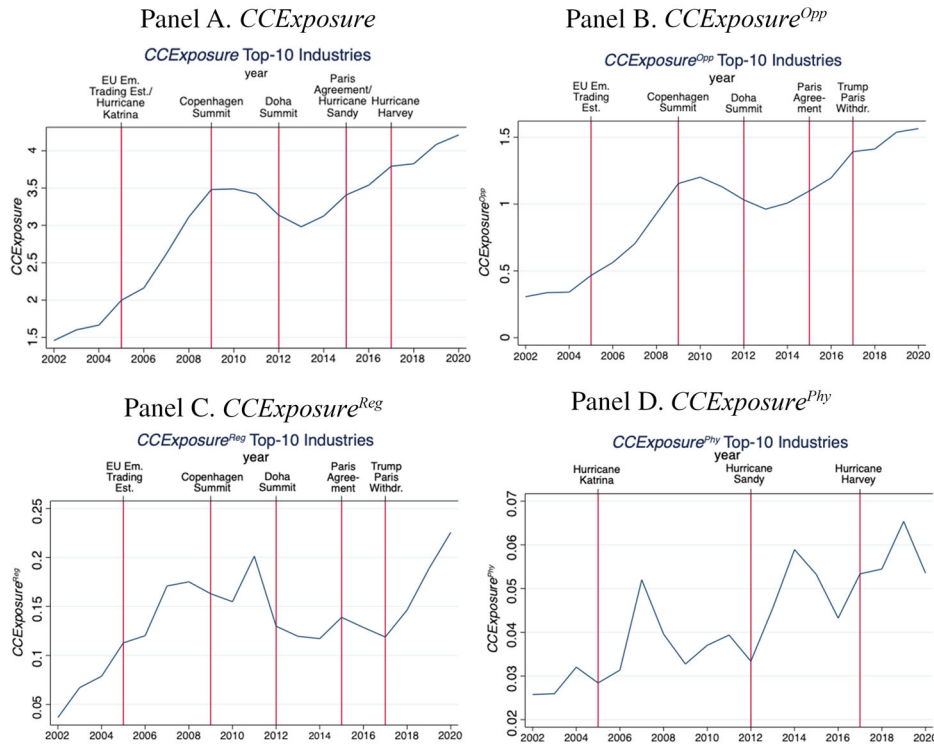
**Figure 2. Climate change exposure calculated with initial bigrams.** This figure shows how frequently  $CCExposure^{Initial}$  signals zero climate change exposure, while  $CCExposure$  instead reveals that such exposure exists. Results are reported by  $CCExposure$  decile.  $CCExposure^{Initial}$  is a measure of climate change exposure based on the initial seed bigrams only. Panel A reports results for the overall climate change exposure measure, and Panel B for the topic-based measures. In the figure, the exposure measures are calculated at the quarterly (transcript) level. (Color figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com))

increase in  $CCExposure$  since around 2013. At the end of the sample,  $CCExposure$  peaks with earnings calls exhibiting about four climate change bigrams per 1,000 bigrams; this compares to about 0.1 political bigrams per 1,000 bigrams in Hassan et al. (2019).

In Panel B, the time series for  $CCExposure^{Opp}$  is similar to that of the overall measure:  $CCExposure^{Opp}$  trends upward, especially at the beginning of the sample. In Panel C,  $CCExposure^{Reg}$  increases between 2002 and 2008, varies around a markedly lower level between 2011 and 2013, spikes in 2015 (Paris Agreement), and follows an increasing trend since 2017. This is consistent with intensified policy discussions about how to achieve the Paris goals. In Panel D,  $CCExposure^{Phy}$  displays more swings than the other measures, albeit also around an upward trend. It appears that  $CCExposure^{Phy}$  does not strongly reflect highly salient climate events. For example, while there is a jump after major U.S. hurricanes (i.e., Katrina, Sandy, and Harvey), the jumps occur with a considerable lag. This pattern indicates that  $CCExposure^{Phy}$  primarily reflects firm-specific exposures to physical climate events, (e.g., local heat waves or droughts).

### B.3. Climate Change Exposure and Carbon Emissions

We explore how well the exposure measures correlate with firms' carbon emissions. Carbon emissions constitute an essential variable to measure firm-level exposure to climate change, especially for regulatory shocks (Bolton and Kacperczyk (2021, 2023)). The analysis of carbon emissions is also the most frequently used climate risk management tool of institutional investors (Krueger, Sautner, and Starks (2020)). A benefit of using carbon emissions is that they



**Figure 3. Climate change exposure over time.** This figure shows firms' average climate change exposures over time.  $CCEXposure$  measures the relative frequency with which climate change bigrams occur in earnings calls.  $CCEXposure^{Opp}$  measures the relative frequency with which bigrams that capture opportunities related to climate change occur in earnings calls.  $CCEXposure^{Reg}$  measures the relative frequency with which bigrams that capture regulatory shocks related to climate change occur in earnings calls.  $CCEXposure^{Phy}$  measures the relative frequency with which bigrams that capture physical shocks related to climate change occur in earnings calls. For each exposure measure, we construct the time series for firms in the top-10 industries (see Table III). Table A.1 provides detailed variable definitions. (Color figure can be viewed at [wileyonlinelibrary.com](https://onlinelibrary.wiley.com))

are easy to understand and compute, readily available for subscribers of environmental, social, and governance (ESG) databases, and genuinely related to changes in the global climate.

We expect that regulatory climate topics arise more frequently in earnings calls of large carbon emitters, as they are more strongly affected by carbon taxes or related regulations. At the same time, regulatory threats related to emissions may also spur technological innovation that provides firms with opportunities in the marketplace.<sup>25</sup> Furthermore, some firms' emissions may be "good" in supporting the transition to a greener economy; these firms,

<sup>25</sup> For example, utilities with a large carbon footprint may have strong incentives to develop low-carbon alternatives (e.g., wind farms, solar farms), which provide future opportunities. Indeed, as

called “climate enablers,” include, for example, manufacturers of building materials that help houses become more energy-efficient. Finally, carbon emissions should be unrelated to the exposure to physical shocks at the firm level.

We test these predictions by regressing the exposure measures on lagged emission values (we use lagged values because emissions covering year  $t - 1$  are reported in year  $t$ ). Table IV, Panel A, reports the results. In column (1), we observe a strong positive association between *Total Emissions* and *CCExposure*. As predicted, this association originates from positive correlations between emissions and both *CCExposure<sup>Opp</sup>* (column (2)) and *CCExposure<sup>Reg</sup>* (column (3)). A one-standard-deviation increase in the emissions variable is associated with an increase in *CCExposure<sup>Reg</sup>* equal to 23% of its standard deviation (using values for the regression sample). In column (4), we find no association between emissions and physical exposure.

#### B.4. Climate Change Exposure and Public Attention to Climate Change

Time-series variation in public attention to climate change, as proxied by *WSJ CC News Index*, has been shown to affect financial market participants (e.g., Choi, Gao, and Jiang (2020) or Ilhan, Sautner, and Vilkov (2021)). Accordingly, we expect earnings call discussions to react to the salience of climate topics in the public arena. Indeed, Table IV, Panel B, shows that measured climate change exposure is higher at times when public climate attention rises. In column (1), a one-standard-deviation increase in *WSJ CC News Index* is associated with an increase in *CCExposure* of 0.05 (5% of the mean within the regression sample). This effect reflects a positive association between *WSJ CC News Index* and both *CCExposure<sup>Opp</sup>* and *CCExposure<sup>Reg</sup>*. Hence, when public climate attention is high, earnings calls discuss regulatory shocks and climate opportunities more extensively. Higher values of *WSJ CC News Index* do not translate into more discussions of physical shocks. This suggests that *CCExposure<sup>Phy</sup>* mostly captures firm-specific physical shocks, rather than economy-wide shocks that make it to the WSJ (this conclusion is consistent with the time-series evidence in Figure 3).

### IV. Variance Decomposition and Role of Measurement Error

#### A. Variance Decomposition

We conduct a variance analysis to examine the extent to which *CCExposure* and its components quantify firm-level variation in climate change exposure. Table V reports the incremental explanatory power from conditioning the exposure measures on fixed effects that plausibly drive the variation. Time fixed effects (i.e., economy-wide changes in aggregate exposure) explain little

mentioned above, Cohen, Gurun, and Nguyen (2021) demonstrate that some of the largest carbon emitters produce more and better green innovation than other firms.



Table IV  
Climate Change Exposure Measures: Effects of Carbon Emissions  
and Climate Change News

This table reports regressions that relate carbon emissions and climate change news to the climate change exposure measures. Regressions are estimated at the firm-year level.  $CCExposure$  measures the relative frequency with which bigrams related to climate change occur in earnings calls.  $CCExposure^{Opp}$  measures the relative frequency with which bigrams that capture opportunities related to climate change occur in earnings call transcripts.  $CCExposure^{Reg}$  measures the relative frequency with which bigrams that capture regulatory shocks related to climate change occur in earnings calls.  $CCExposure^{Phy}$  measures the relative frequency with which bigrams that capture physical shocks related to climate change occur in earnings calls. For all measure, we average values of the four earnings calls during the year. *Total Emissions* is the sum of a firm's Scope 1 and Scope 2 carbon emissions. *WSJ CC News Index* is a time-series index developed in Engle et al. (2020) that captures climate change news in the *Wall Street Journal*. We divide the coefficient on *WSJ Climate Change News Index* by 100. The regressions control for  $Log(Assets)$ ,  $Debt/Assets$ ,  $Cash/Assets$ ,  $PP\&E/Assets$ ,  $EBIT/Assets$ ,  $CAPEX/Assets$ , and  $R\&D/Assets$  (all in  $t - 1$ ). In Panel B, we do not include time-varying industry fixed effects, as *WSJ CC News Index* varies only in the time series. Standard errors, clustered at the industry-year level, are in parentheses. Table A.1 defines all variables in detail. \* $p < 0.1$ ; \*\* $p < 0.05$ ; \*\*\* $p < 0.01$ .

