An EDHEC-Risk Climate Impact Institute Publication

The Impact of Climate Change News on Low-minus-High Carbon Intensity Portfolios

June 2023

with the support of

Amundi
ASSET MANAGEMENT

EDHEC-RISK
Climate Impact
INSTITUTE

Table of Contents

Authors and acknowledgments
This publication was prepared by Jean-Michel Maeso and Dominic O'Kane. The authors thank Frédéric Ducoulombier, Teodor Dyakov, Lionel Martellini, Irene Monasterolo, Riccardo Rebonato and Olivier David Zerbib for their very helpful comments.

Funding/support
The present publication was prepared as part of the «Measuring and Managing Climate Risks in Investment Portfolios» Research Chair at EDHEC-Risk Climate Impact Institute, established with the support of Amundi.

Scientific integrity and competing interests
The authors confirm that they abide by the general principles and requirements of the European Charter for Researchers (2005/251/EC) and declare no competing interests.

Citing this report
Maeso, J. M. and D. O'Kane (2023) The Impact of Climate Change News on Low-minus-High Carbon Intensity Portfolios. EDHEC-Risk Climate Impact Institute, EDHEC Business School (June).

Foreword

This study is the first output of the Research Chair «Measuring and Managing Climate Risks in Investment Portfolios» established at EDHEC-Risk Climate Impact Institute with the support of Amundi.

Several pioneering papers have already examined the link between climate news and equity market returns with a view to isolating "climate beta" that could be used to construct climate-risk hedging portfolios with easy-to-trade assets. However, this study applies the latest natural language processing methods to construct climate news indices from newspaper articles. Linguistic dictionary, lexical sentiment-based techniques, and state-of-the-art transformer-based models are used to capture climate change concerns in leading newspapers over the 2005-2021 period.

Daily indices are built for each source and for the aggregation of all five major news sources. The authors isolate the unexpected changes in these indices and study the relationship between these innovations and the daily changes in value of portfolios of large capitalisation companies listed in the United States.

Sorting by the carbon intensity of issuers, the authors then create a high carbon intensity portfolio from the top tercile, a low carbon intensity portfolio from the bottom tercile, and a long-short portfolio of low-minus-high carbon intensity stocks. Portfolio returns are regressed against each climate news innovation index (while controlling for exposure to standard return factors).

When single news sources are used to construct each index, the authors observe statistical significance (at the commonly used 5% level) for only a quarter of the indices for high carbon intensity portfolios and a sliver of the indices for long-short portfolios. However, when the aggregate index is used, statistical significance is observed for almost all the language models for the high carbon intensity portfolios and two thirds of these models for the long-short portfolios – aggregation reduces the weight of idiosyncratic articles.

These results are consistent with previous studies documenting overperformance of low over high carbon intensity portfolios in reaction to innovations in climate change concerns, but the authors find that this outperformance arises from the fall in the value of the high carbon intensity portfolios. Interestingly, the value added by the most advanced approaches relative to the simple attention-based model is modest.

The authors' documentation of a negative reaction of high carbon intensity equity prices to climate change concerns (as reflected in the news) is consistent with the intuition that higher carbon intensity activities have higher exposure to the risks arising from climate change and climate change action. The finding that this relationship is statistically significant and of the right sign is good news; however, as in previous studies, the economic impact does not appear significant enough for comfort – one disquieting explanation for investors is that equity prices have yet to adjust to the climate emergency.

I wish to commend Doctor Jean-Michel Maeso and Professor Dominic O'Kane for the considerable data and modeling work underpinning their interesting study and to thank Alice James and Laurent Ringelstein for their editing and publishing work.

We wish you a useful and informative read.

Frédéric Ducoulombier
Director of EDHEC-Risk Climate Impact Institute

About the Authors

Jean-Michel Maeso is a Senior Quantitative Researcher at EDHEC-Risk Climate Impact Institute, where he is working on the application of natural-language processing to the production of climate metrics and the study of their relationships with asset prices. After joining EDHEC-Risk Institute in 2015, he conducted research along three main avenues: (1) machine learning techniques applied to quantitative finance problems; (2) portfolio construction for retirement investing; and (3) factor investing in equity and fixed income universes. He holds a PhD in Applied Mathematics from Université Nice Côte d'Azur and an engineering degree from the Ecole Centrale de Lyon with a specialisation in applied mathematics. Previously, he spent five years in the financial industry specialising in the research, development and implementation of investment solutions (structured products and systematic strategies) for institutional investors. His research has been published in such journals as *Quantitative Finance*, *The Journal of Portfolio Management*, and *The Journal of Fixed Income*.

Dominic O'Kane is a Research Director at EDHEC-Risk Climate Impact Institute and a Professor of Finance at EDHEC Business School. His current focus is on climate metrics and the impact of climate change on investment. His previous focus was on asset pricing and financial stability in the derivatives market. He holds a PhD in Theoretical Physics from Oxford University which he followed with a post-doctoral position at Imperial College London. Prior to joining EDHEC Business School, he spent three years at Salomon Brothers and nine years at Lehman Brothers, where he became a managing director and head of quantitative research in Europe. Dominic has taught at Oxford University, Yale University and the London Business School. His research has been published in such journals as *Quantitative Finance*, *Finance and Stochastics*, and *the Journal of Derivatives*. He is the author of a book on the valuation and risk management of credit derivatives.

Executive Summary

In recent years, several papers, beginning with Engle et al. (2020), have examined the link between climate news and equity market returns, hoping to identify a measure of a so-called "climate beta". Using a variety of language models and high-quality English language newspaper sources, including the Financial Times and New York Times, we construct a climate news index (CNI) for each model and source and an aggregate index from all the news sources. We measure the impact of each of these indices on a range of Low Carbon Intensity (LCI), High Carbon Intensity (HCI), and Low-minus-High Carbon Intensity (LmHCI) equity portfolios, constructed by sorting the top 500 US-listed firms by capitalisation (US500) based on their carbon intensity. We find that the relationship between the news indices and the LCI, HCI, and LmHCI portfolio returns is overall not statistically significant for individual news sources across all the language models. However, it does become significant for all the aggregate news indices, suggesting that combining different news sources enables the detection of news events that are covered by multiple news sources, and which are more likely to impact the market and the climate beta. For most of the language models considered, the sensitivity of returns to an increase in the corresponding aggregate news index is negative and statistically significant at 1% for HCI portfolio returns, but it is not significant for LCI portfolio returns. This is for the period from July 2012 to November 2021. These results suggest that the factor extracted from an aggregate news index is a climate risk proxy whose beta coefficient can explain the returns of HCI and hence LmHCI portfolios. We also generate climate news indices for physical risk and transition risk, and find that the statistical significance of these indices is close to that of the other climate change indices for the portfolios analysed.

News Data

To perform this analysis, we require digitised, daily, English language, highcirculation, and high-quality news sources with a European and US perspective, including at least one with a dedicated financial market focus. For these reasons we have chosen the following news data sources:

- The Financial Times (FT) digital archive. The Financial Times is widely recognised as the leading European English language financial newspaper.
- The Lexis Nexis (LN) database of newspapers. This provides access to many thousands of newspapers internationally. From these we selected the New York Times (NYT), the Los Angeles Times (LAT), The Guardian (UKG) and The Daily Telegraph (DT).

We use daily articles from these news sources over the period from 3 January 2005 to 3 November 2021. We assume that an article appears in the morning of the first publication date. This can be a different time depending on whether the news source is in the US or Europe. As Europe is several hours ahead of the US, the arrival of news from US and European news sources will impact the US equity market on that same day. To align the arrival of weekend news with the financial markets, we adjust the publication date of news stories that appear on a Saturday or Sunday to the following Monday, the earliest date on which this news can impact the US equity market.

We consider news in the form of newspaper articles where each article has both a headline and content. The headline is typically added by a sub-editor who has read the article and wishes to summarise the key message of the article to attract the reader. Consequently, the headline will usually reflect the most important part of the article and any associated positive or negative sentiment. For this reason, we will examine both the article headline and the article content to see if the headline can provide a clearer measure of article focus and sentiment than an analysis of its longer and more complex content.

Figure 1: We measure newspaper media attention to the subject of climate change from 2005-2021 by calculating the percentage of daily FT articles that contain the term 'climate change', the percentage that contain the term 'global warming' and the percentage containing either term. We apply a 30-day moving average. The vertical lines are climate change-related events - see Table 1 for the corresponding numbered list of events.

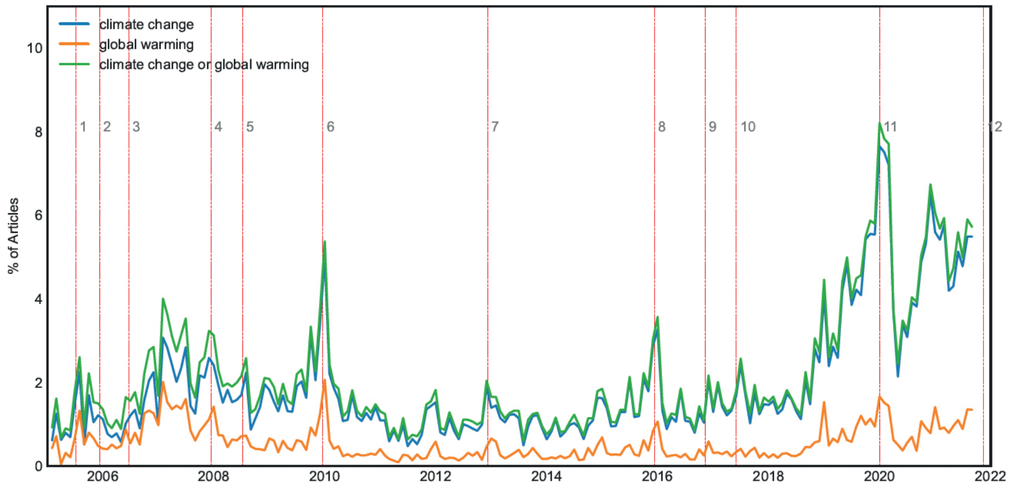

Table 1: Identification of the most active climate news events seen in Figure 1. For conferences we have used the conference end date when the final agreement is usually announced.

Event #	Date	Description
1	8 July 2005	G8 Summit Scotland
2	9 Dec 2005	Montreal CC conference
3	21 Jun 2006	UK CC and Sustainable Energy Act 2006
4	14 Dec 2007	Bali UN CC Conference
5	9 Jul 2008	G8 Summit
6	19 Dec 2009	Copenhagen UN CC Conference
7	7 Dec 2012	Doha UN CC Conference
8	12 Dec 2015	Paris Agreement Signed
9	7 Nov 2016	Marrakech UN CC Conference
10	1 Jun 2017	US President Trump Withdraws from Paris Agreement
11	Dec 2019 - Jan 2020	Australian wildfires, High temperatures
12	13 Nov 2021	Build up to COP 26 CC Glasgow

To quantify the newspaper media's attention to climate change, we first need to identify a 'climate change article'. Using single-word search terms can lead to false positives. Work by Engle et al. (2020) and others has thus tried to identify the relevant news by searching for 'climate change'. However, 'global warming' is also used to describe the same phenomenon. To determine whether we should include it, we perform searches for articles that contain (i) 'climate change', (ii) 'global warming', and (iii) one or both of 'climate change' or 'global warming'. Using the Financial Times as news source, we calculate the fraction of daily climate change articles using each of these 'bigrams'.

Using the fraction of articles rather than the number corrects for the fact that the total number of daily Financial Times articles has varied significantly over this period. All three time-series are shown in Figure 1. We observe that use of the bigram 'global warming' has declined in relative terms over time but remains significant. The bigram 'climate change' has clearly become the dominant bigram. To ensure that we capture as many climate change articles as possible, and especially those in the earlier period of analysis, we include both in our definition of a 'climate change article'. It is reassuring to note that the peaks in Figure 1 align closely with the major climate change news events listed in Table 1.

Table 2 shows the number of articles per newspaper, per year. We see that the UK Guardian is the leading publisher of climate change-related articles over time among our corpus of news sources, followed by the Financial Times and Daily Telegraph.

Climate Change News Indices Overview

We explore several approaches for constructing a CNI from newspaper articles. If there is a link between climate change news and market price movements then we would expect that the link will be strongest for the index that best captures the quantity, content, and sentiment of the climate change news. The six index construction approaches we will use, in order of increasing level of sophistication, are as follows:

Table 2: Number of climate change articles by year and news source for the Financial Times, The Daily Telegraph, The Guardian, Los Angeles Times and The New York Times news sources. Note that 2021 only includes articles up to 3 November 2021.

Year	FT	DT	UKG	LAT	NYT
2005	339	413	1,127	313	294
2006	523	645	1,673	611	385
2007	1,304	1,241	2,676	1,053	1,198
2008	1,113	723	2,499	762	997
2009	1,495	885	3,914	673	1,169
2010	1,093	1,023	2,480	548	666
2011	706	783	1,932	329	334
2012	739	899	1,853	306	287
2013	800	822	1,914	339	332
2014	800	639	2,284	468	399
2015	1,212	715	5,278	771	600
2016	892	429	4,881	660	532
2017	1,003	465	2,002	776	491
2018	1,062	566	2,570	697	389
2019	2,040	1,393	4,051	984	618
2020	1,909	1,180	3,091	697	485
2021	2,218	1,866	3,718	1,076	701

1. Attention - The number of climate change articles published each day.
2. Similarity - The similarity between each day's climate change articles and a representative climate change document.
3. Concern - Climate change concern using word-frequencies based on text analysis

lexicons designed to measure psychological and social processes. Concern is high if the number of 'negative words' in a climate change article is higher than the number of 'positive words' and the fraction of 'risk words' is high.

4. VADER - We use a rules-based lexical approach called VADER that assigns a sentiment polarity score to specific words to determine if the climate change article sentiment is positive or negative.
5. BERT with Fine-Tuned Sentiment - We use a 'language model', i.e., an artificial intelligence system that can understand and generate human language, and fine-tune it to identify sentiment using human-labelled, finance-related training examples (we rely on the BERT language model of Devlin et al. (2018)).
6. ClimateBERT with Fine-Tuned Sentiment - We take the domain-specific Climate-BERT model by Webersinke et al. (2021) and fine-tune it to identify sentiment using human-labelled, finance-related training examples.

The first Attention-based measure simply counts the daily number of climate change articles published. Such an approach was analysed in the media bias model of Gentzkow and Shapiro (2010), who noted that the number and length of the articles reflect reader interest. Our second approach is based on the Natural Language Processing (NLP) metric called Term Frequency - Inverse Document Frequency (TF-IDF) which measures how important a word is to a document located in a collection of documents; here we follow the approach of Engle et al. (2020). For the next level of sophistication, we use an approach that detects actual semantic meaning within the news articles by quantifying the degree of emotional concern. This is a challenging task as the emotional sentiment of long and highly articulate articles is not always simple to extract. For this reason, we first use a sentiment-analysis tool named VADER and then turn to state-of-the-art language models such as the BERT Transformer-based model from Devlin et al. (2018), and the CBERT model by Webersinke et al. (2021) which has been specifically designed to better understand climate-related texts. We develop two approaches for each of these sophisticated tools: one using headlines (VAD-H, BERT-H and CBERT-H) and another using article content (VAD-S, BERT-S and CBERT-S) for sentiment analysis.

We use these different approaches to construct a family of CNIs, each a daily time series from 2005 to 2021. To perform a market analysis using our climate change news indices we must isolate the unexpected component of the daily climate change news index changes. To do this, we assume that the CNI obeys a simple process (a first-order autoregressive 'AR(1)' model), whereby the current value of the index is predicted by its previous value and the error in prediction is the innovation. Calibrating the CNI to this process allows us to extract a family of unexpected climate change news (UCNI) indices.

In addition to the set of CNIs for each news source, we also wish to construct a single aggregate index by averaging across all five news sources. Doing this increases the total number of articles being used. It might also be expected to increase the significance of climate news-related effects in the article counts. For example, a newspaper may decide on a specific date to publish a feature article on climate change that is not linked to a specific news event. It would almost certainly not coincide with a climate change article in any other news source. However a news story in all five news sources on a specific date is almost certainly driven by a common climate news event. By averaging over the news indices on each date, the aggregate index is more able to identify actual climate news events.

Rather than simply average the individual newspaper indices, we first standardise them so that each index has a unit standard deviation over a three-year period of T dates. This ensures that a newspaper index that experiences a high level of variability in both article number and sentiment score is adjusted to be more comparable to a newspaper index that has a lower variability.

Climate Change News and Equity Portfolios

To study the impact of news on portfolios, we first extract the top 500 US stocks by capitalisation from the CRSP database and allocate each stock to a carbon intensity group. As per the definition given by the Task Force on Climate-related Financial Disclosures, carbon intensity is defined as the ratio of direct ('Scope 1') and energy-related ('Scope 2') emissions of the issuing company to its revenues. We rely on ISS ESG data sourced via FactSet. A stock is labelled 'LCI'/'HCI' if its issuer is in the lower/upper three deciles by carbon intensity and 'neutral' otherwise. As noted by Vaucher et al. (2023)[1], the use of this metric as single sorting criterion is commonplace in the literature and the added value of more complex sorting approaches remains moot (e.g., Roncalli et al. (2020) show that the composite indicator built from 10 metrics by Görgen et al. (2020) is well captured by a factor based on the emissions intensity only).