Panel A: Carbon Emissions

	$CCExposure_{i,t}$ (1)	$CCExposure_{i,t}^{Opp}$ (2)	$CCExposure_{i,t}^{Reg}$ (3)	$CCExposure_{i,t}^{Phy}$ (4)
$Log(1 + Total\ Emissions_{i,t-1})$	0.169*** (0.023)	0.036*** (0.009)	0.023*** (0.003)	-0.000 (0.001)
Model	OLS	OLS	OLS	OLS
Sample	All	All	All	All
Controls	Yes	Yes	Yes	Yes
Industry $\times$ Year Fixed Effects	Yes	Yes	Yes	Yes
Industry Fixed Effects	No	No	No	No
Country Fixed Effects	Yes	Yes	Yes	Yes
$N$	30,905	30,905	30,905	30,905
Adj. $R^2$	0.390	0.267	0.145	0.035

Panel B: Public Attention to Climate Change

	$CCExposure_{i,t}$ (1)	$CCExposure_{i,t}^{Opp}$ (2)	$CCExposure_{i,t}^{Reg}$ (3)	$CCExposure_{i,t}^{Phy}$ (4)
$WSJ\ CC\ News\ Index_t$	0.427** (0.168)	0.154* (0.089)	0.034*** (0.010)	0.002 (0.004)
Model	OLS	OLS	OLS	OLS
Sample	All	All	All	All
Controls	Yes	Yes	Yes	Yes
Industry $\times$ Year Fixed Effects	No	No	No	No
Industry Fixed Effects	Yes	Yes	Yes	Yes
Country Fixed Effects	Yes	Yes	Yes	Yes
$N$	54,824	54,824	54,824	54,824
Adj. $R^2$	0.298	0.185	0.090	0.024

**Table V**  
**Variance Decomposition of Climate Change Exposure Measures**

This table provides a variance decomposition of the climate change exposure measures. Regressions are estimated at the firm-year level. In Panel A, the table reports the incremental  $R^2$  from adding a specific fixed effect. In Panel B, the table decomposes the variation into a firm fixed effect and a residual component.  $CCExposure$  measures the relative frequency with which climate change bigrams occur in earnings calls.  $CCExposure^{Opp}$  measures the relative frequency with which bigrams that capture opportunities related to climate change occur in earnings calls.  $CCExposure^{Reg}$  measures the relative frequency with which bigrams that capture regulatory shocks related to climate change occur in earnings calls.  $CCExposure^{Phy}$  measures the relative frequency with which bigrams that capture physical shocks related to climate change occur in earnings calls. For all measures, we average values of the four earnings calls during the year. Table A.1 defines all variables in detail.

	$CCExposure_{i,t}$	$CCExposure_{i,t}^{Opp}$	$CCExposure_{i,t}^{Reg}$	$CCExposure_{i,t}^{Phy}$
	(1)	(2)	(3)	(4)
Panel A: Incremental $R^2$				
Year Fixed Effect	0.7%	0.7%	0.5%	0.05%
Industry Fixed Effect	27.1%	16.9%	7.8%	2.0%
Industry $\times$ Year Fixed Effect	1.9%	2.6%	1.4%	1.5%
Country Fixed Effect	0.6%	0.7%	0.4%	0.3%
"Firm Level"	69.7%	79.1%	89.9%	96.2%
Sum	100.0%	100.0%	100.0%	100.0%
Panel B: Fraction of Variation				
Firm Fixed Effect:				
Permanent differences across firms				
within sector and countries		51.6%	56.4%	44.7%
Residual:				45.1%
Variation over time in the identity				
of firms within industries and countries				
most affected by exposure variable		48.4%	43.7%	55.3%
Sum		100.0%	100.0%	100.0%

variation, yielding an incremental  $R^2$  below 1% for each measure. For industry fixed effects, the same observation holds only for  $CCExposure^{Phy}$ . In contrast, exposures to opportunity or regulatory shocks have a sizeable industry component (17% and 8%, respectively), which might stem from regulation targeting specific industries or technological developments affecting entire sectors. The interaction between industry and time fixed effects accounts for, at most, an additional 2.6% of the variation (in the case of  $CCExposure^{Opp}$ ). Country fixed effects provide little additional explanatory power, which mitigates concerns that our measures are strongly affected by the native language in a country or how distant this language is from English. Depending on the measure, between 70% and 96% of the variation is *unexplained* by these sets of fixed effects. Thus, variation plays out at the firm level, rather than at the level of the country, industry, or over time. (The high unexplained variation for  $CCExposure^{Phy}$

is unsurprising given that exposure to physical shocks depends highly on the location of a firm's production sites or insurance policies.) Adding firm fixed effects, permanent differences across firms in an industry and country account for 52%, 56%, 45%, and 45% of the variation of  $CCExposure$ ,  $CCExposure^{Opp}$ ,  $CCExposure^{Reg}$ , and  $CCExposure^{Phy}$ , respectively. The remaining 48%, 44%, 55%, and 55%, respectively, come from variation over time in the identity of firms in industries and countries most affected by the respective climate change variables.

### B. Assessing Measurement Error

We interpret the large share of variance within the firm-year as capturing economically meaningful heterogeneity. Under this view, a firm's idiosyncratic exposure to climate change is the key determinant of the measured variation. A plausible alternative explanation is that part of the firm-level variation reflects idiosyncratic measurement error. We conduct several tests that dispel this alternative. First, we note that we find robust associations between  $CCExposure$  and important real and financial outcomes (as do other papers). These findings suggest that the variation reflected in firm-level  $CCExposure$  is not simply noise.

Second, following Hassan et al. (2019), we quantify the amount of measurement error contained in the firm-level variation by assuming that a firm's "true" exposure follows a first-order autoregressive (AR) process. We then assume that  $CCExposure$  measures this true exposure with classical (i.i.d.) measurement error.<sup>26</sup> Suppose a valid instrument for (lagged)  $CCExposure_{i,t-1}$  were available. In this case, one could back out the share of its variation consisting of measurement error by comparing the OLS and instrumental variable (IV) coefficients. Intuitively, the idea is that candidate IVs measure true climate change exposure with error. Under the i.i.d. assumption, the measurement error in the IV is uncorrelated with that in  $CCExposure_{i,t}$  and thus can be used to "purge" the latter's measurement error. For this procedure to work, we do not assume that the IV has lower measurement error—indeed, it is likely to have higher measurement error. We assume only that the measurement error in the IV and in measured climate change exposure are statistically independent.

Table VI shows three implementations of this idea. One implementation uses an alternative exposure measure constructed by applying our algorithm to the "Management Discussion and Analysis" (MD&A) section in firms' annual 10K filings. The two other implementations use lags of this alternative measure

<sup>26</sup> Under these assumptions, if the correlation between two different lags of the firm-year data is known, the AR(1) parameter and the estimated measurement error can be backed out. For example, if the first lag has a correlation of 0.45 ( $=0.5*0.9$ ) and the second lag a correlation of 0.41 ( $=0.5*0.9*0.9$ ), that would imply measurement error of 50% of the variation and an AR coefficient of 0.9. If the first lag has a correlation of 0.9 and the second 0.8, this implies no measurement error and an AR coefficient of 0.9.

Table VI  
Quantifying Measurement Error in Climate Change Exposure Measures

This table shows AR(1) regressions of climate change exposure. Regressions are estimated at the firm-year level.  $CCExposure$  measures the relative frequency with which climate change bi-grams occur in earnings calls. We average values of the four earnings calls during the year.  $CCExposure^{10K}$  measures climate change exposure by applying our algorithm to the “Management Discussion and Analysis” (MD&A) section in firms’ annual 10K filings. Following Hassan et al. (2019),  $CCExposure$  and  $CCExposure^{10K}$  in this table are standardized by subtracting the sample mean and dividing by the sample standard deviation. *Implied Share Measurement Error* is calculated as  $1 - (\hat{\beta}_{OLS}/\hat{\beta}_{IV})$ , where  $\hat{\beta}_{OLS}$  is the estimated coefficient in  $CCExposure_{i,t} = \alpha + \beta CCExposure_{i,t-1} + \epsilon$  and  $\hat{\beta}_{IV}$  is the coefficient on the instrumented  $CCExposure_{i,t}$  in the same specification. To obtain bootstrapped standard errors for *Implied Share Measurement Error*, we repeat the following procedure 500 times: draw a random sample of the same sample size (with replacement and clustered by firm) from our regression sample, run the two regressions, and obtain the implied share of measurement error. These standard errors are clustered at the firm level. Table A.1 defines all variables in detail. \* $p < 0.1$ ; \*\* $p < 0.05$ ; \*\*\* $p < 0.01$ .

Panel A: Overall Variation

	$CCExposure_{i,t}$			
	(1)	(2)	(3)	(4)
$CCExposure_{i,t-1}$	0.922*** (0.002)	1.008*** (0.003)	0.991*** (0.003)	0.958*** (0.002)
Model	OLS	IV	IV	IV
Instrument		$CCExposure^{10K}_{i,t-1}$	$CCExposure^{10K}_{i,t-2}$	$CCExposure_{i,t-2}$
Sample	U.S.	U.S.	U.S.	U.S.
Industry × Year Fixed Effects	No	No	No	No
N	47,589	47,589	41,794	41,794
Implied Share Measurement Error		0.085 (0.007)	0.069 (0.007)	0.037 (0.005)

Panel B: Firm-Level Variation

	$CCExposure_{i,t}$			
	(1)	(2)	(3)	(4)
$CCExposure_{i,t-1}$	0.886*** (0.002)	0.992*** (0.004)	0.966*** (0.002)	0.932*** (0.003)
Model	OLS	IV	IV	IV
Instrument		$CCExposure^{10K}_{i,t-1}$	$CCExposure^{10K}_{i,t-2}$	$CCExposure_{i,t-2}$
Sample	U.S.	U.S.	U.S.	U.S.
Industry × Year Fixed Effects	Yes	Yes	Yes	Yes
N	47,502	47,502	41,712	41,712
Implied Share Measurement Error		0.107 (0.002)	0.083 (0.012)	0.050 (0.007)

and *CCExposure* itself as instruments. While the estimates of the share of measurement error in *CCExposure* vary somewhat across the three approaches, approximately 5% to 10% of the variation in measured *CCExposure* is due to measurement error.<sup>27</sup> The implied measurement error at the firm level (in Panel B) is about 2 percentage points higher than in the overall variation (Panel A). Although we interpret these results with due caution, they suggest that measurement error in the firm-level dimension is higher than in the overall panel, but only modestly. Thus, concerns that the variation displayed at the firm level is subject to more measurement error than the overall climate change exposure measure (before any fixed effects) are not substantiated.