Given a set of climate change news indices described above, the next step is to determine whether these indices have an impact on equity returns. We examined three liquid US stock portfolios, engaged respectively in a LCI strategy, a HCI strategy, and a Low-minus-High Carbon Intensity (LmHCI) strategy, that we expect to be sensitive to climate change risk. The LCI and HCI portfolios equally weight the LCI and HCI stocks, respectively - as such they consist of the 30% of stocks with respectively the lowest and highest carbon intensities[2].

The CNI indices that we have calculated may embed some auto-correlation effects and these must be removed if we are to correctly capture the unexpected changes in the climate news index. The unexpected climate news index at time t, $UCNI_t$, are defined by

$$UCNI_t = CNI_t - \mathbb{E}\left[CNI_t | I_{t-1}\right]$$

Where I_{t-1} is the information to time $t-1$. Each value of the $UCNI_t$ is calculated as the residual of an AR(1) process calibrated to the CNIt over the previous three years. The Aggregate $UCNI_t$ for the different language models are shown in Figure 2.

Then we examine whether differences in exposure to the climate news index help us to explain expected returns of LCI, HCI, and LmHCI portfolios. To do so, we look at whether the return in excess of the risk-free rate can be explained as a linear function of the UCNI

1 - Vaucher et al. (2023) emphasise that sorting stocks according to carbon intensity may lead to large weights toward sectors that are not relevant from a climate policy standpoint according to the classification of economic activities developed after Battiston et al. (2017) to assess climate transition risk.

2 - We also tested an alternative approach for the LmHCI portfolio construction developed by Vaucher et al. (2023). The long ("LCI") leg of the LmHCI portfolio is built as an equally-weighted (EW) portfolio of the 50% of the stocks with the lowest carbon intensity selected within each of six sectors with highest exposure to stranding risk in the event of a disorderly low-carbon transition (according to the classification of Battiston et al. (2017)). Conversely, the short ("HCI") leg is built as an EW portfolio of the 50% of the stocks with the highest carbon intensity selected within each of these sectors. The findings generated via this alternate methodology concur with those presented here.

- we do so while controlling for the return for exposure to standard factors, in the sense of Fama and French (2015):

$$\begin{aligned} \widetilde{r_t} &= \alpha + \beta_{UCNI} \cdot UCNI_t + \beta_{MKT} \cdot MKT_t + \beta_{HML} \cdot HML_t + \beta_{SMB} \cdot SMB_t \\ &+ \beta_{CMA} \cdot CMA_t + \beta_{RMW} \cdot RMW_t + \beta_{WML} \cdot WML_t + \varepsilon_t \end{aligned} \tag{1}$$

where *MKT* is the excess market return, *SMB* the small minus big (size) factor, *HML* the high minus low (value) factor, *RMW* the robust minus weak (high profitability) factor, *CMA* the conservative minus aggressive (low investment) factor, and *WML* the winners minus losers (momentum) factor.

The left-hand side is the daily excess return of the portfolio under study. We want to determine if, for the linear regression, the factor loading on the Aggregate $UCNI_t$ is statistically significant. This will test the model of Pastor et al. (2021) which predicts a strictly positive β_{UCNI} coefficient for the LCI and LmHCI portfolios and a strictly negative β_{UCNI} coefficient for the HCI portfolio. We will also study the statistical significance of β_{UCNI} coefficients.

Figure 2: Climate change news indices (CNI) for the period July 2012 to November 2021 and Unexpected Climate News Innovations (UCNI) for the aggregate news source over the period July 2012 to November 2021.

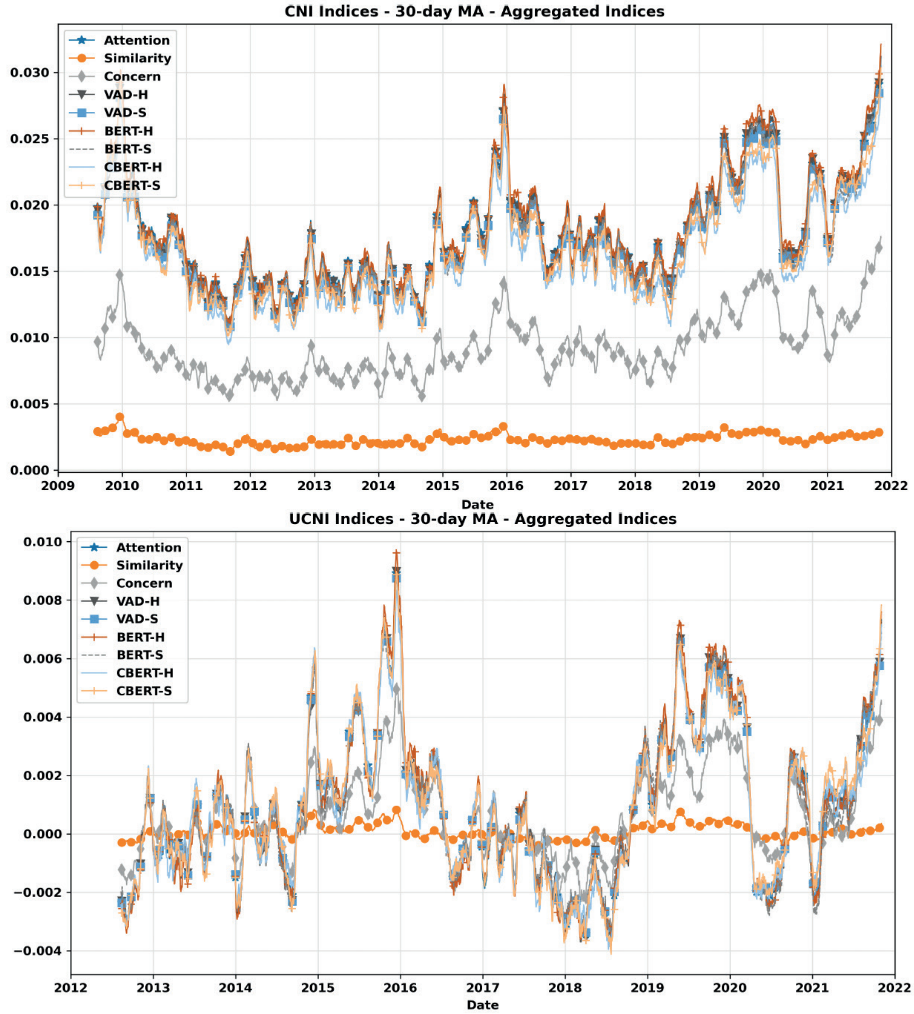

We apply the linear regression written in Equation (1) to the LCI, HCI and LmHCI portfolios. Table 3 reports the β_{UCNI} coefficient for the LCI portfolio. These results have low significance and hence the impact of climate news on LCI stock returns is low.

*Table 3: Value of the corresponding UCNI beta and significance for the LCI portfolio by news source and index construction methodology. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	FT	Daily T.	Guardian	NYT	LAT	Aggregate
Attention	0.0042	-0.001	-0.004	-0.0037	-0.0002	-0.0006
Similarity	-0.0048	0.0379	0.0096	-0.0098	0.0381	0.0514
Concern	0.0027	0.0047	-0.0021	-0.0068	0.0039	0.0058
VAD-H	0.005	-0.0012	-0.0037	-0.0034	-0.0003	-0.0003
VAD-S	0.0063	-0.0017	-0.0042	-0.0046	0.0004	-0.0005
BERT-H	0.0017	-0.0006	-0.0054	-0.0039	-0.0015	-0.0023
BERT-S	0.0054	-0.0017	-0.003	-0.0027	0.0016	0.0009
CBERT-H	0.0016	-0.0044	-0.0043	-0.0095*	0.0012	-0.0051
CBERT-S	0.0089*	-0.0012	-0.0038	-0.0036	-0.0007	0.0007

Table 4 shows the equivalent UCNI beta for the HCI portfolio for the different news sources, including the aggregate news source. We note that the betas are all negative, implying that a day with high unexpected negative concerns is, on average, always associated with a negative impact on the returns of HCI stocks. Significance for the individual news source betas is reasonably high with some at 10%, 5% and one at 1% significance (CBERT-H). However, for the aggregate index, the significance improves substantially, and five of the nine indices have significance at 1%, three having significance at 5% amongst the other four.

Table 5 shows the UCNI beta coefficient for the LmHCI portfolio for the linear regression performed over the period from July 2012 to November 2021. The first observation is that the betas are mostly positive: unexpected negative climate sentiment is associated with a positive return due to a fall in value of the short position in the HCI stocks. Once again the climate beta statistical significance is greatest for the Aggregate indices, with six out of nine climate betas showing statistical significance at 5%.

For the VADER, BERT, and Climate BERT models, the (H) approach looking at the titles of articles achieves the same level of significance for the HCI portfolios as the (S) approach considering the contents of the articles (S). Nevertheless, for the LmHCI portfolios, the (S) approach yields more significant results (1% for BERT-S and CBERT-S versus 5% for BERT-H and CBERT-H).

*Table 4: Value of the corresponding UCNI beta and significance for the HCI portfolio by news source and index construction methodology. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	FT	Daily T.	Guardian	NYT	LAT	Aggregate
Attention	-0.0135	-0.0165**	-0.0157**	-0.0102	-0.0139*	-0.0295***
Similarity	-0.0296	-0.0807	0.059	-0.0522	-0.1022** -	0.1735*
Concern	-0.0202*	-0.0087	-0.0111	-0.0231*	-0.0186	-0.0435**
VAD-H	-0.0127	-0.0142*	-0.0158**	-0.0099	-0.0144*	-0.0296***
VAD-S	-0.0142	-0.0142*	-0.0156**	-0.0084	-0.014*	-0.0296***
BERT-H	-0.0151*	-0.0112*	-0.0117*	-0.0099	-0.0097	-0.0257**
BERT-S	-0.0154*	-0.0102	-0.014**	-0.0099	-0.0087	-0.0272**
CBERT-H	-0.0116	-0.0179***	-0.014**	-0.0026	-0.016**	-0.0311***
CBERT-S	-0.011	-0.0185**	-0.016**	-0.0145	-0.0149*	-0.0334***

*Table 5: Value of the corresponding UCNI beta and significance for the LmHCI portfolio by news source and index construction methodology. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	FT	Daily T.	Guardian	NYT	LAT	Aggregate
Attention	0.0177*	0.0156*	0.0117	0.0065	0.0136	0.0289**
Similarity	0.0248	0.1187*	-0.0494	0.0424	0.1403**	0.2249*
Concern	0.0229	0.0134	0.009	0.0162	0.0225	0.0493**
VAD-H	0.0177	0.013	0.0121	0.0065	0.0141	0.0293**
VAD-S	0.0205*	0.0125	0.0114	0.0038	0.0144	0.0291**
BERT-H	0.0168	0.0106	0.0064	0.0059	0.0082	0.0235*
BERT-S	0.0208*	0.0085	0.011	0.0072	0.0102	0.0282**
CBERT-H	0.0132	0.0136	0.0098	-0.0069	0.0172*	0.026*
CBERT-S	0.0199*	0.0174**	0.0122	0.0108	0.0141	0.0341**

Portfolio Performance Conditional on the UCNI

We examine whether unexpected climate news innovations can help predict the conditional performance of HCI versus LCI stocks. To do this, we study the average daily performance of the LmHCI portfolio over the period from July 2012 to November 2021, conditional on the level of the daily UCNI index. We define three different regimes which we call low, medium, and high. The low regime corresponds to days when the UCNI time series is below the first tercile, the medium regime to days when the UCNI time series is between the first and the second terciles, and the high regime to daily periods when the UCNI time series is above the second tercile.

Table 6 reports the conditional performance of the LCI, HCI and LmHCI portfolios with respect to UCNI calculated from the Aggregate news source. The average annualised return of the LmHCI portfolio over the high regime is greater than that over the medium and low regime for all the methodologies under study. We also find that the average annualised return of the LmHCI portfolio in the medium regime is greater than that over the low regime for all the methodologies under study. For example, choosing the ClimateBERT sentences index it is 13.6% in the high regime versus 4.0% in the low regime. We see that when the UCNI is in the high tercile, the return of the HCI portfolio is negative for all the indices except the similarity index.

Conclusion and Extensions

We constructed nine climate change news indices (CNI) using news data from five different high-quality news sources, specifically the Financial Times, The Daily Telegraph, The Guardian, Los Angeles Times and The New York Times. We extracted the corresponding unexpected climate news index (UCNI) after assuming an auto-regressive model for the indices. We then performed several regressions of these UCNI using a LCI portfolio, a HCI portfolio and a LmHCI portfolio. We found that none of these indices displays a statistically significant climate beta at the 5% level for the LCI portfolio, 11 out of the 45 indices display a statistically significant climate beta at the 5% level for the HCI portfolio, and only two out of the 45 indices displayed a statistically significant climate beta at the 5% level for the LmHCI portfolio.

We then constructed nine different indices based on aggregating the UCNI across the five news sources. This resulted in greater significance when regressed against the returns of the set of asset portfolios. Out of the nine aggregate news indices obtained with that protocol, eight displayed a statistically significant climate beta at the 5% level for the HCI portfolio, six displayed a statistically significant climate beta at the 5% level for the LmHCI portfolio, and five displayed a statistically significant climate beta at the 1% level for the HCI portfolio. None displayed a statistically significant climate beta at the 5% level for the LCI portfolio. These results suggest that the climate news (UCNI) factor built on an aggregate index of high-quality newspapers has an explanatory power over the High Carbon Intensity portfolio returns, and hence over Low-minus-High Carbon Intensity portfolio returns. The improved significance of an aggregate index over individual indices suggests that individual newspapers may publish climate change articles even when there is no climate change event to report. By aggregating these newspapers, we reduce the importance of these idiosyncratic articles while retaining the importance of the climate articles which they all publish - those which report on actual unexpected climate news events.

Our findings agree with the general observation by Ardia et al. (2022) who find that LCI firms outperform HCI firms when there are unexpected increases in climate change concern. Unlike Ardia et al. (2022), we find that this is not because LCI stocks rise in value, but because HCI stocks fall in value. This result is consistent with the results of Bua et al. (2020) and is true for simple attention measures which do not study the article sentiment. This implies that unexpected climate news is generally bad for High Carbon Intensity assets, perhaps because related firms may have more to lose from climate change news than Low Carbon Intensity assets have to gain.

Furthermore, we find that the average return of LmHCI portfolios over the period from July 2012 to November 2021 consistently rises with the UCNI regime (low, medium, high) for all index types. This provides further evidence supporting the hypothesis that climate change concern plays a significant role in the long-term performance of Low-minus-High Carbon Intensity portfolios.

Out of all the language models used, the most advanced domain-specific ClimateBERT model did not materially outperform the simpler Attention-based model. This suggests that it is the number of articles, rather than their content, that drives climate risk awareness. It may also suggest that the ability of these state-of-the-art language models to extract sentiment from high quality newspaper articles is limited. This may be due to the complexity of the language found in these newspapers' articles. It may also be due to the desire of serious newspapers to be even-handed and moderate in tone.

There are several possible extensions of our paper. First, we may wish to add more individual news sources to the aggregate index to see the impact on the significance of the aggregate index. Second, it would be of interest to explore aspect-based sentiment approaches such as Peng et al. (2020). Using such an approach, we can ensure that the target of the expressed sentiment is indeed a climate change related matter. It may then be possible to distinguish between LCI and HCI targets. Third, it would be of considerable interest to determine whether the out-of-sample performance of the aggregate UCNI is sufficient to enable us to use it for portfolio hedging as initially proposed by Engle et al. (2020).

Table 6: UCNI-Conditional annualised performance of LCI, HCI and LmHCI portfolios for the aggregate news source and for headlines (H) and sentences (S).

Attention	LCI	HCI	LmHCI
low	21.8%	19.4%	2.4%
medium	33.3%	25.0%	8.2%
high	-0.5%	-10.4%	9.9%
Similarity	LCI	HCI	LmHCI
low	31.2%	24.6%	6.6%
medium	3.4%	3.5%	-0.2%
high	20.8%	5.4%	15.4%
Concern	LCI	HCI	LmHCI
low	24.7%	21.0%	3.7%
medium	25.2%	21.8%	3.4%
high	5.5%	-9.3%	14.8%
VAD-H	LCI	HCI	LmHCI
low	22.7%	20.3%	2.4%
medium	30.2%	23.3%	6.9%
high	2.5%	-10.1%	12.6%
VAD-S	LCI	HCI	LmHCI
low	23.1%	19.6%	3.5%
medium	30.0%	23.2%	6.8%
high	2.3%	-9.3%	11.6%
BERT-H	LCI	HCI	LmHCI
low	28.6%	24.6%	4.0%
medium	23.0%	15.0%	8.1%
high	3.8%	-6.1%	9.9%
BERT-S	LCI	HCI	LmHCI
low	20.4%	19.5%	0.9%
medium	34.0%	22.2%	11.8%
high	1.0%	-8.2%	9.2%
CBERT-H	LCI	HCI	LmHCI
low	21.1%	18.1%	3.0%
medium	28.0%	19.6%	8.4%
high	6.3%	-4.2%	10.5%
CBERT-S	LCI	HCI	LmHCI
low	27.6%	23.2%	4.4%
medium	27.1%	22.5%	4.5%
high	0.7%	-12.3%	13.0%

1. Introduction

Climate change is expected to have an impact on the value of financial assets over this century. It is therefore of interest to both investors and hedgers to be able to quantify their climate change exposure via a *climate beta*, just as an exposure to market risk is measured by the traditional market beta. However, measuring the *climate beta* is a challenging task as it is a sensitivity with respect to some proxy variable that captures changes in climate risk. For example Dietz et al. (2018) define the climate beta as the elasticity of the net benefit of an investment with respect to a change in aggregate consumption and estimate it using an extension of the 2013 version of the Dynamic Integrated Climate Economy Model (Nordhaus and Sztorc (2023)). Huij et al. (2022) define the climate beta with respect to the returns of a portfolio of so-called 'pollutive minus clean' assets. A recent and innovative approach to examining the impact of climate change on asset prices has been the work of Engle et al. (2020) who seek to use the unexpected arrival of climate news as an explanatory driver for the returns of climate-sensitive asset portfolios and hence as a proxy variable for a new climate beta measure. This has been followed by a number of related papers including Apel et al. (2021), Ardia et al. (2022), Bua et al. (2020), Faccini et al. (2021) and Meinarding et al. (2020).

In this paper, we also seek to estimate a climate beta using climate news. We construct climate change news indices based on high-quality newspaper sources and state-of-the-art language models. Our aim is to quantify the arrival of unexpected climate change news and to determine whether or not this has an effect on asset returns. The causal mechanism is straightforward - asset managers are influenced by the arrival of unexpected climate change news.[3] This could, for example, be a report of an extreme global warming event linked to climate change or the passing of some climate change legislation. This updates the information set of asset managers and may result in decisions to invest in a 'Low Carbon Intensity' asset, i.e., an asset that contributes positively to climate change mitigation or adaptation goals or disinvest from a 'High Carbon Intensity' asset, i.e. a climate damaging asset, with some associated market price impact. The nature of the media coverage in terms of the number of articles, their length and content may also determine the market impact. As described in McCombs and Shaw (1972) the amount of media exposure given to a topic has an influence on the public's view of the importance of that topic.

To analyse the content of climate news, we employ advanced Natural Language Processing (NLP) techniques. There has been a significant growth in the use of such techniques within finance and economics in recent years. As described in Gentzkow et al. (2019), these have been applied to a range of tasks including market prediction, fraud detection, asset allocation, credit scoring, the analysis of corporate earnings and the detection of sentiment. For our purposes, we require news sources that S&P 500 investors would be most likely to read. We therefore focus on high-quality, high-circulation, English-language news sources which publish daily. Ideally we should favour sources that focus on financial markets. We have therefore chosen to use the *Financial Times*, the *New York Times*, the *UK Daily Telegraph*, the *UK Guardian* and the *US Los Angeles Times*. Taken together, these provide us with a European and a US perspective. The Financial Times, with a daily circulation of around 500,000 copies and an estimated global readership of 1.9 million[4]

3 - Financial theory states that expected news is already embedded in market prices.
4 - See https://www.gale.com/intl/c/financial-times-historical-archive.

plays an especially important role in the dissemination of news information to the financial markets.

We can construct a climate change news index in several different ways. Our simplest Attention-based measure counts the daily number of climate change articles published. Such an approach was analysed in the media bias model of Gentzkow and Shapiro (2010). Our second approach, based on the Natural Language Processing (NLP) metric called Term Frequency - Inverse Document Frequency (TF-IDF), follows the approach of Engle et al. (2020). For the next level of sophistication, we use an approach that detects actual semantic meaning within the news articles by quantifying the degree of emotional concern. This is a challenging task as the emotional sentiment of long and highly articulate articles is not always simple to extract. For this reason, we turn to state-of-the-art language models such as the BERT[5] Transformer-based model from Devlin et al. (2018), and finally to the ClimateBert model by Webersinke et al. (2021) which has been specifically designed to better understand climate-related texts. We use these different approaches to construct a family of climate news indices (CNI), each a daily time series from 2005 to 2021. We do this for both the article content and the article headline. The headline is short and should indicate the sentiment of the article. As such it may be easier to extract the article sentiment from the headline than the entire article. To perform a market analysis using our climate change news indices we must isolate the unexpected component of the daily climate change news index changes. To do this, we assume that the CNI obey an auto-regressive AR(1) process where the changes in unexpected climate change news are the innovations. Calibrating to this process enables us to extract a family of unexpected climate change news innovation (UCNI) indices.

In addition to the set of CNI for each news source, we also construct a single aggregate index across all five news sources. Doing this increases the total number of articles being used in the construction of this index. It might be expected to increase the significance of climate news-related effects in the article counts. For example, a newspaper may decide on a specific date to publish a feature article on climate change that is not linked to a specific news event. It would almost certainly not coincide with a climate change article in any other news source. However a news story in all five news sources on a specific date is almost certainly driven by a common climate news event. By averaging over the news indices on each date, the aggregate index is more able to identify actual climate news events.

We examine the connection between these climate news indices and the performance of asset prices, specifically the top 500 US-listed firms by capitalisation (US500). To identify climate change risk in portfolios, we first assign each asset to one of the three categories of 'Low Carbon Intensity', 'Medium Carbon Intensity' and 'High Carbon Intensity'. A 'Low Carbon Intensity' asset is a security, which can be equity or debt, issued by a company that has low greenhouse gas (GHG) emissions per unit of revenue. Similarly, a 'High Carbon Intensity' asset is a security, equity or debt, issued by a company with high GHG emissions per unit of revenue. The remaining assets are labelled as 'Medium Carbon Intensity'. In this work, our aim is to determine whether the unexpected arrival

5 - BERT stands for Bidirectional Encoder Representations from Transformers.

of climate-related news has a differing impact on the returns of Low Carbon Intensity (LCI), High Carbon Intensity (HCI) and Low-minus-High Carbon Intensity (LmHCI) stock portfolios. By doing so, we test the two-factor model of Pástor et al. (2021) which predicts that LCI stocks have a lower expected return than HCI stocks due to investors having a taste for Low Carbon Intensity stocks and because they provide a better hedge against climate risk. This model also predicts that Low Carbon Intensity assets can have higher unexpected returns when *agents' demands shift unexpectedly in the LCI direction*. This can be a result of investors increasing demand for Low Carbon Intensity assets, consumers increasing their demand for Low Carbon Intensity products, and investors reducing their demand for High Carbon Intensity assets.

To categorise a company as LCI, HCI or Medium Carbon Intensity, we need to know the amount of major greenhouse gases[6] generated by the company per unit of revenue, also known as carbon intensity. Greenhouse gas emissions related to the activities and products of corporates are typically categorised into Scope 1, 2 and 3. Scope 1 emissions are those from sources owned or controlled by the company; Scope 2 are indirect emissions from the generation of purchased energy; and Scope 3 are other indirect emissions in the company's value chain. Obtaining accurate Scope 3 data across a broad asset universe is difficult as it requires a complete and detailed knowledge of the entire supply chain, plus a calculation of the carbon emissions associated with every type of intermediate output product. As discussed in Klaaßeen and Stoll (2021), reporting inconsistencies, boundary incompleteness and activity exclusions make reported numbers unsystematic and not comparable. Data providers who offer value chain emissions estimates typically take insufficient consideration of corporate circumstances to support comparisons, see Ducoulombier (2021), and are divergent, see Busch et al. (2022).

For this reason, most investment managers currently rely on (see Busch et al. (2022)) Scope 1 and 2 emissions to compute carbon intensity for decision making, and we will do the same. In this work we rely on asset portfolios constructed using ISS carbon intensities[7]. We analyse the interaction between the UCNI and various combinations of these company stocks including (i) Low Carbon Intensity (LCI) long-only portfolios, (ii) High Carbon Intensity (HCI) long-only portfolios, (iii) Low-minus-High Carbon Intensity (LmHCI) longshort portfolios and (iv) assets grouped by industry.

The first result of our analysis is that the relationship between the *individual* news source UCNI indices and the LCI, HCI, and LmHCI equity portfolio returns is overall not statistically significant. However the relationship does becomes significant for the HCI and LmHCI portfolios when we combine these indices across all five news sources to construct an *aggregate* UCNI.

We find that the *climate change beta*, i.e. the sensitivity of portfolio returns to an increase in the aggregate UCNI, is negative and statistically significant at the 5% level for HCI portfolio returns for eight out of nine language models and significant at the 1% level for five out of nine language models. It is small with varying sign and not significant at the 5% level for LCI portfolio returns. This result does not concur with the model of Pástor

6 - This quantity is expressed in CO2 equivalent according to their global warming potential over a reference period.
7 - We divide the carbon emissions by company revenue.

et al. (2021) which finds that the LCI and HCI stocks have an opposite exposure to an ESG risk factor. Our results suggest that it is the impact of unexpected climate news on the HCI portfolio that is the primary driver of the performance of a LmHCI portfolio. We also perform a conditional analysis in which we calculate the terciles of the daily UCNI for each news source and classify each day as low, medium or high. We also calculate the average return of the LmHCI portfolio under each regime. We find that conditional on the value of the daily UCNI index, the LmHCI portfolio average return increases with the level of the UCNI. Once again we find that this is driven by the fall in value of the HCI portfolio.

Not all climate change risks have the same characteristics. Here we divide climate change risks into two main categories - *physical* and *transition* risk. Physical risk relates to the impact on a company of event driven (acute) or longer-term shifts (chronic) in climate patterns. Examples include hurricanes, wildfires, extreme rainfall, droughts and sea-level changes. Such events can be value-destructive for industries such as agriculture. Flooding, storms and excessive temperatures may also threaten the functioning of firm assets such as real estate and industrial plants. The IPPC, see Pörtner et al. (2022), expects these risks to increase over time as the expected impact of climate change becomes more severe. Transition risk relates to the extensive policy, legal, technological and market decisions needed to address mitigation and adaptation requirements related to climate change. Such decisions can include the imposition of a carbon tax, revised investment mandates requiring funds to divest of assets with a high carbon intensity, and requiring that physical reserves of fossil fuels should remain untapped, becoming stranded. The transition may be orderly or disorderly. Using a set of lexicons that we generated, we construct both physical and transition risk indices. However, we find that these indices do not perform better than the pure climate change indices in terms of their stock return prediction performance.

This paper is organised as follows. In section 2 of this paper we set out a literature review of recent applications of NLP methods to constructing CNIs. In section 3 we describe the various index construction approaches that we have implemented. In section 4 we describe the tests we have performed to see if climate change news has an impact on the pricing of asset portfolios, and section 5 contains our conclusions.

2. Literature Review

Several studies have looked at the impact of news on prices (see Bollen et al. (2011), Chen et al. (2014), Hillert et al. (2014), Heston and Sinha (2017) inter alia). In terms of natural language processing (NLP), considerable recent progress in this field has made it possible to use models which go beyond the identification of key words, but which also have some internal representation of word semantics, and can perform well at tasks such as sentiment analysis and information extraction. This has had a number of benefits, including increasing the sophistication of financial analysis.

In recent years, a number of papers have examined the link between climate news and markets. Roughly speaking, each of these papers divides into two sections. The first section describes how to extract a climate news index (CNI) from digital news data; the second section attempts to determine whether this index has explanatory power for stock returns of companies with some measurable exposure to the physical and/or transition risks of climate change.

Although earlier papers such as Cody et al. (2015) construct a model to extract climatechange sentiment from climate news, the first attempt to link climate news information to the performance of climate-sensitive assets is by Engle et al. (2020). In this paper, the authors seek to create asset portfolios whose returns hedge the innovations in climate news out-of-sample. To do so, these authors create two climate news indices. The first uses the *Wall Street Journal* (WSJ) to calculate a time series of the similarity between a daily issue of the WSJ and a climate-change document compiled from a set of climate change-related sources. This is done using a TF-IDF metric[8]. Their second news index is a sentiment-based index generated by a third party[9] that measures the negative sentiment of climate-change news articles across time - it uses over 1,000 different news sources. Then the authors construct market, size, value and climate characteristic portfolios for a US equity universe. The climate characteristic portfolio is built using firm-level E-score data from MSCI and Sustainalytics. Companies with high E-scores are expected to have a lower exposure to climate risk while companies with low E-score are expected to have a higher exposure to climate risk.[10] The market, size and value characteristic portfolios are respectively built from the share of total market value, from standardised market value and from standardised book-to-market. Next, they regress the WSJ climate news index innovations CC_t on these portfolio characteristics. Using Sustainalytics E-scores, they find that, in-sample, these portfolios have a positive and significant relationship with CC_t. The in-sample R^2, which is a goodness of fit indicator of the linear model, is between 15% and 19%. Switching to MSCI E-scores reduces the R^2 to 8% – 9%. They also build hedge portfolios by regressing news index innovations on these characteristic portfolios and find that out-of-sample, using the WSJ climate news index (CNI) and Sustainalytics E-scores, the correlation between the CNI innovations and the hedging portfolio returns is 17%. However, with MSCI E-scores this correlation falls to 1%. The results improve slightly if they use the climate change negative-sentiment index instead of the WSJ news index.

Meinarding et al. (2020) propose and implement a text-based method to identify shocks to climate-change transition risk. They identify transition risk shocks as instances where

8 - TF-IDF is the Term Frequency Inverse Document Frequency metric. It measures how important a word is to a document relative to its importance in a collection of documents.
9 - Crimson Hexagon which has now merged with Brandwatch.
10 - Engle et al. (2020) acknowledge that the type of climate risk exposures captured by the E-score metric is not evident and that it might be more related to regulatory risks than physical risks.

a strong differential valuation of low carbon footprint versus high carbon footprint firms coincides with significant information on climate change. For that purpose, they combine information from long-short equity portfolios sorted on firms' carbon footprints with information from textual analysis of newspaper archives. The news is based on an analysis of news articles from 10 US newspapers, and the articles are filtered using the terms 'climate change' and 'economy' or 'economic'. They are able to identify seven transition risk shocks. They find that shocks increasing transition risk (negative abnormal returns of high carbon footprint firms) induce a decline in aggregate and sectoral industrial production, and that they significantly affect financial stability, as measured by the excess bond premium. Finally, they document a pronounced asymmetry in the economy's response to shocks increasing or decreasing transition risk.

A recent important theoretical contribution to our understanding of the relationship between LCI and HCI stocks, and the impact of unexpected information is provided by Pástor et al. (2021). They analyse the effects of sustainable investing using a two-factor equilibrium model, the factors being a market factor and an ESG factor. They argue that ESG preferences should move asset prices and that LCI stocks have negative alphas while HCI stocks have positive CAPM alphas. The negative CAPM alphas are linked to investor tastes for LCI holdings and the use of these assets for hedging climate risk. However the key conclusion that emerges from their approach is that LCI stocks can outperform HCI stocks if there are unexpected increases in ESG concerns.

In Ardia et al. (2022), the authors test the predictions of the model of Pástor et al. (2021). To do this, they seek to use news media to measure the unexpected changes in climate change concerns. To measure 'concerns', the authors rely on two lexicons - a lexicon of risk-related words and lexicon of sentiment-related words.[11]. Using US newspapers, they calculate an index based on the excess of negative over positive words weighted by risk concerns[12]. They then construct from the S&P 500 universe LCI, HCI and Low-minus-High Carbon Intensity portfolios[13] to see if LCI firms outperform HCI firms when unexpected media concerns (UMC) increase. To buid these portfolios, stocks are ranked on their carbon intensity as taken from the Asset4/Refinitiv database. LCI firms are those in the top quartile and HCI portfolios are in the bottom quartile. These equally-weighted portfolios of LmHCI assets are re-balanced daily. They perform a factor analysis using the five Fama and French (2015) factors to control for other factors that may drive stock returns. They find that on days with an unexpected increase in climate change concerns, LCI firm stock prices tend to increase while HCI firm stock prices tend to decrease.

In Apel et al. (2021) the authors create an index based on news article headlines extracted via the data provider Ravenpack from a range of reputable news sources. Their index measures the daily overlap between the news headline and a domain-specific vocabulary that can change over time. Using the TF-IDF metric, they calculate a 'climate weight' for each headline. They use the BERT machine learning model for natural language processing introduced by Devlin et al. (2018), and fine-tune it for sentiment classification on 25,000 news headlines from the year 2000 to 2008. These have been hand-annotated with

11 - These lexicons are taken from the LIWC(2015) software project
12 - They capture risk concerns by counting the number of risk words contained in each article.
13 - In contrast to Ardia et al. (2022) who refer to Low Carbon Intensity portfolios as 'Green' portfolios, refer to High Carbon Intensity portfolios as 'Brown' portfolios and label long/short Low-minus-High Carbon Intensity portfolios as Green-minus-Brown (GMB) portfolios, we refer to them as 'Low Carbon Intensity' (LCI) portfolios, 'High Carbon Intensity' (HCI) portfolios and Low-minus-High Carbon Intensity (LmHCI) portfolios, respectively.

respect to their implied impact on transition risk. They train and test their model on their own labelling and obtain an F1 score of 0.82 with 0.74 on the fine-tuned BERT model[14]. This model outputs a value for negative, neutral, and positive sentiment for each headline. The result is a transition risk index (TRI). They then calculate innovations in their TRI as residuals from a simple autoregressive moving average model ARMA(1,1) model[15]. They regress the returns of decarbonised and pure-play market indices[16] on TRI and the 5 Fama and French (2015) factors and find that transition risk tends to affect stock prices based on firms' business activity but not their emissions.