## V. Economic Applications

### A. Real Outcomes: Green-Tech Jobs and Green Patents

Significant climate-related innovation is required to reach net-zero emissions by 2050 (Stern and Valero (2021)), implying huge investments by firms in human capital and R&D. According to some estimates, incremental investments of \$50 trillion are needed in solar technology, decarbonization, energy efficiency, or carbon capture by 2050 (World Economic Forum (2021)). To illustrate that our exposure measures help predict real outcomes related to the net-zero transition, we relate next year's creation of disruptive green-tech jobs and green patents to this year's values of climate change exposure. Among the sampled U.S. firms, for firm  $i$  and year  $t$  we estimate

$$\text{Green Outcome}_{i,t+1} = \exp(\alpha_i + \beta \log(1 + \text{CCExposure}_{i,t}) + \gamma \mathbf{X}_{i,t} + \delta_j \times \delta_t + \epsilon_{i,t+1}), \quad (5)$$

where *Green Outcome* <sub>$i,t+1$</sub>  is #*Green-Tech Jobs* <sub>$i,t+1$</sub>  or #*Green Patents* <sub>$i,t+1$</sub>  in year  $t + 1$  and *CCExposure* <sub>$i,t$</sub>  is the climate change exposure measure in year  $t$  (we include the overall and topic-based measures). The vector  $\mathbf{X}_{i,t}$  includes *Log(Assets)*, *Debt/Assets*, *Cash/Assets*, *PP&E/Assets*, *EBIT/Assets*, *CAPEX/Assets*, and *R&D/Assets*. The variables  $\delta_j \times \delta_t$  represent industry-year fixed effects. We account for industry shocks that vary over time, as firm-level innovation-related activity contains a large time-varying industry component (Aghion et al. (2005)). As demonstrated in Table V, such variation is also an important determinant of climate change exposure, making it important to identify effects *within* industry-year pairs. We cluster standard errors at the industry-year group level.

We estimate equation (5) using Poisson regressions, which have two advantages (Cohn, Liu, and Wardlaw (2022)). First, Poisson regressions account for the distributional characteristics of our count-based outcomes (they provide unbiased estimates for dependent variables with a large mass of values at zero combined with severe skewness). Second, Poisson regressions allow use to

<sup>27</sup> These estimates compare favorably to the amount of measurement error found using similar assumptions in firm-level variables measured using accounting data (e.g., measures of total factor productivity constructed by Bloom et al. (2018) and Collard-Wexler (2011)).

include industry-year fixed effects without biasing the estimation. They thus address the issue of separable group fixed effects (in our case at the industry-year level) by basing the estimation only on observations with at least one nonzero value within a group. This is desirable, as it restricts the usable sample to those groups that are informative about the effects of *CCExposure*.<sup>28</sup> For robustness, we also estimate linear and log1plus-linear models (with and without industry-year fixed effects) on the unrestricted sample (we interpret these models' estimates with caution).

The estimation results for *#Green-Tech Jobs* are reported in Table VII. In column (1), the estimates show that firms with higher overall exposure post more vacancies for jobs in disruptive green technologies over the subsequent year. A one-standard-deviation increase in *CCExposure* is associated with a 109% increase in the number of green-tech jobs over the next year.<sup>29</sup> Columns (2) to (4) consider the topic-based measures. As expected, the overall exposure effect is due in large part to high-opportunity firms (column (2)). Firms with higher regulatory exposure also plan to hire more green-tech workers than firms with lower exposure (column (3)). We do not find that firms with larger physical exposure post more green-tech jobs (column (4)). In column (5), we continue to find that *CCExposure* positively predicts green-tech hiring if we replace *#Green-Tech Jobs* with *I(Green-Tech Jobs)*, an indicator for whether a firm posts a green-tech job (we estimate a linear model with the same observations as in columns (1) to (4)). Similarly, in column (6) estimates are robust to using the ratio of green-tech jobs to all tech jobs (*Green-Tech Ratio*). Column (7) addresses the concern that high-exposure firms may simply recruit more personnel in disruptive technologies across the board, without a specific focus on green jobs per se (for example, because these firms happen to be more innovative). Specifically, we replace *#Green-Tech Jobs* with *#Nongreen-Tech Jobs* and reestimate the regression in column (1). We do not find positive predictive effects of the exposure measure, which mitigates concerns of spurious relationships. In fact, firms with higher climate change exposure hire less, not more, nongreen-tech jobs. Overall, the data are more consistent with a recruiting shift from nongreen-tech jobs to green-tech jobs, rather than a general expansion of tech-related hiring at high-exposure firms.

The results for green-tech jobs broadly extend to green patents in Table VIII. In columns (1) to (3), firms with greater climate change exposure show more green patenting in the next year. A one-standard-deviation

<sup>28</sup> Cohn, Liu, and Wardlaw (2022) show that log1plus-linear models may be biased in our context. The admission of separable group fixed effects in Poisson regressions differs from that in other nonlinear count-data models. These alternative models are subject to the incidental parameter problem, which leads to biased and inconsistent estimates (Lancaster (2000)).

<sup>29</sup> In a Poisson model, for a regression coefficient  $\beta$ , the magnitude of a one-standard-deviation change in the independent variable is calculated as  $e^{\beta \times STD} - 1$ . This effect size (when multiplied by 100%) represents the percentage change in the dependent variable. We use the within-fixed-effects (rather than overall-panel) standard deviation to capture plausible variation. The large magnitude of the effect also indicates that the average number of disruptive green-tech jobs is relatively low.



Table VII  
Green-Tech Jobs and Climate Change Exposure Measures

This table reports regressions that relate green-tech jobs to the climate change exposure measures. Regressions are estimated at the firm-year level. *#Green-Tech Jobs* is the number of job postings for disruptive green-tech jobs. *I(Green-Tech Jobs)* is an indicator that equals one if *#Green-Tech Jobs* is positive, and zero otherwise. *#Nongreen-Tech Jobs* is the number of job postings for nongreen disruptive tech jobs. *Green-Tech Ratio*<sub>*i,t*+1</sub> is the number of job postings for disruptive green jobs relative to the total number of all disruptive job postings. *CCEexposure*, *CCEexposure*<sub>*Opp*</sub>, *CCEexposure*<sup>Reg</sup>, and *CCEexposure*<sup>Phy</sup> are defined as in previous tables. The regressions control for *Log(Assets)*, *Debt/Assets*, *PP&E/Assets*, *EBIT/Assets*, *CAPEX/Assets*, and *R&D/Assets* (all in *t*). In columns (5) to (7), we use the same observations as in columns (1) to (4). In columns (1) to (4) and (7), the economic effects are computed as the percentage change in the dependent variable for a one-standard-deviation change in the exposure variable of interest. In columns (5) and (6), the economic effect is computed as the effect of a one-standard-deviation change in the exposure variable relative to the standard deviation of the dependent variable. We use the within-fixed-effect standard deviations. Standard errors, clustered at the industry-year level, are in parentheses. Table A.1 defines all variables in detail. \**p* < 0.1; \*\**p* < 0.05; \*\*\**p* < 0.01.

	#Green-Tech Jobs <sub><i>i,t</i>+1</sub>				<i>I</i> (#Green-TechJobs) <sub><i>i,t</i>+1</sub>	<i>Green-TechRatio</i> <sub><i>i,t</i>+1</sub>	#Nongreen-TechJobs <sub><i>i,t</i>+1</sub>
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
<i>Log</i> (1 + <i>CCEexposure</i> <sub><i>i,t</i></sub> )	1.564*** (0.199)				0.077*** (0.006)	0.015*** (0.003)	−0.204*** (0.060)
<i>Log</i> (1 + <i>CCEexposure</i> <sub><i>i,t</i></sub> <sup><i>Opp</i></sup> )		1.833*** (0.229)					
<i>Log</i> (1 + <i>CCEexposure</i> <sub><i>i,t</i></sub> <sup><i>Reg</i></sup> )			1.458*** (0.445)				
<i>Log</i> (1 + <i>CCEexposure</i> <sub><i>i,t</i></sub> <sup><i>Phy</i></sup> )				1.079 (1.217)			
Model	Poisson	Poisson	Poisson	Poisson	OLS	OLS	Poisson
Sample	U.S.	U.S.	U.S.	U.S.	U.S.	U.S.	U.S.
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry x Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	23,870	23,870	23,870	23,870	23,870	23,870	23,870
Adj./ps. <i>R</i> <sup>2</sup>	0.754	0.767	0.687	0.684	0.116	0.049	0.526
Dep. Variable: Mean	2.82	2.82	2.82	2.82	0.07	0.003	845.09
Dep. Variable: STD	89.56	89.56	89.56	89.56	0.26	0.042	3613.42
Economic Effect, %	108.7	79.5	20.0	6.8	14.0	16.9	−9.1

Table VIII  
Green Patents and Climate Change Exposure Measures

This table reports regressions that relate green patents to the climate change exposure measures. Regressions are estimated at the firm-year level. #Green Patents is the number of green patents.  $I(\text{Green Patents})$  is an indicator that equals one if #Green Patents is positive, and zero otherwise.  $\text{Green Patents Ratio}_{i,t+1}$  is the number of green patents relative to the total number of all patents. #Nongreen Patents is the number of nongreen patents.  $\text{CCExposure}^{\text{Opp}}$ ,  $\text{CCExposure}^{\text{Reg}}$ , and  $\text{CCExposure}^{\text{Phy}}$  are defined as in previous tables. The regressions control for  $\text{Log}(\text{Assets})$ ,  $\text{Debt}/\text{Assets}$ ,  $\text{Cash}/\text{Assets}$ ,  $\text{PP\&E}/\text{Assets}$ ,  $\text{EBIT}/\text{Assets}$ ,  $\text{CAPEX}/\text{Assets}$ , and  $\text{R\&D}/\text{Assets}$  (all in  $t$ ). In columns (5) to (7), we use the same observations as in columns (1) to (4) (the Poisson estimation in Column (7) drops some observations). In columns (1) to (4) and (7), the economic effects are computed as the percentage change in the dependent variable for a one-standard-deviation change in the exposure variable of interest. In columns (5) and (6), the economic effect is computed as the effect of a one-standard-deviation change in the exposure variable relative to the standard deviation of the dependent variable. We use the within-fixed-effect deviations. Standard errors, clustered at the industry-year level, are in parentheses. Table A.1 defines all variables in detail. \* $p < 0.1$ ; \*\* $p < 0.05$ ; \*\*\* $p < 0.01$ .