Bua et al. (2020) study the pricing of climate risk in European equity markets. They use NLP methods to construct a physical risk (PR) and transition risk (TR) indicator using Reuters News as source. They look at physical and transition risk using separate lexicons. They add the PR and TR indices to a Fama and French (2015) five factor model and seek to determine if there is a climate risk premium. They find that when stocks are sorted according to carbon intensity scores, increases in transition risk increase excess returns of LCI stocks, while decreasing the excess returns of HCI stocks.

Finally, Faccini et al. (2021)study whether market-wide physical or transition climate risks are priced in US stocks. They perform a textual and narrative analysis of Reuters climate change news over the period 2000-2018. They use Latent Dirichlet Allocation (LDA)[17] to identify the main topics in the articles (see Blei et al. (2003) for more details on the LDA approach). Initially, 25 topics are detected, and these can be loosely associated with themes such as natural disasters or fossil fuel emissions. They identify four topics with a clear interpretation which are related to natural disasters, global warming, international summits, and US climate policy, respectively. Only the climate policy factor is priced, especially post-2012. They find that the documented risk premium is consistent with investors hedging the imminent transition risks from government intervention, rather than the direct risks from climate change itself. They also find that firms that are most exposed to transition risks tend to be polluting businesses which show no strong intention of becoming greener.

14 - The F1 score is a measure of a model's accuracy. It is calculated by taking the harmonic mean of the model's precision and recall. Precision is the fraction of the positive predictions made by the model that are correct: it is calculated by dividing the number of true positive predictions made by the model by the total number of positive predictions made by the model. Recall is the fraction of the total number of relevant items (for instance number of spams, number of LCI stocks,...) that the model is able to predict. It is calculated by dividing the number of true positive predictions made by the model by the total number of relevant items. The F1 score is a balance between precision and recall. It is generally used when one wants to weigh both of these metrics equally. A high F1 score indicates that the model is accurate and able to predict a high number of relevant items.
15 - An ARMA(1,1) is a statistical model that combines both an autoregressive model and a moving average model, so it uses both past values of the time series and the residual errors from past predictions to make a prediction for the current time step. The order of the autoregressive model, represented by the '1' in ARMA(1,1), refers to the number of lag terms included in the model. So, an AR(1) model would use the value of the time series at the previous time step to predict the current value. The order of the moving average model, represented by the '1' in ARMA(1,1), refers to the number of lag terms included in the model. So, an MA(1) model would use the residual error from the previous time step to predict the current value.
16 - The authors define directly pure-play approaches as allocations toward "companies that generate a significant share of revenues from products and services related to environmental challenges". See Table A.1. of their paper for more details.
17 - LDA is a generative probabilistic model used to discover latent topics in a collection of documents. It assumes that each document is a mixture of a small number of topics, and each topic is a distribution over words.

3. Constructing Climate Change News Indices

We now specify the news sources used in our analysis and then detail the full construction details of our climate change news indices (CNI) which are then used in Section 4.1 to generate our unexpected climate news indices (UCNI).

3.1 News Data

We require digitised, daily, English-language, high-circulation and high-quality news sources with a European[18] and US perspective, including at least one with a dedicated financial market focus. For these reasons we have chosen the following news data sources:
1. *The Financial Times* (FT) digital archive. The Financial Times is widely recognised as the leading European English-language financial newspaper. It is published daily from Monday to Saturday and covers not just business news, but also world politics and current affairs.
2. The Lexis Nexis (LN) database of newspapers. This provides access to many thousands of newspapers internationally. From these we selected the *The New York Times* (NYT), the *Los Angeles Times* (LAT), *The Guardian* (UKG) and *The Daily Telegraph* (DT).

We use daily articles from these news sources over the period from 3 January 2005 to 3 November 2021. For each article, we have the date of first publication, the headline and text body. We assume that an article appears in the morning of the first publication date. This can be a different time depending on whether the news source is in the US or Europe. As Europe is several hours ahead of the US, the arrival of news from US and European news sources will impact the US equity market on that same day. To align the arrival of weekend news with the financial markets, we adjust the publication date of news stories that appear on a Saturday or Sunday to the following Monday, the earliest date on which this news can impact the US equity market.

We consider news in the form of newspaper articles hence each article has both a headline and content. The headline is typically added by a sub-editor who has read the article and wishes to summarise the key features in order to attract the reader's interest. The headline will usually reflect the most important part of the article and any associated positive or negative sentiment. For this reason, we will examine both the article headline and the article content to see if the headline can provide a clearer measure of article focus and sentiment than an analysis of its longer and more complex content.

To quantify the newspaper media's attention to climate change, we first need to identify a 'climate change article'. Using unigram search terms can lead to false positives. Hence we search for bigrams - combinations of two words - that ensure the subject matter is related to climate change. Work by Engle et al. (2020) and others has done this using the search term 'climate change'. However the bigram 'global warming' has also been widely used as a synonym. To determine whether we should include it, we perform searches for articles that contain (i) 'climate change', (ii) 'global warming', and (iii) one or both of 'climate change' or 'global warming'. We focus on the *Financial Times* news source and calculated the fraction of daily articles that are returned by these search terms. Using the fraction of articles rather than the number corrects for the fact that the total number of daily *Financial Times* articles has varied significantly over this period.[19] All three time series are shown in Figure 3. We observe that use of the bigram 'global warming' has

18 - Even though the study focuses on a US stock universe, it can be influenced by global events and news from Europe.
19 - Over the 2005-2021 time period the monthly number of Financial Times articles ranged from around 2,000 in 2006 to over 6,000 in 2014 and then back down to around 3,000 in 2020.

declined in relative terms over time but it is still used. The bigram 'climate change' has clearly become the dominant bigram. To ensure that we capture as many climate change articles as possible, and especially those in the earlier period of analysis, we include both in our definition of a 'climate change article'.

Figure 3: We measure newspaper media attention to the subject of climate change from 2005-2021 by calculating the percentage of daily FT articles that contain the term 'climate change', the percentage that contain the term 'global warming' and the percentage containing either term. We apply a 30-day moving average. The vertical lines are climate change-related events - see Table 7 for the corresponding numbered list of events.

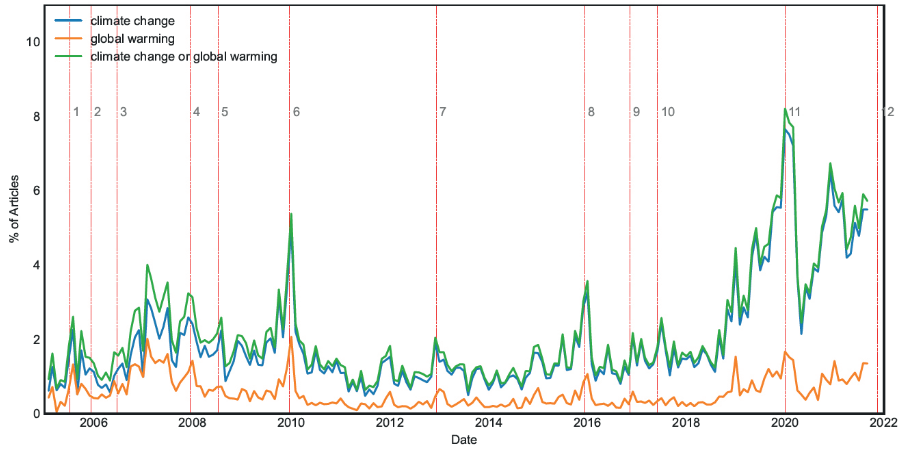

Table 7: Identification of the most active climate news events seen in Figure 3. For conferences we have used the conference end date when the final agreement is usually announced.

Event #	Date	Description
1	8 July 2005	G8 Summit Scotland
2	9 Dec 2005	Montreal CC Conference
3	21 Jun 2006	UK CC and Sustainable Energy Act 2006
4	14 Dec 2007	Bali UN CC Conference
5	9 Jul 2008	G8 Summit
6	19 Dec 2009	Copenhagen UN CC Conference
7	7 Dec 2012	Doha UN CC Conference
8	12 Dec 2015	Paris Agreement Signed
9	7 Nov 2016	Marrakech UN CC Conference
10	1 Jun 2017	US President Trump Withdraws from Paris Agreement
11	Dec 2019 - Jan 2020	Australian Wildfires, High Temperatures
12	13 Nov 2021	Build up to COP 26 CC Glasgow

As we might expect, Figure 3 also shows that a high level of climate attention is closely linked to highly newsworthy climate-related events such as IPCC conferences, G8 summits and the signing of climate change treaties. A list of such events is provided in Table 7. It is evident that climate attention has increased significantly since 2016, driven by a number of events including greater climate change activism, Australian wildfires and preparations for the UN Climate Change Conference (COP26) in Glasgow.

We finally analyse the difference between the different news sources. Table 8 shows the counts of the total number of selected articles per newspaper, per year. We see that the UK *Guardian* is the leading publisher of climate change-related articles over time among our corpus of news sources, followed by the *Financial Times* and the *Daily Telegraph*.

3.2 Constructing Climate Change News Indices (CNIs)

We explore several approaches for constructing a climate change news index (CNI) from newspaper articles. If there is a link between climate change news and market price movements, then we would expect that the link will be strongest for the index that best captures the quantity, content and sentiment of the climate change news. The index construction approaches we will use, in order of increasing level of sophistication, are as follows:

1. Attention - The number of climate change articles published each day.
2. Similarity - The similarity between the newspaper's climate change articles of the day and a representative climate change document (TF-IDF cosine similarity).
3. Concern - Climate change concern using word-frequencies based on the Linguistic Inquiry and Word Count (LIWC) lexicons, which are designed to capture social and psychological states.[20]. Concern is high if the number of 'negative words' in a climate change article is higher than the number of 'positive words' and the fraction of 'risk words' is high.
4. VADER - We use a lexicon and ruled-based sentiment analysis tool called VADER (for Valence Aware Dictionary and sEntiment Reasoner) that assigns a sentiment polarity score to specific words to determine if the climate change article sentiment is positive or negative.
5. BERT model with Fine-Tuned Sentiment - We take a BERT language model as described in Devlin et al. (2018), and fine-tune it to identify sentiment using humanlabelled, finance-related training examples.
6. ClimateBERT with Fine-Tuned Sentiment - We take the domain-specific Climate-BERT model by Webersinke et al. (2021) and fine-tune it to identify sentiment using human-labelled, finance-related training examples.

Table 8: Number of climate change articles by year for the Financial Times, The Daily Telegraph, The Guardian, Los Angeles Times and The New York Times news sources. Note that 2021 only includes articles up to 3 November 2021.

Year	FT	DT	UKG	LAT	NYT
2005	339	413	1,127	313	294
2006	523	645	1,673	611	385
2007	1,304	1,241	2,676	1,053	1,198
2008	1,113	723	2,499	762	997
2009	1,495	885	3,914	673	1,169
2010	1,093	1,023	2,480	548	666
2011	706	783	1,932	329	334
2012	739	899	1,853	306	287
2013	800	822	1,914	339	332
2014	800	639	2,284	468	399
2015	1,212	715	5,278	771	600
2016	892	429	4,881	660	532
2017	1,003	465	2,002	776	491
2018	1,062	566	2,570	697	389
2019	2,040	1,393	4,051	984	618
2020	1,909	1,180	3,091	697	485
2021	2,218	1,866	3,718	1,076	701

20 - See https://www.liwc.app/.

We now explain the detailed construction of these indices. To help facilitate a comparison between the indices, we construct them such that a higher value will be associated with greater negative concern or greater negative sentiment, and vice-versa. In the first instance we will compute indices which look at climate-change risk. We will then consider indices that seek to differentiate between physical and transition risk. To make clear the index construction methodology, we define the following notation:

- The number of articles in newspaper *b* on day *t* is given by $N^b(t)$.
- The number of articles in newspaper *b* on day *t* containing the term *w* is $N^b_w(t)$).

For example, the three series consisting of the daily fraction of climate change articles where *b* denotes the *Financial Times* news source, and which are shown in Figure 3, are defined as:

$$CCA^{FT}_{cc}(t) = \frac{N^{FT}_{cc}(t)}{N^{FT}(t)} \text{ and } CCA^{FT}_{gw}(t) = \frac{N^{FT}_{gw}(t)}{N^{FT}(t)} \text{ and } CCA^{FT}_{cc,gw}(t) = \frac{N^{FT}_{cc,gw}(t)}{N^{FT}(t)},$$

where the subscripts *gw* and *cc* refers to articles containing the bigrams 'global warming' and 'climate change' respectively, and the superscript *FT* refers to the source as the *Financial Times*.

3.2.1 Climate Attention Index

Our first CNI is the climate attention index. This counts the number of climate change articles published each day by each source. Such an index is pertinent if readers are influenced by the number of articles on climate change, even if the total number of articles is growing or declining. In an age of online as opposed to physical newspapers, a climateaware reader is not easily able to assess the total number of articles published daily by a news source, but can be aware that the number of climate change-related articles is rising or falling. Mathematically, the daily attention index is given by

$$VI^b(t) = f(N^b_{cc,gw}(t)).$$

The function $f(x)$ is used to model the 'market impact' of the number of daily articles. It may be that a doubling of articles does not result in a doubling of market-perceived climate attention - the marginal impact of each additional article is decreasing. In this case we would expect $f(x)$ to be a concave function. We will examine this by trying two functional forms, (i) $f(x)$ =√ x and (ii) $f(x) = x$. Figure 4 shows the evolution of the attention index over time for each of the news sources. We see, for example, that *The Guardian* is a more frequent publisher of climate change articles than the other newspapers, and with a greater volatility.

Figure 4: 30-Day moving average of the climate attention CNI for the Financial Times, The Daily Telegraph, The Guardian, Los Angeles Times and The New York Times news sources from January 2005 to November 2021. The y-axis corresponds to the square root of the number of daily climate change articles published by each news source.

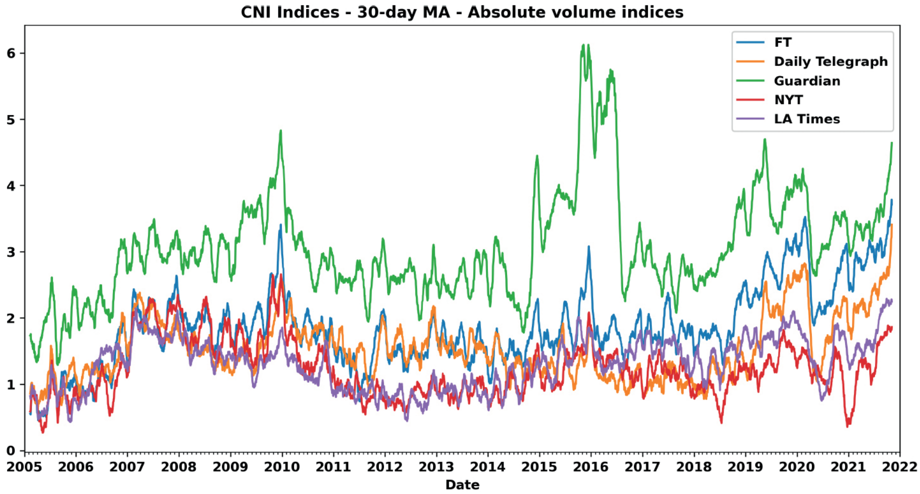

3.2.2 Similarity-Based Index

An attention-based index simply counts articles, but it cannot capture the fact that two climate change articles may not have the same degree of focus on the issue of climate change, and so should not be weighted equally in a CNI. For this reason, we adopt a TF-IDF cosine similarity-based approach to measure the climate news importance of an article. The Term Frequency - Inverse Document Frequency (TF-IDF) approach uses a vector-based representation of a document. It is intended to capture the importance of individual words in that document relative to the importance of those words across the entire corpus of documents - it is explained in greater detail in Appendix 6.2. Two similar documents should have a similar TF-IDF vector. So if we use the TF-IDF approach to compare a news article to a reference climate change document (CCD), we can measure their similarity and hence the degree to which the news article is about climate change.

For this to work, the climate change document must be representative of a document that discusses a broad range of climate change-related issues, capturing the entire vocabulary used within this topic domain. To construct such a climate change document, we extract 26 authoritative articles on the subject of climate change. These include IPCC[21] reports ranging from 2000 to 2021, as well as some documents produced by the US Environmental Protection Agency (EPA) and the International Monetary Fund (IMF). The climate change document was constructed by concatenating all of these individual documents, which are listed in Table 24, to make one large document[22].

In this case, we use a vocabulary of climate-related bigrams that cover both physical risk and transition risk. To calculate the similarity index at date t, we combine the climate change related articles for each day into a single document and consider its cosine similarity with the CCD. The resulting indices for each news source are shown in Figure 5. The approach we use to calculate the cosine similarity measure is in line with Engle et al. (2020) as we do not compute it individually for each article and then sum over the daily count. Additionally, we do not incorporate a function $f(x)$ to adjust the resulting index

21 - IPCC stands for Intergovernmental Panel on Climate Change
22 - The pdf files were downloaded, then converted to plain text format ensuring that we correct for any special characters, and then concatenated into a single text format document with a plain text format size of approximately 50MB.

value as we do not consider the number of daily articles published on climate change in our index calculation.

Figure 5: 30-Day moving average of the similarity-based CNI for the Financial Times, The Daily Telegraph, The Guardian, Los Angeles Times and The New York Times news sources from January 2005 to November 2021.

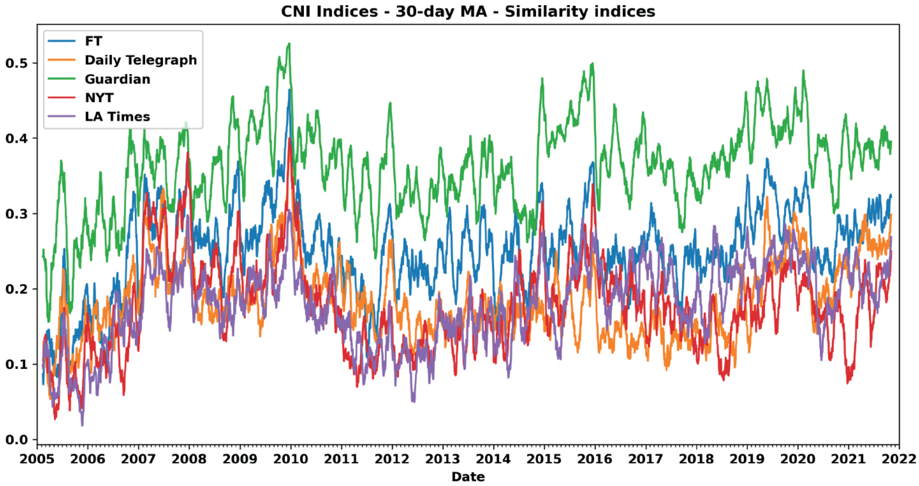

Specifically, for each news source b we construct a corpus $\mathcal{D}^b$, containing each climate change-related article from 2005 to 2021 plus the climate change document (CCD). On each date t, we calculate the TF-IDF vector for each daily climate article issue and the CCD using a vocabulary of climate-related bigrams that cover both physical risk and transition risk. We then compute[23] the cosine similarity $CS^b_{\mathcal{V}}(t)$ for news source b between the climate change related issue on date t and the climate change document (CCD), using the vocabulary V. We define this as the similarity-based CNI.

3.2.3 Concern-Based Index

It has long been argued that markets are influenced by sentiment, for example by Nobel Prize-winning economist Shiller (2000). While humans can easily detect sentiment, the machine learning of sentiment is a very challenging linguistic task. Synonyms, bad syntax, double negatives and sarcasm are just four of the complicating factors. However, the considerable growth of social media and a wish to analyse the content of the resulting blogs, tweets, Facebook posts, film reviews and product-reviews has led to significant progress in language modelling and analysis. The method used for this index is based on the use of specific human-created lists of words, known as 'lexicons' that convey emotional information.

The 'concern' measure is the emotion that drives our first semantic climate change news index. This index seeks to quantify the degree of concern of climate change news articles. This approach is similar to that used by Ardia et al. (2022), who defined *concern* as 'the perception of risk and related negative consequences associated with this risk'. We also use the same Linguistic Inquiry and Word Count (LIWC) sentiment lexicon to detect the sentiment. This is a dataset consisting of over 100 lexicons, designed to contain words that identify someone's social and psychological state. These lexicons encompass a range of human experiences including positive emotions and negative emotions, anxiety,

23 - The cosine similarity between the respective TF-IDF vectors $\mathbf{T}^b_{\mathcal{V}}(t)$ and $\mathbf{T}_{CCD,\mathcal{V}}$ of the daily climate article issue of news source b and the CCD is given by

$$CS^b_{\mathcal{V}}(t) = \frac{\mathbf{T}^b_{\mathcal{V}}(t) \cdot \mathbf{T}_{CCD,\mathcal{V}}}{\|\mathbf{T}^b_{\mathcal{V}}(t)\| \times \|\mathbf{T}_{CCD,\mathcal{V}}\|}$$

risk, anger and sadness. Example positive emotion words are *love*, *nice*, *sweet*, negative emotion words are *hurt*, *ugly* and *nasty*, and risk words are *danger* and *doubt*[24]. One benefit of this approach over more sophisticated machine learning approaches is that the results can be easily understood as the lexicons can be inspected.

We calculate the negative concern index (*NCI*) using the same method as Ardia et al. (2022). We count the number of negative emotion words (*NW*) minus the number of positive emotion words (*PW*) divided by the sum of positive and negative emotion words.

This is then scaled by the fraction of risk words (*RW*) divided by the number of words in the article (*N*). The fraction of risk words divided by the total number of words in the article provides a measure of the proportion of text that contains language related to potential harm, danger, or uncertainty. Scaling the difference between negative and positive emotion words by this fraction serves to weight the emotional valence of the text in relation to the presence of risk language. So the concern of article *i* published on day *t* from a single source *b* is given by

$$NCI_i^b(t) = 100 \cdot \left(\frac{RW_i(t)}{N_i^b(t)}\right) \cdot \frac{1}{2} \cdot \left(\frac{NW_i^b(t) - PW_i^b(t)}{NW_i^b(t) + PW_i^b(t)} + 1\right)$$

Figure 6: 30-Day moving average of the negative concern-based CNI for the Financial Times, The Daily Telegraph, The Guardian, Los Angeles Times and The New York Times news sources from January 2005 to November 2021 with $f(x)$=√x.

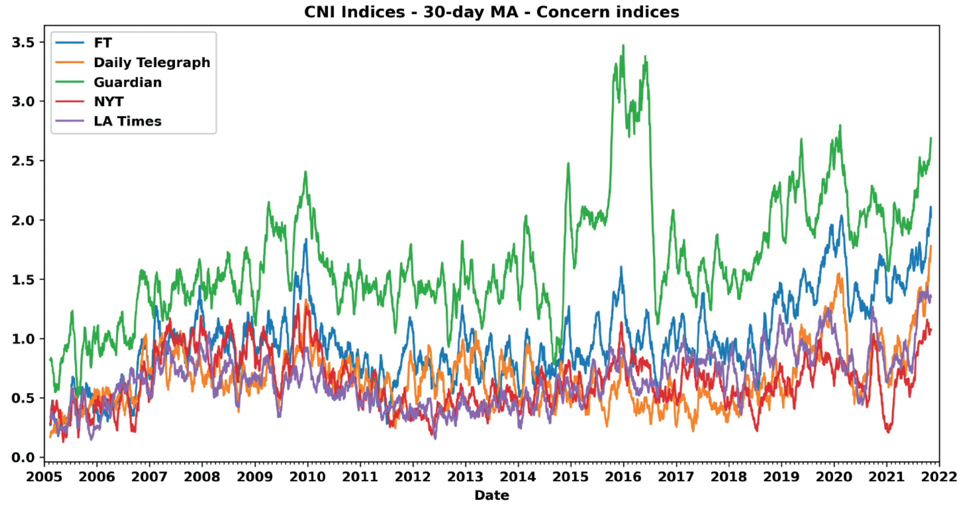

The concern measure is based on word counts at an article level. To make this a daily index, we sum the article-level concern measure over all $N_{cc}^b(t)$ climate change articles on day *t* and pass it through a function $f(x) = \sqrt{x}$ or $f(x) = x$ as previously defined. The resulting index value is given by:

$$NCI^b(t) = f\left(\sum_{i=1}^{N_{cc}^b(t)} NCI_i^b(t)\right)$$

We end up with a daily negative concern index for each day *t* and news source *b*. Note that on a day with no climate change news articles, when $N_{cc}^b(t) = 0$, the value of the index is zero, corresponding to the lowest level of negative concern. The resulting concern indices are shown in Figure 6.

24 - These lexicons have been created by experts in the field of psychology and have been extensively employed in over 20,000 published research papers. See https://www.liwc.app/ and https://mcrc.journalism.wisc.edu/files/2018/04/Manual_LIWC.pdf for more details.

3.2.4 Rule-Based Sentiment Index

The previous approach examined concern based on the number of positive and negative emotion words. However, counting positive words and negative words, without context, may produce an incorrect result. This is because negations, sarcasm and other grammatical complexities can weaken or even reverse the actually sentiment implied by the preponderance of positive over negative words. Another issue is that not all positive or negative words have the same emotional intensity, and it would be helpful to measure the strength or *valence* of the emotion. We therefore choose to apply the VADER model of Hutto and Gilbert (2014). VADER is a rules-based model that goes beyond the simple dictionary approach of LIWC - it has its own lexicon of several thousand words, each of which are each labelled with a valence score[25]. Furthermore, it also implements a set of heuristic rules that includes more sophisticated language features such as negation, punctuation, capitalisation and modifiers. The valence scores of the words range from −4 to +4, with −4 corresponding to extremely negative sentiment words, and +4 to extremely positive sentiment words. The authors of VADER show that it outperforms lexicon-based approaches such as LIWC on sentiment in a social setting.

Another criticism of previous approaches is that a sentence in a long climate change news article that contains emotional words may not be referring to climate change, but to another subject. For this reason we wish to ensure that the expressed emotion is in close proximity to the discussion of climate change within the article. To do so, we no longer determine the sentiment of the entire article as a whole, but we examine the sentiment of only those sentences that we have identified as being climate-change related, i.e. they contain at least one of the terms in our climate change lexicon. This a list of climate change bigrams we have constructed. It is formed by the union of the physical and transition risk lexicons described in more detail in section 3.3. As a result, we use VADER to determine the sentiment of each climate change sentence, rather than the entire article. We then aggregate these sentence sentiments as described below to compute the article sentiment.

This is the first CNI index for which we have calculated a sentiment for both the headline and the content of an article. Although VADER's word valences are in the range [−4, +4], the valence score for each sentence is normalised to be in the range [−1, +1], with +1 being the most positive sentiment. To ensure consistency with our earlier indices, we map[26] this valence score so that it is in the range [0, 2], with 2 representing the most negative sentiment. We calculate both a headline index and a sentence-based index. The headline sentiment is calculated as follows:

1. For each date t and news source b we select all $N^b_{cc}(t)$ climate change articles.
2. For each of $i = 1, \dots, N^b_{cc}(t)$ articles, we use its headline to obtain the corresponding VADER headline sentiment $S^H_{i,b}(t)$.
3. We calculate the sum of the headline sentiments over all the climate change articles on that date and then scale it using the function f(x) as follows:

$$S^H_b(t) = f\left(\sum_{i=1}^{N^b_{cc}(t)} S^H_{i,b}(t)\right)$$

25 - A valence score is a measure of the positive or negative emotional content of a word, phrase, or sentence. It is a numerical representation of the emotional tone conveyed by a piece of text, ranging from negative to positive.
26 - We simply map $s \rightarrow 1 - s$.

To build the sentence sentiment index time series $S_b^S(t)$, the process is slightly more complicated and we proceed as follows:
1. For each date t and news source b we select all $N_{cc}^b(t)$ climate change articles.
2. For each climate change article i we identify the $N_{cc,i}^b(t)$ climate change sentences that contain words in our climate change lexicon[27]. For each of these sentences j we calculate the corresponding VADER sentiment $S_{i,j,b}^S(t)$
3. We calculate the average article-level sentiment index by summing over the individual climate-related sentences

$$S_{i,b}^S(t) = \frac{1}{N_{cc,i}^b(t)} \sum_{j=1}^{N_{cc,i}^b(t)} S_{i,j,b}^S(t).$$

4. We then calculate daily VADER sentence sentiment for news source *b* by summing over all of the article sentiments and then apply the function $f(x)$.

$$S_b^S(t) = f\left(\sum_{i=1}^{N_{cc}^b(t)} S_{i,b}^S(t)\right)$$

The result is a daily VADER sentiment index for headlines and sentences for each news source. This is shown in Figure 7.

Figure 7: 30-Day moving average of the VADER sentence CNI for the Financial Times, The Daily Telegraph, The Guardian, Los Angeles Times and The New York Times news sources from January 2005 to November 2021.

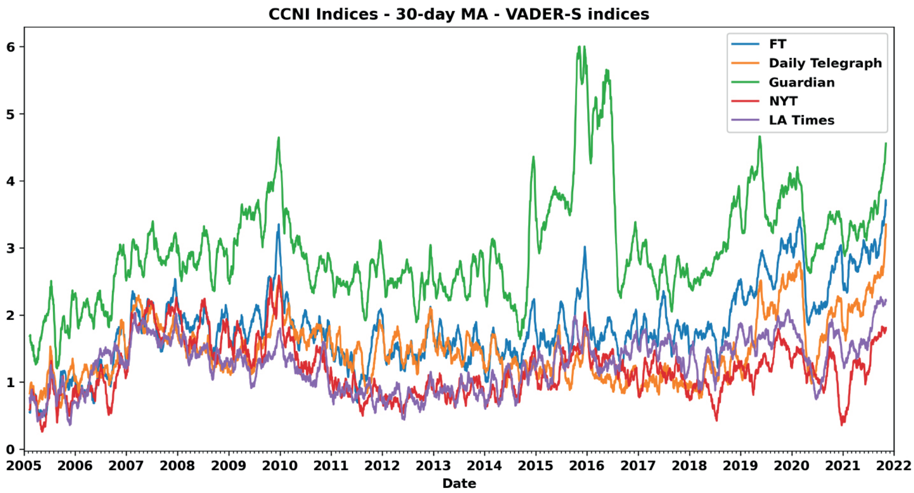

3.2.5 Deep Sentiment-Based Index

Extracting accurate sentiment from high quality news articles is challenging for two reasons. First, the high quality of writing in reputable news sources such as the Financial Times, which often consists of long semantically sophisticated sentences, makes it hard for rulesbased approaches to determine an overall sentiment. Second, the style of the writing in such newspapers tends to be sober and balanced, without significant expressions of strongsentiment. To address the challenge of extracting sentiment from such challenging texts, we apply Google's BERT[28] model, first presented in Devlin et al. (2018). BERT is a sophisticated language model that produced state of the art performance on a number of language understanding tasks.

27 - This is a single set of words (mostly bigrams) that relate to climate change. We constructed this lexicon by concatenating the physical and transition risk lexicons whose construction we describe in detail in section 3.3.
28 - BERT stands for Bidirectional Encoder Representations from Transformers.

BERT is a deep learning neural network model with an internal *transformer* architecture that has been specifically designed for the purpose of contextual language comprehension. It has been trained on a dataset of 11,038 unpublished books, plus the entire corpus of English Wikipedia (about 2.5 billion words). This is described in more detail in Appendix 6.3.

Indeed, BERT is a large model. Even the simplest BERT_{BASE} uncased[29] model has 110 million parameters and so training it is computationally challenging, typically requiring many hundreds of hours of CPU time. Fortunately, we are not required to train BERT ourselves, and can instead access a pre-trained model provided by Hugging Face[30].

However, a pre-trained BERT model cannot immediately perform text sentiment analysis. We are required to *fine-tune* the BERT model for this task. Specifically we must add a further layer to the BERT_{BASE} architecture and then fine-tune the model for the task of sentiment analysis by training it to learn as set of example input-output pairs. Since the BERT model has already been trained to understand language and context, it is found that the number of examples required to learn an extra language task to an acceptable accuracy should not be as many as we would need to train a model from scratch. Known as *transfer learning*, this is a very powerful feature of deep learning models.

Table 9: Excerpts from three examples sentences from the Financial Phrase-Bank of Malo et al. (2014) with the associated sentiment label and our assigned sentiment score. The higher the score the more negative the sentiment.

Sentences	Sentiment (Score)
However, the growth margin slowed down due to the GFC	Negative (2)
... generated net sales of 7.5 mln euro ... in 2005	Neutral (1)
... revenues have risen on an average by 40%	Positive (0)

For our training set we used the human-annotated Financial Phrase-Bank of Malo et al. (2014) to fine-tune BERT for sentiment. This consists of 4,840 sentences of English financial news categorised by sentiment. It has been divided into subsets according to the degree of agreement between the label human annotators and from these we use only the 4,217 sentences for which at least 66% of the human annotators agreed on the sentiment score. See Table 9 for an illustrative example of three sentences from this dataset and their labelled sentiment. We discuss the details of the fine-tuning in Appendix 6.3. We fine-tune the BERT model to output a numerical sentiment that is highest when detected sentiment is most negative. Following the approach used previously, we define three output states which we number as 2 (negative), 1 (neutral) and 0 (positive). We performed the fine-tuning using the PyTorch deep learning library3[31]. After training for 10 epochs[32] on the training set, we found that the cross-validated $F1$ score on the test set equaled 88.75%, which means the program did a good job at classifying the sentiment of the new examples.