	#Green Patents <sub>i,t+1</sub>				<i>I</i> (Green Patents) <sub>i,t+1</sub>	Green Patents Ratio <sub>i,t+1</sub>	#Nongreen Patents <sub>i,t+1</sub>
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Log(1 + CCE_{i,t}^{exposure})$	1.102*** (0.231)				0.025*** (0.003)	0.006*** (0.001)	−0.436*** (0.118)
$Log(1 + CCE_{i,t}^{exposure_{Opp}})$		0.854*** (0.312)					
$Log(1 + CCE_{i,t}^{exposure_{Reg}})$			3.061*** (0.272)				
$Log(1 + CCE_{i,t}^{exposure_{Phy}})$				−1.155 (2.865)			
Model	Poisson	Poisson	Poisson	Poisson	OLS	OLS	Poisson
Sample	U.S.	U.S.	U.S.	U.S.	U.S.	U.S.	U.S.
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Industry × Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	21,914	21,914	21,914	21,914	21,914	21,914	21,776
Adj./ps. <i>R</i> <sup>2</sup>	0.617	0.603	0.614	0.598	0.078	0.023	0.752
Dep. Variable: Mean	0.28	0.28	0.28	0.28	0.03	0.003	22.10
Dep. Variable: STD	4.07	4.07	4.07	4.07	0.18	0.040	224.23
Economic Effect, %	71.7	32.0	47.3	−6.9	7.0	7.4	−19.3

increase in  $CCExposure$  is associated with a 72% increase in the number of green patents over the next year. The effect for  $CCExposure^{Opp}$  is intuitive, as green innovation provides business opportunities during the net-zero transition. To illustrate the intuition behind the effects for  $CCExposure^{Reg}$ , the case of Caterpillar is insightful. This firm is not only the top green patent producer in our sample (see Section I.E), but it also exhibits high measured regulatory exposure. This latter feature stems from its legacy business related to mining and diesel engines (sample mean of  $CCExposure^{Reg}_{Caterpillar} = 0.09$ , in the top decile of  $CCExposure^{Reg}$ ). We do not find that firms with larger physical exposure generate more green patents (column (4)). In columns (5) and (6), we continue to find that  $CCExposure$  predicts green patenting if we replace  $\#Green\ Patents$  with an indicator for whether a firm created green patents (column (5)) or with the green patents ratio as in Cohen, Gurun, and Nguyen (2021) (column (6)). Column (7) shows that high-exposure firms are not simply more innovative in general; the estimates indicate fewer, not more, nongreen patents by firms with high values of  $CCExposure$ .

Table IA.XIII shows that the results in Tables VII and VIII are robust to controlling for carbon emissions. This finding demonstrates that our measures contain additional information beyond what is reflected in emissions (the sample size is reduced in the panel due to the lower number of observations on carbon emissions).

In Table IA.XIV, a series of alternative specifications continue to show that  $CCExposure$  predicts green-tech job creation. In column (1), we dispel concerns related to strategic disclosure in earnings calls (Mayew (2008), Hassan et al. (2019)). One specific potential concern is that managers may want to distract attention from poor performance and strategically “cheap talk” about climate change (Hail, Kim, and Zhang (2021)). Following Hassan et al. (2019), we test for this possibility by adding a control for the firm’s overall sentiment (share of positive and negative tone words across the earnings call transcript) and two proxies for recent performance.<sup>30</sup> The estimates show that our results are robust to adding these controls. In column (2), we restrict the sample to firm-years within the BG database to ensure that the results are unaffected by how we classify the firms missing in BG; recall that we assume no green-tech job creation for these firms (BG may systematically miss scraping some firms’ postings). In column (3), exposure is based on a count of bigrams in the Q&A session, that is, the part of the call that is less under management control and in turn less subject to concerns of strategic (non)disclosure and greenwashing. In column (4),  $CCSentiment^{Pos}$  strongly predicts next-year green-tech job creation, while  $CCSentiment^{Neg}$  is insignificant (albeit marginally). In column (5),  $CCRisk$  is positively associated with green-tech job creation. In column (6) to

<sup>30</sup> We measure performance as the precall stock return accumulated over the seven days prior to the earnings call and the earnings surprise. Earnings surprise is defined as earnings per share before extraordinary items minus earnings per share in the same quarter of the prior year, divided by the price per share at the beginning of the quarter (Ball and Bartov (1996)). We average the two variables across the earnings calls of a firm-year to obtain an annual measure.

(9), results hold if we estimate OLS specifications to address potential concerns with the Poisson specification. We estimate models with and without industry-year fixed effects, and with *#Green-Tech Jobs* or  $\text{Log}(1 + \text{\#Green-Tech Jobs})$ . We also provide estimates that replace the log1plus version of *CCEXposure* with an unlogged version. Table IA.XV applies the same alternative specifications to green patenting. The estimates show that our results continue to hold.

Table IA.XVI reports regressions for the subsamples in which the exposure measures that rely exclusively on the initial bigrams indicate zero exposure. In these estimations, our exposure measures continue to predict green outcomes. This finding corroborates the performance gain from using more subtle and less visible climate change bigrams, as the estimation is identified from the bigrams obtained through the keyword search algorithm.

Finally, Table IA.XVII documents the covariate balance of observations that are either included or excluded from the estimations in Tables VII and VIII. Excluded firm-years exhibit lower climate change exposure, implying that our estimates are obtained within the set of firms for which climate change issues are most pressing.

## B. Financial Market Outcomes

### B.1. Options Market Risks and Risk Premiums

Firms with higher regulatory exposure are more strongly affected by future regulations to combat global warming, and uncertainty over such regulations should be priced in the options market (Kelly, Pastor, and Veronesi (2016)). Likewise, climate opportunities are risky, with plenty of uncertainty surrounding investments in green technologies or renewable energy. We therefore test whether climate change exposure is related to option-implied risks and risk premiums. We consider three sets of risk variables. First, to quantify general risks, we use three implied central moments, namely, variance (*IVar*), skewness (*ISkew*), and kurtosis (*IKurt*). Second, we calculate two heuristic measures quantifying the relative expensiveness of protection against left (*SlopeD*) and right (*SlopeU*) tail risks.<sup>31</sup> Third, we use the variance risk premium (*VRP*) to measure the premiums that investors are willing to pay to hedge against general climate-related variance risk (or uncertainty, as suggested in Bali and Zhou (2016)). Using each of these variables, we run the regression:

$$OI\ Outcome_{i,t+1} = \alpha_i + \beta \text{Log}(1 + CCEXposure_{i,t}) + \gamma \mathbf{X}_{i,t} + \delta_j \times \delta_t + \epsilon_{i,t+1}, \quad (6)$$

where *OI Outcome*<sub>*i,t+1*</sub> is an option-implied measure for firm *i* measured at the end of quarter *t* (i.e., a conditional expectation of some quantity over the period

<sup>31</sup> The variable *SlopeD* increases when the cost of left-tail protection goes up (relative to the cost of at-the-money [ATM] options), and *SlopeU* decreases (becomes more negative) when the relative cost of obtaining upside growth increases. Note that Sautner et al. (2022) define their measure of *SlopeU* as minus one times *SlopeU*.

$t + 1$ ), and  $CCExposure_{i,t}$  is firm  $i$ 's climate change exposure in quarter  $t$ .<sup>32</sup> The vector  $\mathbf{X}_{i,t}$  includes the same controls as before (delayed to be available in the third quarter after the annual close of the fiscal period). The variables  $\delta_j$  and  $\delta_t$  represent industry and year fixed effects, respectively. We cluster standard errors at the industry-year level.

Table IX, Panel A, documents that  $CCExposure$  is strongly linked to forward-looking risks and risk premiums. In columns (2) and (3),  $CCExposure$  predicts a more negatively skewed return distribution ( $ISkew$ ) and fatter tails ( $IKurt$ ). Furthermore, tail exposure is more costly for firms with higher climate change exposure. More specifically, downside protection in column (4) (positive and significant coefficient on  $SlopeD$ ) and upside potential in column (5) (negative and significant coefficient on  $SlopeU$ ) become more expensive when  $CCExposure$  is higher. In terms of magnitudes, the effects are strongest in column (3) for  $IKurt$ . A one-standard-deviation change in  $CCExposure$  is associated with a change in  $IKurt$  equivalent to 7% of its standard deviation. The effects for  $SlopeD$  and  $SlopeU$  are 4.5% and 4.1%, respectively.