Using this fine-tuned BERT_{BASE} sentiment model, we can then construct sentiment indices for each news source *b*, with one for the overall content and one for the headline.
To build the headline sentiment index time series $S_H^b(t)$, we proceed as follows:
1. For each date *t*, source *b* and article *i*, we use the headline to obtain the corresponding ERT_{BASE} headline sentiment $S_{i,b}^H(t)$.

29 - The model does not distinguish between lower and upper case text.
30 - https://huggingface.co/
31 - For the choice of gradient descent algorithm, we chose *AdamW* with a learning rate of 1e − 5 which is a measure of how quickly the algorithm adjusts the parameters of the model, and a numerical stability term equal to 1e−8 which helps to prevent numerical errors that can occur when working with very small numbers. We chose a batch number of 32, which means that the model is trained on 32 examples at a time.
32 - An epoch corresponds to the complete pass of the training dataset through the algorithm and the resulting updates of the network weights and biases using backpropagation.

2. We take the square root of the sum of the headline sentiments over all the climate change articles on that date as follows

$$S_b^H(t) = f\left(\sum_{i=1}^{N_{cc}^b(t)} S_{i,b}^H(t)\right)$$

To build the sentence-level sentiment index time series $S_b^S(t)$, the process is slightly more complicated and we proceed as follows:

1. For each date *t*, source *b*, and article *i*, we identify the $N_{cc,i}^b(t)$ climate change sentences. For each sentence *j* we calculate the corresponding BERT$_{BASE}$ sentiment $S_{i,j,b}^S(t)$.

2. We average over all the climate change sentences in each article to obtain the article level sentiment index

$$S_{i,b}^S(t) = \frac{1}{N_{cc,i}^b(t)} \sum_{j=1}^{N_{cc,i}^b(t)} S_{i,j,b}^S(t)$$

3. Finally, we calculate theBERT$_{BASE}$ sentence sentiment index on day *t* for news source *b* as

$$S_b^S(t) = f\left(\sum_{i=1}^{N_{cc}^b(t)} S_{i,b}^S(t)\right).$$

From this we end up with a daily time series of BERT$_{BASE}$ sentiment for both headlines and sentences. The sentence CNIs are displayed in Figure 8.

Figure 8: 30-Day moving average of the BERT sentiment-based CNI for the Financial Times, The Daily Telegraph, The Guardian, Los Angeles Times and The New York Times news sources from January 2005 to November 2021.

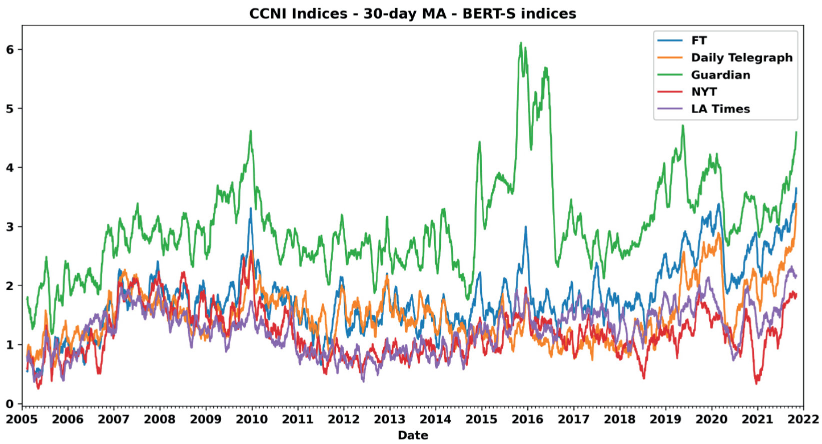

3.2.6 Domain-Specific BERT with Fine-Tuned Sentiment

Although the BERT model has been trained on a huge corpus of English-language texts, only a very small proportion of these could be characterised as 'climate change' texts. It has therefore no special expertise in the climate change domain, and so may not fully understand words with ambiguous context-dependent meanings such as 'climate' and 'atmosphere'. There is therefore some potential benefit to be gained from using a *domain-specific* BERT model - one has been trained on a corpus that relates to the subject domain within which it will be used. One such model is ClimateBERT by Webersinke et al. (2021). This has been trained on 1.6 million climate-related paragraphs and is claimed to be 46% more accurate on climate-related tasks than the general BERT model.

A pre-trained ClimateBERT model has been made available at Hugging Face[33]. It is based on the DistilRoBERTa base model architecture from Sanh et al. (2019). Its architecture is 40% smaller than the $BERT_{BASE}$ model and is estimated to be 3% less accurate but 60% faster to train and predict. As with $BERT_{BASE}$, we must specifically fine-tune ClimateBERT to perform sentiment analysis. Once again we do this using the *Financial Phrase-Bank* of Malo et al. (2014). We also build headline and sentence sentiment indices for each news source b and day t using the same method as the one described in section 3.2.5. The resulting CBERT sentiment indices are shown in Figure 9.

Figure 9: 30-Day moving average of the ClimateBERT sentence sentiment index for the Financial Times, The Daily Telegraph, The Guardian, Los Angeles Times and The New York Times news sources from January 2005 to November 2021.

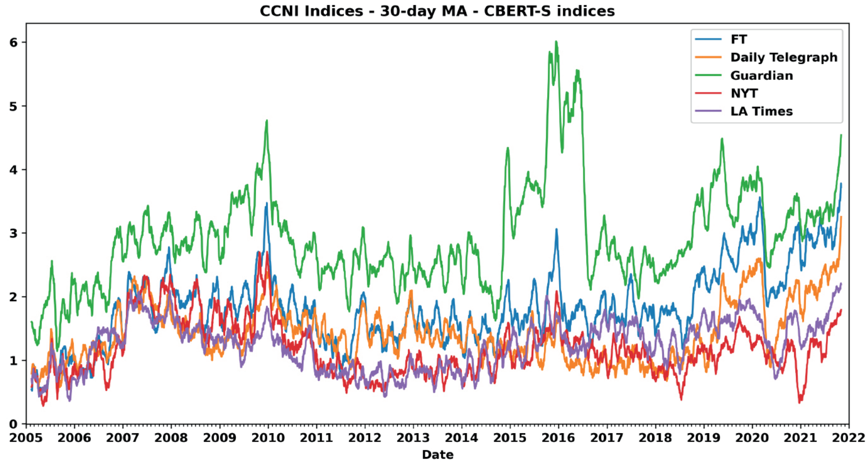

3.3 Climate News Indices for Physical and Transition Risk

We wish to quantify to what extent a climate change article concerns physical risk, transition risk, or both. To begin, we construct a lexicon of physical risk terms, $\mathcal{P}$, and a lexicon of transition risk terms, $\mathcal{T}$, using climate change (CC) glossaries produced by the Intergovernmental Panel on Climate Change (IPCC) and the US Environmental Protection Agency (EPA). This was done manually and validated by another independent researcher. Care was taken to avoid using single words with an ambiguous meaning, e.g. a word like 'warming' is liable to return lots of false positives. Instead we chose to use bigrams such as 'Paris accord'. The resulting physical risk lexicon consists of 182 terms and the transition risk lexicon consists of 133 terms. To provide some transparency on the contents of these lexicons, Figure 10 shows two word clouds for the single words most commonly found in each of the physical and transition risk lexicons.

To calculate each of the physical and transition indices, we once again adopt a TF-IDF approach in which we determine the TF-IDF cosine similarity of each climate change article with the climate change document (CCD) described in section 3.2.2. We differentiate the physical and transition indices by using the physical and transition risk lexicons respectively as the vocabulary $\mathcal{V}$. This is a better approach than calculating the cosine similarity between each news article and each lexicon as a lexicon has only one occurrence of each bigram - the word frequencies do not resemble a typical document on physical or transition risk. As this approach is done at an article level, it captures information about both the

33 - See https://huggingface.co/climatebert/distilroberta-base-climate-f for more details.

number of climate change articles and the degree to which each reflects physical or transition risk.

Specifically, for each news source b we construct a corpus $\mathcal{D}^b$, containing each climate change-related article from 2005 to 2021 plus the climate change document. On each date t, we calculate the TF-IDF vector for each article and for the CCD using either the physical or transition risk lexicon as the vocabulary. We then compute[34] the cosine similarity $CS^b_{i,\mathcal{V}}(t)$ for news source b, article i, using vocabulary $\mathcal{V}$ on date t, and use this to compute a daily climate change news index for physical risk using

$$CNI^b_{\mathcal{V}}(t) = f\left(\sum_{i=1}^{N_b(t)} CS^b_{i,\mathcal{V}}(t)\right)$$

where $f(x)$ was defined and discussed earlier and where $\mathcal{V} = \mathcal{P}$ if we wish to calculate the physical risk index and $\mathcal{V} = \mathcal{T}$ if we wish to calculate the transition risk cosine similarity.

3.4 An Aggregate Climate Change News Index

In addition to the set of CNI for each news source, we also wish to construct an index that has been aggregated across all five news sources. Doing this increases the total number of articles being used in the construction of this index, see Table 8, and might be expected to reduce any statistical noise in the article counts and so enhance any signal that may exist across the individual indices. Table 10 displays the correlation between the five individual news source indices for the Attention methodology over the 3 January 2005 - 3 November 2021 period: the correlation coefficients range between 0.327 (for the Financial Times and the New York Times) and 0.493 (for the Financial Times and the Guardian).

Table 10: Correlation matrix between all news sources indices for the attention methodology over the 3 January 2005 - 3 November 2021 period

	FT	Daily T.	Guardian	NYT	LAT
FT	1.000	0.423	0.493	0.327	0.399
Daily T.	0.423	1.000	0.365	0.294	0.377
Guardian	0.493	0.365	1.000	0.342	0.397
NYT	0.327	0.294	0.342	1.000	0.355
LAT	0.399	0.377	0.397	0.355	1.000

Rather than simply average the individual newspaper indices, we first standardise them so that each index has a unit standard deviation over a three-year period of T dates. This ensures that a newspaper index that experiences a high level of variability in both article number and sentiment score is adjusted to be more comparable to a newspaper index that has a lower variability. Hence, for each news source b, and index $CNI^b(t)$, the standard deviation of the index is calculated as

$$\sigma_b = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\left(CNI^b(t) - \overline{CNI}^b\right)^2}$$

34 - The cosine similarity between the respective TF-IDF vectors $\mathbf{T}^b_{i,\mathcal{V}}(t)$ and $\mathbf{T}_{CCD,\mathcal{V}}$ given by $CS^b_{i,\mathcal{V}}(t) = \frac{\mathbf{T}^b_{i,\mathcal{V}}(t) \cdot \mathbf{T}_{CCD,\mathcal{V}}}{\|\mathbf{T}^b_{i,\mathcal{V}}(t)\| \times \|\mathbf{T}_{CCD,\mathcal{V}}\|}$.

where

$$\overline{CNI}^{b} = \frac{1}{T}\sum_{t=1}^{T} CNI^{b}(t).$$

The aggregate index is given as follows

$$CNI^{Agg}(t) = \frac{\overline{\sigma}}{n_B}\sum_{b=1}^{n_B}\frac{CNI^{b}(t)}{\sigma_b}$$

where $\overline{\sigma} = 1/n_B \sum_b \sigma_b$. The aggregate index at date t has been standardised by a volatility estimated over a three-year rolling window prior to t, so there is no look-ahead bias.

Figure 10: 'Word clouds' showing single words from the terms in the physical risk (top) and transition risk (bottom) lexicons. Word size relates to the number of times a word occurs within each lexicon - the same word may occur several times across the many bigrams in each lexicon.

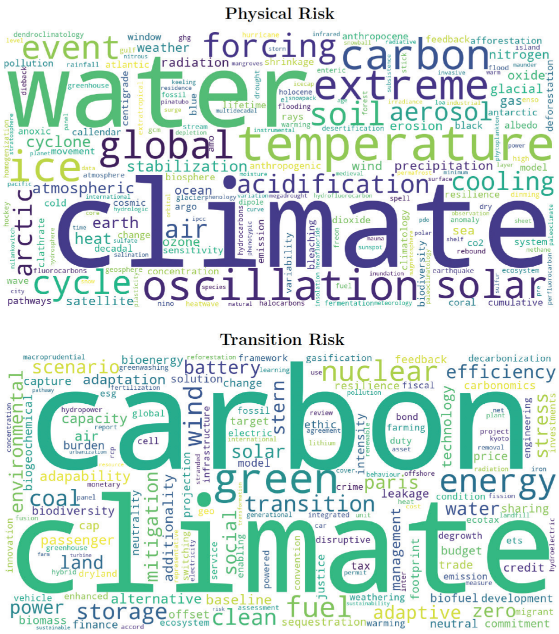

Table 11: Summary description of the CNI Index construction methodologies.

Index	Construction Methodology
Attention	Counts the number of daily climate change articles.
Similarity	Calculates the TF-IDF cosine similarity between each daily newspaper issue and the climate change document.
Concern	Uses LIWC lexicons to measure the excess of positive sentiment versus negative sentiment words weighted by the concern words in each article.
VAD-H	Uses the VADER rules-based sentiment model to determine the sentiment of each article by calculating the sentiment of the headline.
VAD-S	Uses the VADER rules-based sentiment model to determine the sentiment of each article by calculating the average sentiment of each climate change-related sentence.
BERT-H	Uses the fine-tuned BERTBASE language model to determine the sentiment of each article by calculating the sentiment of the headline.
BERT-S	Uses the fine-tuned BERTBASE language model to determine the sentiment of each article by calculating the average sentiment of each climate change-related sentence.
CBERT-H	Uses the fine-tuned domain-specific CBERT language model to determine the sentiment of each article by calculating the sentiment of the headline.
CBERT-S	Uses the fine-tuned domain-specific CBERT language model to determine the sentiment of each article by calculating the average sentiment of each climate change-related sentence.
Physical Similarity	Calculates the TF-IDF cosine similarity between each daily newspaper article and the climate change document using a Physical risk lexicon.
Transition Similarity	Calculates the TF-IDF cosine similarity between each daily newspaper article and the climate change document using a Transition risk lexicon.

4. Climate Change News and Equity Portfolios

Given a set of climate change news indices (CNI) described in Table 11, the next step is to determine whether or not these indices have an impact on equity returns. We examine two different investment universes that we expect to be sensitive to climate change risk. These are:

1. Three liquid US stock portfolios engaged respectively in a Low Carbon Intensity (LCI) strategy, a High Carbon Intensity (HCI) strategy and a low-minus-high Carbon Intensity (LmHCI) portfolio strategy,
2. An industry-specific portfolio of liquid US stocks.

These will allow us to investigate the different ways in which climate change news can impact the pricing of US stock portfolios. For this analysis, we are reliant on good quality carbon emissions data which only begin in July 2012. As a result the period under study runs from July 2012 until November 2021 for the LCI, HCI and LmHCI portfolios, and for the industry portfolios.

4.1 Unexpected Climate News Innovations

The CNI indices that we have calculated may embed some auto-correlation effects[35] and these must be removed if we are to correctly capture the *unexpected changes* in the climate news index. The unexpected climate news innovations index, $UCNI_t$, are defined by

$$UCNI_t = CNI_t - \mathbb{E}\left[CNI_t | I_{t-1}\right]$$

where I_{t-1} is information to time $t - 1$. Each value of the $UCNI_t$ is calculated as the residual ϵ_t of an AR(1) process calibrated to the CNI_t over the previous three years. The Aggregate $UCNI_t$ for the different language models are shown in Figure 11.

We report in Table 12 the pairwise correlation between the $UCNI_t$ climate news aggregate innovations index time series over the period from 3 January 2005 to 3 November 2021. We see that, apart from the similarity index, the indices are all highly correlated with all pairwise correlations between the unexpected climate news indices greater than 81%, and with many over 95%. We can explain this high correlation by the fact that for all but the similarity index, the daily changes in all of these indices are driven by changes in both the number of climate change articles published and the change in the corresponding sentiment score for these articles. The similarity index is different because rather than sum over articles, we combine the articles for each day into a single document and consider its cosine similarity with the CCD.

Table 12: Correlation matrix of UCNI for the aggregated news source over the period from 3 January 2008 to 3 November 2021.

	Attention	Similarity	Concern	VAD-H	VAD-S	BERT-H	BERT-S	CBERT-H	CBERT-S
Attention	1.000	0.681	0.882	0.994	0.994	0.978	0.975	0.953	0.975
Similarity	0.681	1.000	0.592	0.676	0.677	0.661	0.652	0.654	0.673
Concern	0.882	0.592	1.000	0.884	0.882	0.878	0.877	0.819	0.847
VAD-H	0.994	0.676	0.884	1.000	0.994	0.978	0.977	0.938	0.971
VAD-S	0.994	0.677	0.882	0.994	1.000	0.976	0.982	0.946	0.975
BERT-H	0.978	0.661	0.878	0.978	0.976	1.000	0.963	0.907	0.948
BERT-S	0.975	0.652	0.877	0.977	0.982	0.963	1.000	0.917	0.935
CBERT-H	0.953	0.654	0.819	0.938	0.946	0.907	0.917	1.000	0.945
CBERT-S	0.975	0.673	0.847	0.971	0.975	0.948	0.935	0.945	1.000

35 - Autocorrelation refers to the tendency of data points in a time series to be correlated with their past values. In the context of news, this means that the way news is reported today may be influenced by how it was reported in the past. For example, if there is a news story about a heatwave, it is likely that there will be more stories about heatwaves in the following days or weeks. This is because news outlets may continue to report on the same topic as long as it remains relevant and interesting to their audience.