The three remaining panels consider the topic-based measures. Earnings calls should contain more discussions of climate-related opportunities if a firm is well positioned for the growth potential arising from climate change. The realization of these opportunities could lead to large gains if successful and to large losses if unsuccessful. Investors may in turn trade in the options market to reflect the two-sided effects of climate opportunities. Panel B confirms this intuition: the tail effects for  $CCExposure^{Opp}$  in columns (4) and (5) are similar compared to the corresponding estimates in Panel A. The magnitude of a one-standard-deviation increase in  $CCExposure^{Opp}$  is 4.3% for  $SlopeD$  and 3.9% for  $SlopeU$ , respectively. Thus, it is not only the case that options are more expensive on both tails if climate opportunities are higher, but also that the cost of upside potential grows faster than the cost of downside crash protection. The link between  $CCExposure^{Opp}$  and  $VRP$  in column (6) demonstrates that the wedge between the implied and "historically fair" price of out-of-the-money (OTM) calls increases with opportunity exposure. Thus, investors are ready to pay an extra (volatility) premium when buying options on stocks with climate-related upside potential. However, the effect is small in magnitude and only marginally significant.

In Panel C, the pattern for  $CCExposure^{Reg}$  is similar to that for  $CCExposure^{Opp}$ , though the magnitudes are smaller. While the right-tail option expensiveness increases by 2.6% of its standard deviation (i.e.,  $SlopeU$  diminishes) for a one-standard-deviation change in  $CCExposure^{Reg}$ , the crash protection grows by 2.3%. This confirms our earlier evidence that some firms with high regulatory exposure face downside risks and upside potential due to

<sup>32</sup> When computing quarterly versions of our measures, we encounter the issue that any specific earnings call in a year might not discuss climate change, even though the conversation turns to the issue in surrounding calls. These incidental gaps in the quarterly data (where the measured  $CCExposure = 0$ ) do not reflect business realities. Therefore, we preprocess the quarterly climate change exposure following a method outlined in Sautner et al. (2022), which exponentially smooths each metric for each firm with a half-life of three quarters.

Table IX  
Forward-Looking Risk Measures and Climate Change Exposure Measures

This table reports regressions that relate forward-looking risk measures to the climate change exposure measures. Regressions are estimated at the firm-quarter level. *IVar* is implied variance, *ISkew* is implied skewness, *IKurt* is implied kurtosis, *SlopeD* and *SlopeU* are implied volatility slopes on the left and right of the distribution, and *VRP* is the variance risk premium. Construction of the option-implied measures is detailed in Section II of the Internet Appendix. *CCExposure*, *CCExposure<sup>Opp</sup>*, *CCExposure<sup>Reg</sup>*, and *CCExposure<sup>Phy</sup>* are defined as in previous tables. The regressions control for *Log(Assets)*, *Debt/Assets*, *Cash/Assets*, *PP&E/Assets*, *EBIT/Assets*, *CAPEX/Assets*, and *R&D/Assets* (all in *t*). The economic effect is computed as the effect of a one-standard-deviation change in the exposure variable of interest relative to the standard deviation of the dependent variable (in %). We use the within-fixed-effect standard deviation. Standard errors, clustered at the industry-year level, are in parentheses. Table A.1 defines all variables in detail. \**p* < 0.1, \*\**p* < 0.05, \*\*\**p* < 0.01.

	<i>IVar</i> <sub><i>i,t+1</i></sub> (1)	<i>ISkew</i> <sub><i>i,t+1</i></sub> (2)	<i>IKurt</i> <sub><i>i,t+1</i></sub> (3)	<i>SlopeD</i> <sub><i>i,t+1</i></sub> (4)	<i>SlopeU</i> <sub><i>i,t+1</i></sub> (5)	<i>VRP</i> <sub><i>i,t+1</i></sub> (6)
Panel A: <i>CCExposure</i>						
<i>Log(1 + CCExposure</i> <sub><i>i,t</i></sub> <i>)</i>	−0.002 (0.005)	−0.049*** (0.009)	0.303*** (0.049)	0.033*** (0.007)	−0.026*** (0.006)	0.003 (0.002)
<i>N</i>	42,093	42,093	42,093	42,093	42,093	42,089
Adj. <i>R</i> <sup>2</sup>	0.424	0.140	0.349	0.231	0.236	0.094
Economic Effect, %	−0.42	−4.57	7.01	4.46	−4.14	0.89
Panel B: <i>CCExposure<sup>Opp</sup></i>						
<i>Log(1 + CCExposure<sup>Opp</sup></i> <sub><i>i,t</i></sub> <i>)</i>	0.004 (0.009)	−0.053*** (0.012)	0.403*** (0.067)	0.048*** (0.011)	−0.037*** (0.010)	0.006* (0.003)
<i>N</i>	42,093	42,093	42,093	42,093	42,093	42,089
Adj. <i>R</i> <sup>2</sup>	0.424	0.140	0.348	0.231	0.236	0.094
Economic Effect, %	0.56	−3.27	6.18	4.30	−3.91	1.19
Panel C: <i>CCExposure<sup>Reg</sup></i>						
<i>Log(1 + CCExposure<sup>Reg</sup></i> <sub><i>i,t</i></sub> <i>)</i>	−0.007 (0.014)	−0.075*** (0.024)	0.453*** (0.146)	0.054** (0.027)	−0.053*** (0.019)	0.005 (0.008)
<i>N</i>	42,093	42,093	42,093	42,093	42,093	42,089
Adj. <i>R</i> <sup>2</sup>	0.424	0.139	0.346	0.230	0.235	0.094
Economic Effect, %	−0.46	−2.19	3.28	2.28	−2.64	0.47
Panel D: <i>CCExposure<sup>Phy</sup></i>						
<i>Log(1 + CCExposure<sup>Phy</sup></i> <sub><i>i,t</i></sub> <i>)</i>	−0.033 (0.020)	−0.083 (0.059)	1.336*** (0.319)	0.145*** (0.048)	−0.175*** (0.048)	−0.012 (0.011)
<i>N</i>	42,093	42,093	42,093	42,093	42,093	42,089
Adj. <i>R</i> <sup>2</sup>	0.424	0.139	0.347	0.230	0.236	0.094
Economic Effect, %	−1.02	−1.13	4.51	2.85	−4.06	−0.52

(Continued)



Table IX—Continued

Model	OLS	OLS	OLS	OLS	OLS	OLS
Sample	S&P500	S&P500	S&P500	S&P500	S&P500	S&P500
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Industry $\times$ Year Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Dep. Variable: Mean	0.176	−0.571	4.678	0.317	−0.101	0.042
Dep. Variable: STD	0.199	0.453	1.823	0.312	0.265	0.142

their green innovation activity. In Panel D, the effects for  $CCExposure^{Phy}$  are similar to those of the other measures.<sup>33</sup>

Overall, climate change exposure is priced in the options market. Considering all the evidence, stocks with higher exposure have probability mass shifted to the tails of the distribution, making crash protection and upside potential relatively more expensive. Obtaining protection and upside growth potential comes at a premium, which increases more strongly for firms facing higher opportunities. We acknowledge that the effect magnitudes are modest and hardly tradeable after transaction costs.

### B.2. Cross Section of Stock Returns

Climate change exposure is related to risks and risk premiums in the options market. Consequently, systematic risk related to  $CCExposure$  may be associated with a risk premium in the cross section of returns. That said, testing for the pricing effects of a climate change exposure factor, labeled  $CCEXPOSURE$ , is challenging for several reasons. A conceptual challenge arises because return effects are theoretically more ambiguous to predict compared to the risk measures. On the one hand, firms with high betas for  $CCEXPOSURE$  should be more risky and—in expectation—earn a risk premium.<sup>34</sup> On the other hand, the relations may actually be the opposite, with risks gradually getting priced in during the sample period; as risks emerge, stock prices decline, implying lower realized returns. Pastor, Stambaugh, and Taylor (2021) illustrate this difference between ex ante and ex post returns. An estimation challenge arises because  $CCExposure$  reflects the attention devoted to climate topics at a point in time. This implies that the pricing of  $CCEXPOSURE$  should vary over time, requiring the estimation of conditional risk premiums. Another

<sup>33</sup> Our inference for the pricing of physical exposure is different from the link between hurricane uncertainty and variance pricing in Kruttli, Roth Tran, and Watugala (2021). For example, while we concentrate on the unconditional pricing using the expected  $VRP$ , Kruttli, Roth Tran, and Watugala (2021) study dynamics of the realized  $VRP$ . However, these authors also conclude that, especially in the early sample years, investors underprice variance in options of firms strongly exposed to extreme weather events.

<sup>34</sup> For example, such firms face higher uncertainty related to future developments in climate-related areas, that is, their valuation should include real option value depending on the path of climate-related technologies, regulations, or physical climate shifts.

challenge arises because the number of assets for such tests is large relative to the time points available for the estimation—less than 20 years of data.