Figure 11: Climate change news indices (CNI) for the period July 2012 to November 2021 and Unexpected Climate News Innovations (UCNI) for the Aggregate news source over the period July 2012 to November 2021.

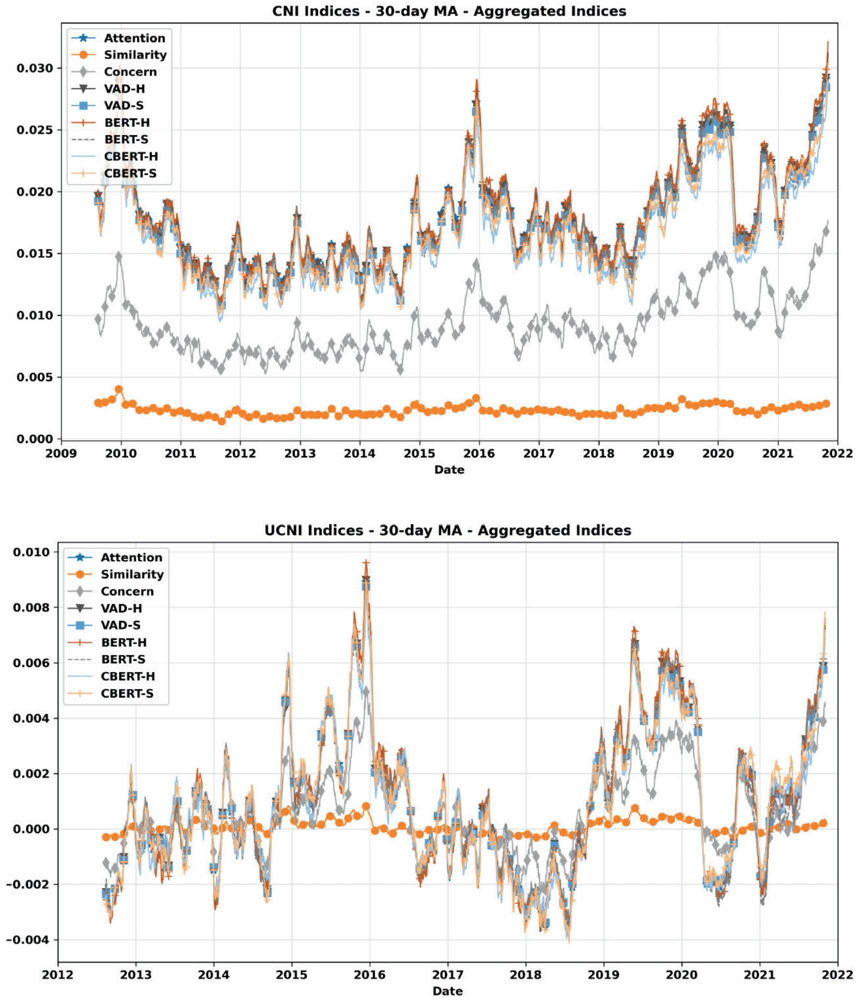

4.2 Is the UCNI Factor Spanned by the Fama-French Factors?

We wish to test whether the UCNI indices can be spanned by the five Fama-French (FF) factors from Fama and French (2015) plus the (WML) momentum factor from Carhart (1997). These are defined as:

- *MKT*: The excess market return,
- *SMB*: The small minus big (size) factor,
- *HML*: The high minus low (value) factor,
- *RMW*: The robust minus weak (high profitability) factor,
- *CMA*: The conservative minus aggressive (low investment) factor, and
- *WML*: The winners minus losers (momentum) factor.

Table 13 shows the linear regression of the UCNI factor for the Aggregate news indices versus the FF factors, plus the Winners Minus Losers (WML) momentum factor from Carhart (1997), for the period from July 2012 to November 2021. Apart from the intercept coefficients, none of the coefficients are statistically significant at the 5% level. This result and the very low R-squared values suggests that the UCNI factor is not spanned by the FF factors for any of the indices, implying that it could perhaps be a new pricing factor.

4.3 The US500 Capitalization-Weighted (US500CW) Portfolio

We identify whether or not the UCNI are a pricing factor that explain differences in expected returns across different stock portfolios. We begin with a simple regression of the UCNI index against the US500CW portfolio return over the period from July 2012 to November 2021. The regression equation is given by:

$$r_t^{US500CW} - r_{ft} = \alpha + \beta_{UCNI} \cdot UCNI_t + \varepsilon_t$$

where $UCNI_t$ is the unexpected climate news innovations index, $r_t^{US500CW}$ is the daily total return of US500CW portfolio, and r_{ft} is the risk-free rate calculated from the three-month US Treasury-Bill rate taken from the St Louis Federal Reserve website[36].

*Table 13: Linear regression results of the UCNI factor against the Fama-French factors plus the WML momentum factor for the aggregate news source. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	Attention	Similarity	Concern	VAD-H	VAD-S	BERT-H	BERT-S	CBERT-H	CBERT-S
Intercept	0.0009***	0.0001***	0.0006***	0.001***	0.001***	0.001***	0.001***	0.0009***	0.001***
	(7.46)	(5.37)	(8.13)	(7.95)	(7.92)	(7.73)	(7.59)	(7.0)	(7.76)
Mkt-RF	0.008	0.0002	-0.0134	0.0076	0.0056	0.0016	0.0052	0.0103	0.0144
	(0.33)	(0.08)	(-0.86)	(0.3)	(0.23)	(0.06)	(0.21)	(0.41)	(0.58)
SMB	-0.0184	-0.002	-0.0085	-0.019	-0.0158	-0.0195	-0.0179	-0.0144	-0.0204
	(-1.46)	(-1.42)	(-1.07)	(-1.5)	(-1.28)	(-1.46)	(-1.41)	(-1.13)	(-1.62)
HML	0.0079	0.0021	-0.0058	0.0049	0.0057	0.0027	0.0074	0.0062	0.0063
	(0.48)	(1.09)	(-0.55)	(0.29)	(0.35)	(0.15)	(0.44)	(0.37)	(0.38)
CMA	0.0342	0.002	0.0067	0.0241	0.023	0.0241	0.0328	0.0198	0.0163
	(1.02)	(0.52)	(0.31)	(0.69)	(0.67)	(0.66)	(0.94)	(0.57)	(0.47)
RMW	0.0345	0.0014	0.0382	0.0211	0.027	0.0192	0.0187	0.0473	0.0298
	(0.77)	(0.27)	(1.34)	(0.47)	(0.61)	(0.4)	(0.41)	(1.04)	(0.66)
Mom	0.0144	0.0039	-0.0037	0.0102	0.0105	0.0123	0.0083	0.0135	0.0124
	(0.64)	(1.49)	(-0.26)	(0.45)	(0.48)	(0.52)	(0.37)	(0.59)	(0.55)
R2	0.0027	0.0021	0.0022	0.0019	0.0017	0.0016	0.0018	0.0024	0.0025

The results in Table 14 are for the Aggregate indices. They suggest that the US500CW portfolio has only a low exposure to any of the UCNI factors. This may be because few individual stocks have any significant beta exposure to these factors. It may also be because the climate change sensitivity of HCI and LCI firms within this index have opposite signed betas and so effectively cancel each other out, leaving no statistically significant signal at the US500CW level.

4.4 LCI, HCI, and Low-minus-High Carbon Intensity Portfolios

The next step is to examine whether differences in exposure to the climate news index help us to explain expected returns of LCI, HCI and Low-minus-High Carbon Intensity (LmHCI) portfolios. Specifically, we want to determine whether the UCNI factor is a new pricing factor, after accounting for the well-documented Fama-French and WML pricing factors. The linear regression we wish to fit is the following:

$$\begin{aligned}\tilde{r}_t \quad &= \alpha + \beta_{UCNI} \cdot UCNI_t + \beta_{MKT} \cdot MKT_t + \beta_{HML} \cdot HML_t + \beta_{SMB} \cdot SMB_t \\ &+ \beta_{CMA} \cdot CMA_t + \beta_{RMW} \cdot RMW_t + \beta_{WML} \cdot WML_t + \varepsilon_t.\end{aligned}$$

36 - See https://fred.stlouisfed.org/series/TB3MS.

The left-hand side is the daily excess return of the portfolio under study. On the right hand side we have the $UCNI_t$ factor, the five Fama-French factors, and the WML momentum factor. The term ϵ_t is an identically and independently distributed error term with zero mean. We will determine if the factor loading on the $UCNI_t$ index is statistically significant. This will test the model of Pástor et al. (2021) which predicts a strictly positive β_{UCNI} coefficient for the LCI and LmHCI portfolios and a strictly negative β_{UCNI} coefficient for the HCI portfolio. We can also study the statistical significance of β_{UCNI} coefficients and repeat this exercise for the other investment universes described at the beginning of Section 4. Finally, we can analyse the sign and magnitude of the β_{UCNI} coefficient of industry portfolios.

To identify which stocks are LCI and which are HCI, we use a selection method that is based on the carbon intensities of individual companies. These have been determined using combined Scope 1 and Scope 2 emissions from Factset's ISS ESG carbon emissions data. Note that the firm-level emissions and revenues used to determine the carbon intensity are calculated at an annual frequency. The initial investment universe starts with the 500 US stocks with the largest market capitalisation (according to CRSP), and these are then sorted by their carbon intensity. The LCI portfolio and HCI portfolio are equallyweighted and consist of the 30% of stocks with respectively the lowest and highest carbon intensity[37]. A snapshot of the industry sector weightings of the LCI and HCI portfolios are shown in Table 15 for 18 March 2022. As we might expect, the HCI portfolio has a significant Energy and Basic Materials component. Returns are daily and the portfolios are re-balanced quarterly, on the third Friday of March, June, September, and December.

*Table 14: Linear regression of the excess annualised return of the US500CW portfolio against the unexpected climate news index computed from the aggregate news source over the period July 2012 to November 2021. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	Attention	Similarity	Concern	VAD-H	VAD-S	BERT-H	BERT-S	CBERT-H	CBERT-S
Intercept	0.0007***	0.0007***	0.0007***	0.0007***	0.0007***	0.0007***	0.0007***	0.0007***	0.0007***
	(3.22)	(3.09)	(3.18)	(3.18)	(3.15)	(3.17)	(3.17)	(3.12)	(3.2)
UCNI	-0.0575*	-0.4229	-0.0878	-0.0568	-0.0507	-0.0526	-0.0546	-0.0454	-0.0605*
	(-1.66)	(-1.38)	(-1.57)	(-1.62)	(-1.42)	(-1.58)	(-1.56)	(-1.3)	(-1.72)
R2	0.0011	0.0008	0.001	0.0011	0.0009	0.0011	0.001	0.0007	0.0013

Table 15: The industry sector breakdown of the LCI and HCI portfolios for 18 March 2022.

Industry Sector	LCI (%)	HCI (%)
Energy	-	12.08
Basic Materials	-	13.42
Industrials	7.38	10.74
Cyclical Consumer	9.40	9.40
Non-Cyclical Consumer	2.01	11.41
Financials	36.24	14.09
Healthcare	15.44	5.37
Technology	29.53	6.71
Telecoms	-	0.67
Utilities	-	16.11

37 - We also tested an alternative approach for the LmHCI portfolio construction developed by Vaucher et al. (2023). The long ("LCI") leg of the LmHCI portfolio is built as an equally-weighed (EW) portfolio of the 50% of the stocks with the lowest carbon intensity selected within each of the six sectors with highest exposure to stranding risk in the event of a disorderly low-carbon transition (according to the classification of Battiston et al. (2017)). Conversely, the short ("HCI") leg is built as an EW portfolio of the 50% of the stocks with the highest carbon intensity selected within each of these sectors. The findings generated via this alternate methodology concur with those obtained through the initial methodology.

We begin the analysis by applying the linear regression written in Equation(2) to the LCI, HCI and LmHCI portfolios described in Section 4.6. Table 16 reports the β_{UCNI} coefficient for the LCI portfolio. These results have low significance and hence the impact of climate news on LCI stock returns is low. Table 17 shows the equivalent UCNI beta for the HCI portfolio for the different news sources, including the aggregate news source. We note that the betas are all negative, implying that a day with high unexpected negative concerns is, on average, always associated with a negative impact on the returns of HCI stocks. Significance for the individual news source betas is reasonably high with some at 10%, 5% and one at 1% significance (CBERT-H). However, for the aggregate index, the significance improves substantially and five of the nine indices have significance at 1%, three having significance at 5% amongst the other four. Table 18 shows the UCNI beta coefficient for the LmHCI portfolio for the linear regression performed over the period from July 2012 to November 2021. The first observation is that the betas are mostly positive - unexpected negative climate sentiment is associated with a positive return due to a fall in value of the short position in the HCI stocks. For the VADER, BERT, and Climate BERT models, the (H) approach looking at the titles of articles achieves the same level of significance for the HCI portfolios as the (S) approach considering the contents of the articles (S). Nevertheless, for the LmHCI portfolios, the (S) approach yields more significant results (1% for BERT-S and CBERT-S versus 5% for BERT-H and CBERT-H).

*Table 16: Value of the corresponding UCNI beta and significance for the LCI portfolio by news source and index construction methodology. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	FT	Daily T.	Guardian	NYT	LAT	Aggregate
Attention	0.0042	-0.001	-0.004	-0.0037	-0.0002	-0.0006
Similarity	-0.0048	0.0379	0.0096	-0.0098	0.0381	0.0514
Concern	0.0027	0.0047	-0.0021	-0.0068	0.0039	0.0058
VAD-H	0.005	-0.0012	-0.0037	-0.0034	-0.0003	-0.0003
VAD-S	0.0063	-0.0017	-0.0042	-0.0046	0.0004	-0.0005
BERT-H	0.0017	-0.0006	-0.0054	-0.0039	-0.0015	-0.0023
BERT-S	0.0054	-0.0017	-0.003	-0.0027	0.0016	0.0009
CBERT-H	0.0016	-0.0044	-0.0043	-0.0095*	0.0012	-0.0051
CBERT-S	0.0089*	-0.0012	-0.0038	-0.0036	-0.0007	0.0007

*Table 17: Value of the corresponding UCNI beta and significance for the HCI portfolio by news source and index construction methodology. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	FT	Daily T.	Guardian	NYT	LAT	Aggregate
Attention	-0.0135	-0.0165**	-0.0157**	-0.0102	-0.0139*	-0.0295***
Similarity	-0.0296	-0.0807	0.059	-0.0522	-0.1022**	-0.1735*
Concern	-0.0202*	-0.0087	-0.0111	-0.0231*	-0.0186	-0.0435**
VAD-H	-0.0127	-0.0142*	-0.0158**	-0.0099	-0.0144*	-0.0296***
VAD-S	-0.0142	-0.0142*	-0.0156**	-0.0084	-0.014*	-0.0296***
BERT-H	-0.0151*	-0.0112*	-0.0117*	-0.0099	-0.0097	-0.0257**
BERT-S	-0.0154*	-0.0102	-0.014**	-0.0099	-0.0087	-0.0272**
CBERT-H	-0.0116	-0.0179***	-0.014**	-0.0026	-0.016**	-0.0311***
CBERT-S	-0.011	-0.0185**	-0.016**	-0.0145	-0.0149*	-0.0334***

*Table 18: Value of the corresponding UCNI beta and significance for the Low-minus-High Carbon Intensity portfolio by news source and index construction methodology. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	FT	Daily T.	Guardian	NYT	LAT	Aggregate
Attention	0.0177*	0.0156*	0.0117	0.0065	0.0136	0.0289**
Similarity	0.0248	0.1187*	-0.0494	0.0424	0.1403**	0.2249*
Concern	0.0229	0.0134	0.009	0.0162	0.0225	0.0493**
VAD-H	0.0177	0.013	0.0121	0.0065	0.0141	0.0293**
VAD-S	0.0205*	0.0125	0.0114	0.0038	0.0144	0.0291**
BERT-H	0.0168	0.0106	0.0064	0.0059	0.0082	0.0235*
BERT-S	0.0208*	0.0085	0.011	0.0072	0.0102	0.0282**
CBERT-H	0.0132	0.0136	0.0098	-0.0069	0.0172*	0.026*
CBERT-S	0.0199*	0.0174**	0.0122	0.0108	0.0141	0.0341**

We completed our analysis by conducting a number of robustness tests. For the first robustness check, we consider two different five-year time periods, the first from December 2012 to December 2017 and the second from December 2015 to December 2020. Due to space limits, we do not include all of the regression tables in this report and simply summarise the results found. For the linear regression performed over the period from December 2012 to December 2017 where the dependent variable is the HCI excess return, we observe that seven out of nine aggregate indices display a statistically significant climate beta at the 1% level. However for the period from December 2015 to December 2020, none of the aggregate indices display a climate beta significant at the 1% level and 3 of them (Concern, CBERT-H and CBERT-S) have a significant climate beta at the 5% level. We find a similar difference in significance when the dependent variable is the LmHCI excess return: over the period from December 2012 to December 2017 where the dependent variable is the LmHCI return, we observe that seven out of nine aggregate indices display a statistically significant climate beta at the 1% level but over the period from December 2015 to December 2020 none of them displays a significant climate beta at the 5% level. This suggests that UCNI changes had a larger market impact in the earlier 2012-2017 period than in the more recent 2015-2020 period. In the second robustness test we repeat the regressions by constructing five new aggregate indices, each with one news source removed and hence with only four news sources, over the period from July 2012 to November 2021. The results obtained for the climate beta of the HCI and LmHCI portfolio are consistent with those obtained when the aggregate news indices are built with five news sources. In a third and final robustness test we changed the percentile cutoff for designating LCI and HCI assets from 30% to first 10% and then 50%. We find that our results for the HCI and LmHCI portfolios remain overall as statistically significant for percentile cutoffs of 10% and 50%. We also observed that for the HCI and LmHCI portfolios, the amplitude of the UCNI beta coefficients increases as the threshold decreases. This is consistent with the expectation that a higher percentile cutoff would result in a greater tilt towards HCI stocks in the HCI portfolio and in the short leg of the LmHCI portfolio.