With these challenges in mind, we investigate the conditional pricing of *CC-EXPOSURE* in the cross section of stocks. We follow Jamilov, Rey, and Tahoun (2021) and construct the factor as an unexpected shock to the aggregate value of *CCExposure*. This involves three primary steps. First, we convert quarterly transcript-level values of *CCExposure*<sub>*i,t*</sub> for U.S.-traded firms to a monthly frequency by propagating the last exposure values for up to three months forward (i.e., we match the month-year of each climate change exposure to the month-year of the respective quarterly transcript). Second, we compute cross-sectional monthly averages of *CCExposure*<sub>*m*</sub>. Third, we take the first differences in these monthly averages as a proxy for unexpected monthly shocks to the aggregate exposure level, and use them as the *CCEXPOSURE* factor.<sup>35</sup>

To examine the conditional pricing of *CCEXPOSURE* among S&P500 firms, we follow Gagliardini, Ossola, and Scaillet (2016, GOS), who provide a conditional extension of the two-pass regression approach (Fama and MacBeth (1973)). We use this approach as it delivers good small-sample performance when—as in our case—the cross section is large relative to the time series. GOS assume a linear conditional factor model for excess returns with time-varying factor exposures and risk premiums. They model the parameters as linear functions of lagged instruments. The factor loadings  $\beta_{i,m}$  depend on stock-specific instruments ( $Z_{i,m-1}$ ) as well as common instruments ( $Z_{m-1}$ ), and the factor expectations only on common instruments. Under this framework, the conditional expected return on stock *i* in month *m* is

$$E[R_{i,m}|Z_{i,m-1}, Z_{m-1}] = \beta_{i,m}^T \lambda_m, \quad (7)$$

where the risk premium  $\lambda_m$  is the sum of the conditional factor expectation  $E[F_m|Z_{m-1}]$  and the process  $v_m$ , estimated from the cross section of stocks. The process  $v_m$  allows the estimated risk premium to deviate from the conditional expectation of a factor due to market imperfections for tradeable factors (Cremers, Petajisto, and Zitzewitz (2013), GOS) and it also reveals an “implicit cost” of projecting a nontradeable factor (like ours) on returns. A similar framework is used, for example, in Barras and Malkhozov (2016). As in GOS, we use as common instruments the term spread and the default spread and as the stock-specific instrument the log of the book-to-market ratio (see Section I.G for definitions). We estimate the time-varying components of the risk premiums with the four-factor model by Carhart (1997) that is augmented with the *CCEXPOSURE* factor.<sup>36</sup>

When performing the estimation, we obtain average conditional risk premiums in line with expectations (risk premiums for the market, size, value, and

<sup>35</sup> The factor is standardized to have zero mean and annual volatility of 10%. Results are robust to using the residuals from an AR(1) process fitted to the monthly exposure series, as implemented in Jamilov, Rey, and Tahoun (2021) (the resulting factors are almost perfectly correlated). However, fitting an AR(1) process may introduce look-ahead bias.

<sup>36</sup> The factor is essentially orthogonal to the other factors, with all unconditional correlations being smaller than 0.05. The results are robust to using three- and five-factor models.

**Table X**  
**Climate Change Exposure Factor: Components of  $F$  and  $\nu$**


This table reports the estimated annualized components of  $F$  and  $\nu$  for the four-factor Carhart (1997) model augmented by a *CCEXPOSURE* factor. The estimation is based on the conditional framework by Gagliardini, Ossola, and Scaillet (2016). The factor is constructed as the monthly change in the cross-sectional average of *CCEXposure* across U.S.-traded sample firms. The factor is standardized to have zero mean and an annual volatility of 10%. All instruments are centered and standardized in the time series. The common instruments are the default spread and the term spread, and the firm-specific instrument is the log of the book-to-market ratio. \* $p < 0.1$ ; \*\* $p < 0.05$ ; \*\*\* $p < 0.01$ .

<i>Factors</i>	<i>Instruments</i>	$F$ (1)	$SE(F)$ (2)	$\nu$ (3)	$SE(\nu)$ (4)
<i>Market</i>	Constant	8.9838***	3.4981	2.3908***	0.7110
	Default Spread	−1.0201	5.4550	2.4676***	0.8715
	Term Spread	−1.9715	3.3962	1.4489**	0.6705
<i>SMB</i>	Constant	2.3669	1.9164	2.6523*	1.3459
	Default Spread	2.5406	2.0404	−1.3983	1.0227
	Term Spread	2.1356	1.8985	−4.6391***	0.9302
<i>HML</i>	Constant	−2.1553	2.0893	−3.5959***	1.0965
	Default Spread	−3.6834	3.9437	3.7360***	0.8545
	Term Spread	4.8748**	2.2504	−0.0444	0.8434
<i>MOM</i>	Constant	1.3199	3.5668	7.2011***	1.6444
	Default Spread	−14.359*	8.2567	7.8356***	1.7552
	Term Spread	2.4766	2.9728	−0.6825	1.2843
<i>CCEXPOSURE</i>	Constant	−0.0032	2.3008	3.7273***	1.1654
	Default Spread	0.0805	2.7644	3.1262***	1.0855
	Term Spread	−0.2941	2.6282	−0.1834	0.9978

momentum factors are 11.4%, 5.0%, −5.8%, and 8.5% per annum (p.a.), respectively). The *CCEXPOSURE* premium is positive, on average (3.7% p.a.), and we obtain positive point estimates for most months. More importantly, the risk premium is not constant over time, and we reject the hypotheses that its two components are constant ( $p$ -values of 0.0137 and 0.0001, respectively).

In Table X, we report the estimated annualized components of the risk premium  $\lambda_m$ , that is, the estimates of  $F$  and  $\nu$ . Similar to the results in GOS, most of the action for the risk premiums comes through the cross-sectional component  $\nu$ . For *CCEXPOSURE*,  $\nu$  has a positive unconditional mean (constant of 3.73%) and a positive link to the default spread (3.13%)—both are highly significant. This indicates that stocks with high exposure to the *CCEXPOSURE* factor are expected to earn higher returns, especially when market-wide default risk increases.

The time series of the estimated risk premium on *CCEXPOSURE* is depicted in Figure 4. The series illustrates significant variability over time, with a large spike around the financial crisis. Further tentative interpretations indicate a temporary spike around the time of Hurricane Sandy (October 2012) and the Doha Climate Summit (November 2012). Another temporary spike occurs just after the Paris Agreement (December 2015). Considering the most recent five years, the risk premium was lowest around the time President Trump took



**Figure 4. Risk premium on the climate change exposure factor.** This figure shows the time series of the risk premium on the *CCEXPOSURE* factor, estimated together with the four-factor Carhart (1997) model using the conditional framework of Gagliardini, Ossola, and Scaillet (2016). The factor is constructed as the monthly change in the cross-sectional average of *CCEXposure* across U.S.-traded sample firms. The factor is standardized to have zero mean and an annual volatility of 10%. (Color figure can be viewed at [wileyonlinelibrary.com](http://wileyonlinelibrary.com))

office (January 2017); it gradually increased thereafter with a drop around the onset of the COVID pandemic.<sup>37</sup>

We emphasize that our objective is not to create an ultimate climate factor to be added to the factor zoo (Feng, Giglio, and Xiu (2020)), but instead to show that attention to climate topics in earnings calls is linked to systematic risk, with shocks to such attention potentially being priced in the cross section (following a narrative as in Shiller (2017)).

## VI. Conclusion

In this paper, we introduce a new method that identifies firm-level climate change exposure from word combinations signaling climate change conversation in earnings conference calls. As these calls reflect the demand side (analysts) and the supply side (management) of a “market for information,”

<sup>37</sup> As in the previous applications, we estimate the risk premiums separately by topic. The topic-based premiums are on average positive, but demonstrate distinct time-series patterns. For example, when the physical risk premium goes up, the opportunity risk premium tends to go down. We do not want to overemphasize the topic-based differences here, as our framework uses the same set of instruments for all topic-based factors.

our measures reflect the combined views of key stakeholders about a firm's climate change exposure. Furthermore, earnings calls are largely forward-looking; while analysts review past results, they also spend much of their time probing management about future plans (Huang et al. (2018)).

Our measures build on recent work that identifies earnings calls as a source for identifying the various risks and opportunities that firms face over time. We adjust the approach of this prior work along several critical dimensions, allowing us to capture aspects of both the opportunities and the (physical and regulatory) risks associated with climate change. For this purpose, we adapt the machine-learning keyword discovery algorithm proposed by King, Lam, and Roberts (2017) to produce several sets of climate change bigrams. Rather than choosing a training library, we start with a short list of initial bigrams that most experts would agree are related to climate change. Our exposure measures capture the proportion of the earnings call related to climate change topics. These measures are available for a global sample of more than 10,000 firms covering the period 2002 to 2020. We demonstrate that our measures are helpful in predicting important real outcomes related to the net-zero transition, notably, green-tech growth and green patenting. We also document that the measures contain information that is priced in the options and equity markets.

# ACKNOWLEDGMENTS

Open access funding enabled and organized by Projekt DEAL.

Initial submission: May 10, 2021; Accepted: June 23, 2022  
 Editors: Stefan Nagel, Philip Bond, Amit Seru, and Wei Xiong

# Appendix A

Table A.1  
 Variable Definitions

Variable	Years	Definition
<i>CCE<sub>exposure</sub></i>	2002 to 2020	Relative frequency with which bigrams related to climate change occur in the transcripts of earnings conference calls. We count the number of such bigrams and divide by the total number of bigrams in the transcripts. Source: Self-constructed.
<i>CCE<sub>exposure</sub><sup>Opp</sup></i>	2002 to 2020	Relative frequency with which bigrams that capture opportunities related to climate change occur in the transcripts of earnings conference calls. We count the number of such bigrams and divide by the total number of bigrams in the transcripts. Source: Self-constructed.

(Continued)

Table A.1—Continued

Variable	Years	Definition
$CCExposure^{Reg}$	2002 to 2020	Relative frequency with which bigrams that capture regulatory shocks related to climate change occur in the transcripts of earnings conference calls. We count the number of such bigrams and divide by the total number of bigrams in the transcripts. Source: Self-constructed.
$CCExposure^{Phy}$	2002 to 2020	Relative frequency with which bigrams that capture physical shocks related to climate change occur in the transcripts of earnings conference calls. We count the number of such bigrams and divide by the total number of bigrams in the transcripts. Source: Self-constructed.
$CCExposure^{Q\&A}$	2002 to 2020	Relative frequency with which bigrams related to climate change occur in the Q&A session part of transcripts of earnings conference calls. We count the number of such bigrams and divide by the total number of bigrams in the Q&A session. Source: Self-constructed.
$CCSentiment^{Pos}$	2002 to 2020	Relative frequency with which bigrams related to climate change are mentioned together with positive tone words that are summarized by Loughran and McDonald (2011) in one sentence in the transcripts of earnings conference calls. We count the number of such bigrams and divide by the total number of bigrams in the transcripts. Source: Self-constructed.
$CCSentiment^{Neg}$	2002 to 2020	Relative frequency with which bigrams related to climate change are mentioned together with the negative tone words that are summarized by Loughran and McDonald (2011) in one sentence in the transcripts of earnings conference calls. Source: Self-constructed.
$CCRisk$	2002 to 2020	Relative frequency with which bigrams related to climate change are mentioned together with the words “risk” or “uncertainty” (or synonyms thereof) in one sentence in the transcripts of earnings conference calls. We count the number of such bigrams and divide by the total number of bigrams in the transcripts. Source: Self-constructed.
$CCExposure^{10K}$	2002 to 2020	Climate change exposure constructed by applying our algorithm to the “Management Discussion and Analysis” (MD&A) section in firms’ annual 10K filings. Source: Self-constructed.
<i>Total Emissions</i>	2004 to 2020	Sum of annual Scope 1 and Scope 2 carbon emissions (metric tons of CO <sub>2</sub> ) at the end of the year. Scope 1 emissions are caused by the combustion of fossil fuels or releases during manufacturing. Scope 2 emissions originate from the purchase of electricity, heating, or cooling. Source: Trucost.

(Continued)



Table A.1—Continued

Variable	Years	Definition
<i>WSJ CC News Index</i>	2002 to 2017	Time-series index of the fraction of the <i>Wall Street Journal</i> dedicated to the topic of climate change. Source: Engle et al. (2020).
<i>#Green-Tech Jobs</i>	2007, 2010 to 2020	Number of job postings for disruptive green-tech jobs in a year according to the Burning Glass (BG) database. Disruptive green-tech job postings relate to jobs in one of four climate-related technology areas identified by Bloom et al. (2021) as having been disruptive (“hybrid vehicle electric car,” “lithium battery,” “solar power,” and “fracking”). We assume that no disruptive green-tech job has been posted if a firm-year is not included in the BG database. Source: Bloom et al. (2021) and BG.
<i>I(Green-Tech Jobs)</i>	2007, 2010 to 2020	Indicator equal to one if <i>#Green – Tech Jobs</i> is positive, and zero otherwise. Source: Bloom et al. (2021) and BG.
<i>Green-Tech Ratio</i>	2007, 2010 to 2020	Number of job postings for disruptive green-tech jobs relative to the total number of all disruptive job postings. Set to zero if the number of disruptive job postings is zero. Source: Bloom et al. (2021) and BG.
<i>#Nongreen-Tech Jobs</i>	2007, 2010 to 2020	Number of job postings for nongreen disruptive tech jobs in a year according to the BG database. Nongreen disruptive tech job postings relate to jobs in one of 25 climate-related technology areas identified by Bloom et al. (2021) as having been disruptive and are unrelated to climate change. We assume that no nongreen disruptive tech job has been posted if a firm-year is not included in the BG database. Source: Bloom et al. (2021) and BG.
<i>#Green Patents</i>	2002 to 2019	Number of green patents obtained in a year according to the Google Patents (GP) database. To identify “green” patents, we follow Cohen, Gurun, and Nguyen (2021) and apply the OECD classification to identify what constitutes a patent with the potential to address environmental problems. We assume that no green patenting has occurred if we are unable to identify a green patent in the GP database for a firm-year. Source: GP.
<i>I(Green Patents)</i>	2002 to 2019	Indicator equal to one if <i>#Green Patents</i> is positive, and zero otherwise. Source: GP.
<i>Green Patents Ratio</i>	2002 to 2019	Number of green patents ( <i>#Green Patents</i> ) relative to the total number of patents. Set to zero if the number of total patents is zero. Source: GP.
<i>#Nongreen Patents</i>	2002 to 2019	Number of nongreen patents obtained in a year according to the GP database. We assume that no patenting has occurred if we are unable to identify a nongreen patent in the GP database for a firm-year. Source: GP.
<i>Assets</i>	2002 to 2020	Total assets (in \$ millions) at the end of the year (Compustat item AT). Winsorized at the 1% level. Source: Compustat NA/Global

(Continued)

Table A.1—Continued

Variable	Years	Definition
<i>Debt/Assets</i>	2002 to 2020	Sum of the book value of long-term debt (Compustat data item DLTT) and the book value of current liabilities (DLC) divided by total assets (Compustat data item AT). Winsorized at the 1% level. Source: Compustat NA/Global.
<i>Cash/Assets</i>	2002 to 2020	Cash and short-term investments (Compustat data item CHE) divided by total assets (Compustat data item AT). Winsorized at the 1% level. Source: Compustat NA/Global.
<i>PPE/Assets</i>	2002 to 2020	Property, plant, and equipment (Compustat data item PPENT) divided by total assets (Compustat data item AT). Winsorized at the 1% level. Source: Compustat NA/Global.
<i>EBIT/Assets</i>	2002 to 2020	Earnings before interest and taxes (Compustat data item EBIT) divided by total assets (Compustat data item AT). Winsorized at the 1% level. Source: Compustat NA/Global.
<i>R&amp;D/Assets</i>	2002 to 2019	R&D expenditures (Compustat data item XRD) divided by total assets (Compustat data item AT). Missing values set to zero. Winsorized at the 1% level. Source: Compustat NA/Global.
<i>CAPEX/Assets</i>	2002 to 2020	Capital expenditures (Compustat data item CAPX) divided by total assets (Compustat data item AT). Winsorized at the 1% level. Source: Compustat NA/Global.
<i>IVar</i>	2002 to 2020	Implied variance of log returns computed from 30-day out-of-the-money options following Bakshi, Kapadia, and Madan (2003). Winsorized at the 1% level. Source: Ivy DB OptionMetrics Volatility Surface File.
<i>ISkew</i>	2002 to 2020	Implied skewness of log returns computed from 30-day out-of-the-money options following Bakshi, Kapadia, and Madan (2003). Winsorized at the 1% level. Source: Ivy DB OptionMetrics Volatility Surface File.
<i>IKurt</i>	2002 to 2020	Implied kurtosis of log returns computed from 30-day out-of-the-money options following Bakshi, Kapadia, and Madan (2003). Winsorized at the 1% level. Source: Ivy DB OptionMetrics Volatility Surface File.
<i>SlopeD</i>	2002 to 2020	Slope of the implied volatility smile on the left side from the at-the-money level (i.e., for negative returns relative to ATM), computed as the slope coefficient from regressing implied volatilities of out-of-the-money puts on the respective option deltas (and a constant). The variable is computed from 30-day options. Winsorized at the 1% level. Source: Ivy DB OptionMetrics Volatility Surface File.

(Continued)

Table A.1—Continued

Variable	Years	Definition
<i>SlopeU</i>	2002 to 2020	Slope of the implied volatility smile on the right side from the at-the-money level (i.e., for positive returns relative to ATM), computed as the slope coefficient from regressing implied volatilities of out-of-the-money calls on the respective option deltas (and a constant). The variable is computed from 30-day options. Winsorized at the 1% level. Source: Ivy DB OptionMetrics Volatility Surface File.
<i>VRP</i>	2002 to 2020	Variance risk premium computed as the difference between the implied variance of log returns ( <i>IVar</i> ) and the realized variance of daily log returns over a historical monthly window. Winsorized at the 1% level. Source: Ivy DB OptionMetrics Volatility Surface File for options data and CRSP for daily returns.

# REFERENCES

Acemoglu, Daron, Philippe Aghion, Lint Barrage, and David Hemous, 2019, Climate change, directed innovation, and energy transition: The long-run consequences of the shale gas revolution, Working paper, MIT.

Addoum, Jawad, David Ng, and Ariel Ortiz-Bobea, 2020, Temperature shocks and establishment sales, *Review of Financial Studies* 33, 1331–1366.

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt, 2005, Competition and innovation: An inverted-U relationship, *Quarterly Journal of Economics* 120, 701–728.

Andersson, Mats, Patric Bolton, and Frédéric Samama, 2016, Hedging climate risk, *Financial Analysts Journal* 72, 13–32.

Baker, Scott R., Nicholas Bloom, and Steven J. Davis, 2016, Measuring economic policy uncertainty, *Quarterly Journal of Economics* 131, 1593–1636.

Bakshi, Gurdip S., Nikunj Kapadia, and Dilip B. Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 101–143.

Bali, Turan G., and Hao Zhou, 2016, Risk, uncertainty, and expected returns, *Journal of Financial and Quantitative Analysis* 51, 707–735.

Ball, Ray, and Eli Bartov, 1996, How naive is the stock market's use of earnings information?, *Journal of Accounting and Economics* 21, 319–337.

Barras, Laurent, and Aytex Malkhozov, 2016, Does variance risk have two prices? Evidence from the equity and option markets, *Journal of Financial Economics* 121, 79–92.

Baz, Salim, Lara Cathcart, Alexander Michaelides, and Yi Zhang, 2022, Firm-level climate regulatory exposure, Working paper, Imperial College London.

Berkman, Henk, Jonathan Jona, and Naomi S. Soderstrom, 2019, Firm-specific climate risk and market valuation, Working paper, University of Auckland Business School.

Bingler, Julia Anna, Mathias Kraus, and Markus Leippold, 2022, Cheap talk and cherry-picking: What ClimateBert has to say on corporate climate risk disclosures, *Finance Research Letters* 47, 102776.

Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen Terry, 2018, Really uncertain business cycles, *Econometrica* 86, 1031–1065.

- Bloom, Nicholas, Tarek A. Hassan, Aakash Kalyani, Josh Lerner, and Ahmed Tahoun, 2021, The diffusion of disruptive technologies, Working paper, Stanford University.
- Bolton, Patrick, and Marcin Kacperczyk, 2021, Do investors care about carbon risk?, *Journal of Financial Economics* 142, 517–549.
- Bolton, Patrick, and Marcin Kacperczyk, 2022, Carbon disclosure and the cost of capital, Working paper, Imperial College London.
- Bolton, Patrick, and Marcin Kacperczyk, 2023, Global pricing of carbon-transition risk, *Journal of Finance* (forthcoming).
- Bolton, Patrick, Marcin Kacperczyk, and Frédéric Samama, 2021, Carbon disclosure and the cost of capital, Working paper, Imperial College London.
- Campello, Murillo, Janet Gao, and Qiping Xu, 2021, Personal income taxes and labor downskilling: Evidence from 27 million job postings, Working paper, Kelley School of Business.
- Carhart, Mark M., 1997, On persistence in mutual fund performance, *Journal of Finance* 52, 57–82.
- Chen, Shuping, Stephan Hollander, and Kelvin Law, 2014, The price of silence: When no one asks questions during conference calls, Working paper, University of Texas at Austin.
- Choi, Garwin, Zhenyu Gao, and Wenxi Jiang, 2020, Attention to global warming, *Review of Financial Studies* 33, 1112–1145.
- Cohen, Lauren, Umit Gurun, and Quoc Nguyen, 2021, The ESG-innovation disconnect: Evidence from green patenting, Working paper, Harvard Business School.
- Cohn, Jonathan B, Zack Liu, and Malcolm Wardlaw, 2022, Count (and count-like) data in finance, Working paper, University of Texas at Austin.
- Collard-Wexler, Allan, 2011, Productivity dispersion and plant selection in the ready-mix concrete industry, Working paper, NYU Stern.
- Cremers, Martijn, Antti Petajisto, and Eric Zitzewitz, 2013, Should benchmark indices have alpha? Revisiting performance evaluation, *Critical Finance Review* 2, 1–48.
- Darendeli, Alper, Kelvin Law, and Michael Shen, 2022, Green new hiring, *Review of Accounting Studies* 27, 986–1037.
- De Haas, Ralph, and Alexander Popov, 2023, Finance and green growth, *Economic Journal* 133, 637–668.
- Di Giuli, Alberta, Alexandre Garel, Roni Michaely, and Arthur Petit-Romec, 2022, Climate change and mutual fund voting on environmental proposals, Working paper, European Corporate Governance Institute.
- Duong, Huu Nhan, Petko S. Kalev, Madhu Kalimipalli, and Saurabh Trivedi, 2021, Do firms benefit from carbon risk management? Evidence from the credit default swaps market, Working paper, Monash University.
- Engle, Robert F., Stefano Giglio, Bryan Kelly, Heebum Lee, and Johannes Stroebel, 2020, Hedging climate change news, *Review of Financial Studies* 33, 1184–1216.
- Fama, Eugene F., and Kenneth R. French, 2008, Dissecting anomalies, *Journal of Finance* 63, 1653–1678.
- Fama, Eugene F., and James D. MacBeth, 1973, Risk return and equilibrium: Empirical tests, *Journal of Political Economy* 81, 607–636.
- Feng, Guanhao, Stefano Giglio, and Dacheng Xiu, 2020, Taming the factor zoo: A test of new factors, *Journal of Finance* 75, 1327–1370.
- Gagliardini, Patrick, Elisa Ossola, and Olivier Scaillet, 2016, Time-varying risk premium in large cross-sectional equity data sets, *Econometrica* 84, 985–1046.
- Gentzkow, Matthew, Bryan Kelly, and Matt Taddy, 2019, Text as data, *Journal of Economic Literature* 57, 535–574.
- Giglio, Stefano, Bryan Kelly, and Johannes Stroebel, 2021, Climate finance, *Annual Review of Financial Economics* 13, 15–36.
- Ginglinger, Edith, and Quentin Moreau, 2022, Climate risk and capital structure, Working paper, Université Paris–Dauphine.
- Gostlow, Glen, 2021, The materiality and measurement of physical climate risk: Evidence from Form 8-K, Working paper, London School of Economics.

- Hail, Luzi, Shawn Kim, and Rachel Xi Zhang, 2021, How do managers greenwash? Evidence from earnings conference calls, Working paper, University of Pennsylvania Wharton School.
- Haščić, Ivan, and Mauro Migotto, 2015, Measuring environmental innovation using patent data, Working paper, OECD.
- Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun, 2019, Firm-level political risk: Measurement and effects, *Quarterly Journal of Economics* 134, 2135–2202.
- Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun, 2021, Firm-level exposure to epidemic diseases: COVID-19, SARS, and H1N1, NBER Working paper 26971.
- Hassan, Tarek A., Stephan Hollander, Laurence van Lent, and Ahmed Tahoun, 2023a, The global impact of Brexit uncertainty, *Journal of Finance* (forthcoming).
- Hassan, Tarek A., Jesse Schreger, Markus Schwedeler, and Ahmed Tahoun, 2023b, Sources and transmission of country risk, *Review of Economic Studies* (forthcoming).
- Heath, Davidson, Daniele Macciocchi, Roni Michaely, and Matthew C. Ringgenberg, 2022, Does socially responsible investing change firm behavior? Working paper, University of Utah.
- Heo, Yuna, 2021, Climate change exposure and firm cash holdings, Working paper, Rutgers Business School.
- Hollander, Stephan, Maarten Pronk, and Erik Roelofsen, 2010, Does silence speak? An empirical analysis of disclosure choices during conference calls, *Journal of Accounting Research* 48, 531–563.
- Hong, Harrison, Frank Weikai Li, and Jiangmin Xu, 2019, Climate risks and market efficiency, *Journal of Econometrics* 208, 265–281.
- Huang, Allen H., Reuven Lehavy, Amy Y. Zang, and Rong Zheng, 2018, Analyst information discovery and interpretation roles: A topic modeling approach, *Management Science* 64, 2833–2855.
- Ilhan, Emirhan, Zacharias Sautner, and Grigory Vilkov, 2021, Carbon tail risk, *Review of Financial Studies* 34, 1540–1571.
- Jamilov, Rustam, Hélène Rey, and Ahmed Tahoun, 2021, The anatomy of cyber risk, NBER Working paper 28906.
- Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, 2021, Measuring technological innovation over the long run, *American Economic Review: Insights* 3, 303–320.
- Kelly, Bryan, Lubos Pastor, and Pietro Veronesi, 2016, The price of political uncertainty: Theory and evidence from the option market, *Journal of Finance* 71, 2417–2480.
- King, Gary, Patrick Lam, and Margaret E. Roberts, 2017, Computer-assisted keyword and document set discovery from unstructured text, *American Journal of Political Science* 61, 971–988.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, 2017, Technological innovation, resource allocation, and growth, *Quarterly Journal of Economics* 132, 665–712.
- Kölbel, Julian F., Markus Leippold, Jordy Rillaerts, and Qian Wang, 2022, Ask BERT: How regulatory disclosure of transition and physical climate risks affects the CDS term structure, *Journal of Financial Econometrics* (forthcoming).
- Krueger, Philipp, Zacharias Sautner, and Laura T. Starks, 2020, The importance of climate risks for institutional investors, *Review of Financial Studies* 33, 1067–1111.
- Kruttili, Mathias S., Brigitte Roth Tran, and Sumudu W. Watugala, 2021, Pricing Poseidon: Extreme weather uncertainty and firm return dynamics, Working paper, Federal Reserve Bank of San Francisco.
- Lancaster, Tony, 2000, The incidental parameter problem since 1948, *Journal of Econometrics* 95, 391–413.
- Lewellen, Jonathan, and Stefan Nagel, 2006, The conditional CAPM does not explain asset pricing anomalies, *Journal of Financial Economics* 82, 289–314.
- Li, Fengfei, Chen Lin, and Tse-Chun Lin, 2022, A one-two punch to the economy: Climate vulnerability and corporate innovation strategies, Working paper, The University of Hong Kong.
- Li, Qing, Hongyu Shan, Yuehua Tang, and Vincent Yao, 2021, Corporate climate risk: Measurement and response, Working paper, University of Florida.
- Loughran, Tim, and Bill McDonald, 2011, When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks, *Journal of Finance* 66, 35–65.
- Mayew, William J., 2008, Evidence of management discrimination among analysts during earnings conference calls, *Journal of Accounting Research* 46, 627–659.

- Mueller, Isabella, and Eleonora Sfrappini, 2022, Climate change-related regulatory risks and bank lending, Working paper, IWH-Halle.
- Pankratz, Nora, and Christoph Schiller, 2021, Climate change and adaptation in global supply-chain networks, Working paper, Board of Governors of the Federal Reserve System.
- Pastor, Lubos, Robert Stambaugh, and Lucian Taylor, 2021, Sustainable investing in equilibrium, *Journal of Financial Economics* 142, 550–571.
- Pickl, Matthias, 2019, The renewable energy strategies of oil majors—From oil to energy?, *Energy Strategy Reviews* 29, 1–7.
- Sautner, Zacharias, Laurence van Lent, Grigory Vilkov, and Ruishen Zhang, 2022, Pricing climate change exposure, *Management Science* (forthcoming).
- Shiller, Robert J., 2017, Narrative economics, *American Economic Review* 107, 967–1004.
- Stern, Nicholas, and Anna Valero, 2021, Innovation, growth and the transition to net-zero emissions, *Research Policy* 50, 104293.
- Tomar, Sorabh, 2023, Greenhouse gas disclosure and emissions benchmarking, *Journal of Accounting Research* (forthcoming).
- van Zaanen, Menno, and Pieter Kanter, 2010, Automatic mood classification using tf\*idf based on lyrics, in *Proceedings of the 11th International Society for Music Information Retrieval Conference*, pp. 75–80.
- Varini, Francesco S., Jordan Boyd-Graber, Massimiliano Ciaramita, and Markus Leippold, 2020, ClimaText: A dataset for climate change topic detection, Working paper, ETH Zurich.
- von Schickfus, Marie-Theres, 2021, Institutional investors, climate policy risk, and directed innovation, Working paper, ifo Institute Munich.
- Webersinke, Nicolas, Mathias Kraus, Julia Anna Bingler, and Markus Leippold, 2021, Climate-Bert: A pretrained language model for climate-related text, Working paper, University of Zurich.
- World Economic Forum, 2021, Financing the transition to a net-zero future, Report.

### Supporting Information

Additional Supporting Information may be found in the online version of this article at the publisher's website:

**Appendix S1: Internet Appendix.  
Replication Code.**