4.5 Industry Portfolios

We now study the performance of equally-weighted industry-specific portfolios that are taken from from Kenneth French's website[38]. In these portfolios each NYSE, AMEX, and NASDAQ stock is assigned to one out of twelve industry portfolio at the end of June of year *t*, based on its four-digit SIC code on that date. We perform the following regression over the July 2012 to November 2021 period

$$\begin{aligned} r_t - r_{ft} &= \alpha + \beta_{UCNI} \cdot UCNI_t + \beta_{MKT} \cdot MKT_t + \beta_{HML} \cdot HML_t + \beta_{SMB} \cdot SMB_t \\ &+ \beta_{CMA} \cdot CMA_t + \beta_{RMW} \cdot RMW_t + \beta_{WML} \cdot WML_t + \varepsilon_t \end{aligned}$$

where $UCNI_t$ is calculated from the aggregate index with the C-BERT-S methodology. The results are shown in Table 19. We find that the beta of the UCNI is negative and significant at 5% for the Consumer Non-Durables (Food, Tobacco, Textiles, Apparel, Leather, Toys), the Manufacturing (Machinery, Trucks, Planes, Off Furn, Paper, Com Printing), the Business Equipment (Computers, Software, and Electronic Equipment) and the Other (Mines, Constr, BldMt, Trans, Hotels, Bus Serv, Entertainment) industries. The Energy (Oil, Gas, and Coal Extraction and Products) industry factor displays the most negative climate beta (–0.0783) of all the 12 industries, but does not have significance at 10%.

4.6 Portfolio Performance Conditional on the UCNI

We examine whether or not unexpected climate news innovations can help predict the *conditional* performance of HCI versus LCI stocks. To do this, we study the average daily performance of the LmHCI portfolio over the period from July 2012 to November 2021, conditional on the level of the daily $UCNI_t$ index. We define three different regimes which we call low, medium and high. The low regime corresponds to days when the UCNI time series is below the first tercile, the medium regime to days when the UCNI time series is between the first and the second terciles, and the high regime to daily periods when the UCNI time series is above the second tercile.

*Table 19: Linear regression of the industry portfolios returns versus the five Fama-French factors plus the WML momentum factor and the UCNI factor built from the CBERT-S aggregated news index. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	NoDur	Durbl	Manuf	Enrgy	Chems	BusEq	Telcm	Utils	Shops	Hlth	Money	Other
Intercept	0.0002***	0.0003**	0.0002**	-0.0004	0.0001	0.0003***	0.0001	0.0001	0.0	0.0002	0.0003***	0.0002**
	(2.6)	(2.16)	(2.15)	(-1.33)	(0.91)	(4.08)	(0.88)	(0.62)	(0.4)	(1.64)	(4.73)	(2.34)
UCNI	-0.0308**	-0.0194	-0.0287**	-0.0783	-0.0369**	-0.0263**	0.0032	-0.0347	-0.0275*	0.0181	0.0154	-0.0316**
	(-2.19)	(-0.95)	(-2.42)	(-1.4)	(-1.98)	(-2.04)	(0.14)	(-1.34)	(-1.81)	(0.86)	(1.43)	(-2.42)
Mkt-RF	0.7996***	0.9668***	0.9866***	1.2817***	0.9756***	0.9561***	0.9047***	0.7506***	0.9292***	0.8917***	0.7947***	0.9381***
	(96.37)	(80.33)	(140.8)	(38.78)	(88.54)	(125.87)	(68.82)	(49.08)	(103.6)	(71.53)	(125.52)	(121.83)
SMB	0.5844***	0.945***	0.8241***	0.6034***	0.6422***	0.6581***	0.6858***	0.0916***	0.8886***	0.9249***	0.5709***	0.7607***
	(39.45)	(43.97)	(65.87)	(10.22)	(32.64)	(48.52)	(29.22)	(3.36)	(55.49)	(41.55)	(50.5)	(55.33)
HML	0.1061***	0.118***	0.2117***	0.5378***	0.1358***	-0.2001***	0.1046***	0.1186***	0.1297***	-0.5242***	0.7021***	0.1715***
	(6.49)	(4.98)	(15.33)	(8.26)	(6.26)	(-13.37)	(4.04)	(3.94)	(7.34)	(-21.34)	(56.29)	(11.3)
CMA	0.1417***	0.1109**	0.1465***	0.4691***	0.3008***	-0.2031***	0.1209**	0.3882***	0.0922***	0.0035	-0.4186***	-0.0973***
	(4.78)	(2.58)	(5.84)	(3.97)	(7.63)	(-7.48)	(2.57)	(7.1)	(2.87)	(0.08)	(-18.49)	(-3.53)
RMW	0.282***	0.1376***	0.0408**	-1.4194***	-0.0988***	-0.3078***	-0.0624*	0.197***	0.3496***	-1.0274***	0.0159	0.039*
	(12.39)	(4.17)	(2.12)	(-15.66)	(-3.27)	(-14.78)	(-1.73)	(4.7)	(14.21)	(-30.06)	(0.92)	(1.85)
Mom	0.0053	-0.0024	-0.0886***	-0.8884***	-0.1138***	-0.0171*	-0.1007***	-0.0288	-0.0249**	-0.1768***	0.0441***	-0.0637***
	(0.49)	(-0.15)	(-9.63)	(-20.48)	(-7.87)	(-1.72)	(-5.84)	(-1.43)	(-2.11)	(-10.8)	(5.31)	(-6.31)
R2	0.86	0.832	0.9379	0.6258	0.8424	0.9145	0.7709	0.5557	0.891	0.8328	0.9373	0.9165

38 - See https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html

Table 20 reports the conditional performance of the LCI, HCI and LmHCI portfolios with respect to UCNI calculated from the Aggregate news source. The average annualised return of the LmHCI portfolio over the high regime is greater than that over the medium and low regime for all the methodologies under study. We also find that the average annualised return of the LmHCI portfolio in the medium regime is greater than that over the low regime for all the methodologies under study. For example, choosing the Climate BERT sentences index it is 13.6% in the high regime versus 4.0% in the low regime. We see that when the UCNI is in the high tercile, the return of the HCI portfolio is negative for all of the indices except the similarity index.

4.7 Results for Physical and Transition Risk

The physical and transition risk cosine similarity indices are shown in Figure 12. We find that the transition risk index is generally higher than the physical risk index, implying that this topic is covered more by the major newspapers. The correlation between the physical and transition risk indices is 63%.

Figure 12: The physical and transition index calculated according to the methodology described in Section 3.3 and in Section 3.4.

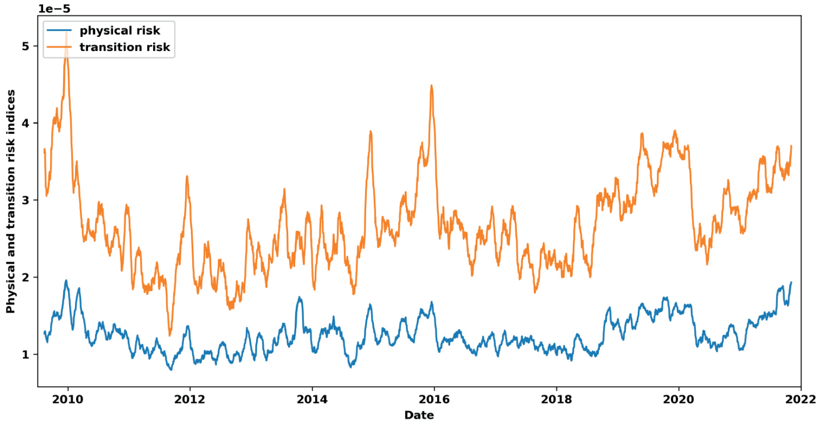

Table 20: UCNI-Conditional annualised performance of LCI, HCI and LmHCI portfolios for the aggregate news source and for headlines (H) and sentences (S).

Attention	LCI	HCI	LmHCI
low	21.8%	19.4%	2.4%
medium	33.3%	25.0%	8.2%
high	-0.5%	-10.4%	9.9%
Similarity	LCI	HCI	LmHCI
low	31.2%	24.6%	6.6%
medium	3.4%	3.5%	-0.2%
high	20.8%	5.4%	15.4%
Concern	LCI	HCI	LmHCI
low	24.7%	21.0%	3.7%
medium	25.2%	21.8%	3.4%
high	5.5%	-9.3%	14.8%
VAD-H	LCI	HCI	LmHCI
low	22.7%	20.3%	2.4%
medium	30.2%	23.3%	6.9%
high	2.5%	-10.1%	12.6%

VAD-S	LCI	HCI	LmHCI
low	23.1%	19.6%	3.5%
medium	30.0%	23.2%	6.8%
high	2.3%	-9.3%	11.6%
BERT-H	LCI	HCI	LmHCI
low	28.6%	24.6%	4.0%
medium	23.0%	15.0%	8.1%
high	3.8%	-6.1%	9.9%
BERT-S	LCI	HCI	LmHCI
low	20.4%	19.5%	0.9%
medium	34.0%	22.2%	11.8%
high	1.0%	-8.2%	9.2%
CBERT-H	LCI	HCI	LmHCI
low	21.1%	18.1%	3.0%
medium	28.0%	19.6%	8.4%
high	6.3%	-4.2%	10.5%
CBERT-S	LCI	HCI	LmHCI
low	27.6%	23.2%	4.4%
medium	27.1%	22.5%	4.5%
high	0.7%	-12.3%	13.0%

We regressed both the physical and transition Aggregate UCNI indices against the LCI, HCI and LmHCI portfolios. The results are shown in Tables 21, 22 and 23. We see that the ability of the Aggregate physical risk index to explain returns for the LCI portfolio is not significant at 10% or less. However the ability of the Aggregate physical risk index to explain returns for the HCI portfolio is significant to 1% for the Aggregate index, falling to 5% for the Low-minus-High Carbon Intensity portfolio.

*Table 21: Value of the corresponding UCNI beta and significance for the LCI portfolio by news source and index construction methodology. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	FT	Daily T.	Guardian	NYT	LAT	Aggregate
Physical	0.0109	0.0039	-0.0092	-0.0078	-0.0008	0.003
Transition	0.0083	0.0001	-0.0068	-0.0036	0.0028	0.003

*Table 22: Value of the corresponding UCNI beta and significance for the HCI portfolio by news source and index construction methodology. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	FT	Daily T.	Guardian	NYT	LAT	Aggregate
Physical	-0.0315*	-0.03*	-0.0376***	-0.0184	-0.0399**	-0.07***
Transition	-0.0212	-0.0277**	-0.018	-0.0152	-0.0252*	-0.0485**

*Table 23: Value of the corresponding UCNI beta and significance for the LmHCI portfolio by news source and index construction methodology. We use *, **, *** to denote statistical confidence at 10%, 5% and 1% respectively.*

	FT	Daily T.	Guardian	NYT	LAT	Aggregate
Physical	0.0424*	0.0339*	0.0284*	0.0107	0.0391*	0.073**
Transition	0.0295	0.0278*	0.0112	0.0116	0.028*	0.0515**

4.8 The Choice of Function $f(x)$

Throughout this analysis, we also investigate the choice of function $f(x)$, testing both $f(x) = x$ and $f(x) = \sqrt{x}$. We find that the second of these two produced more significant results, suggesting that the marginal impact of additional climate change news articles on financial markets is a declining function. Hence we prefer a concave function.

5. Conclusion

Our evidence suggests that the climate news factor built on an aggregate index of highquality newspapers has an explanatory power over the High Carbon Intensity (HCI) portfolio returns, and hence over Low-minus-High Carbon Intensity portfolio returns. The improved significance of an aggregate index over individual indices suggests that individual newspapers may publish climate change articles even when there is no climate change event to report. By aggregating these newspapers, we reduce the importance of these idiosyncratic articles while retaining the importance of the climate articles which they all publish - those which report on actual unexpected climate news events.

Our findings agree with the general observation by Ardia et al. (2022) who find that Low Carbon Intensity (LCI) firms outperform HCI firms when there are unexpected increases in climate change concern. Unlike Ardia et al. (2022), we find that this is not because LCI stocks rise in value, but because HCI stocks fall in value. This result is consistent with the results of Bua et al. (2020). This implies that unexpected climate news is generally bad for HCI assets, perhaps because HCI firms may have more to lose from climate change news than LCI assets have to gain.

Furthermore, we find that the average return of LmHCI portfolios over the period from July 2012 to November 2021 consistently rises with the UCNI regime (low, medium, high) for all index types. This provides further evidence supporting the hypothesis that climate change concern plays a significant role in the long-term performance of Low-minus-High Carbon Intensity portfolios.

Out of all of the language models used, the most advanced domain-specific CBERT model does not materially outperform the simpler Attention-based model. This suggests that it is the number of articles, rather than their content, that drives climate risk awareness. It may also suggest that the ability of these state-of-the-art language models to extract sentiment from high quality newspaper articles is limited. This may be due to the complexity of the language found in these newspapers articles. It may also be due to the desire of serious newspapers to be even-handed and moderate in tone.

There are a number of possible extensions of our paper. First, we may wish to add more individual news sources to the aggregate index to see the impact on the significance of the aggregate index. Second, it would be of interest to explore aspect-based sentiment approaches such as Peng et al. (2020). Using such an approach we can ensure that the target of the expressed sentiment is indeed a climate change related matter. It may then be possible to distinguish between LCI and HCI targets. Third, it would be of considerable interest to determine whether the out-of-sample performance of the aggregate UCNI is sufficient to enable us to use it for portfolio hedging as initially proposed by Engle et al. (2020).

6. Appendix

6.1 Climate News Sources

This table lists the source texts for the climate change document.

Table 24: Source documents used to construct the Climate Change (CC) document used by the similarity index. We show the document title, number of pages and publication year.

Document Title	Source	Num Pages	Year
Emissions Scenarios	IPCC	608	2000
CC Scientific Basis	IPCC	893	2001
CC Synthesis Report	IPCC	409	2001
CC Impacts, Adaptation ...	IPCC	1024	2001
CC Mitigation	IPCC	754	2001
Safeguarding the Ozone layer ..	IPCC	485	2002
Carbon Dioxide Capture and Storage	IPCC	443	2005
CC impacts, adapation and vulnerability	IPCC	987	2007
CC Mitigation	IPCC	863	2007
CC Synthesis Report	IPCC	112	2007
CC Impacts in Europe EU	JRC	132	2009
CC Indicators in the US	EPA	84	2012
Renewable Energy Sources ...	IPCC	1088	2012
Managing Risks of Extreme Events ..	IPCC	594	2012
IPCC WG3 AR5 full document	IPCC	1454	2014
CC Indicators	EPA	112	2014
CC Synthesis Report	IPCC	167	2014
American Climate Prospectus	RBP	206	2014
CC Indicators in the US	EPA	96	2016
The Effects of Weather Shocks ...	IMF	40	2018
Environment and CC Mainstreaming..	EU	39	2018
European State of the Climate	EU	20	2020
IPCC AR6 WGI Technical Summary	IPCC	159	2021
CC and Social Vulnerability in the US	EPA	101	2021
CC The Physical Science Basis	IPCC	42	2021
European Firms and CC	EIB	52	2021

6.2 The TF-IDF Measure

The TF-IDF metric is a powerful document representation approach as it over-weights the important words in a document and under-weights the less important words in a document. Important words are defined as those words in a document that do not occur frequently across the corpus of documents in which the document is found. They are therefore more significant for that document. Less important words in a document are defined as those that occur in the document but also occur in lots of other documents in the corpus. They have no special significance for that document.

The TF-IDF measure is based on two quantities - the *term frequency* and the *inverse document frequency*. The term frequency TF(t, d) is defined as the number of times a term (word) t occurs in document d in a corpus of documents $\mathcal{D}$. The inverse document frequency IDF(t,$\mathcal{D}$) is the logarithm of the number of documents in the corpus N_D divided by the number of documents in the corpus, $n(t,\mathcal{D})$, that contain the term t. Multiplying these two quantities gives a 'TF-IDF' for a specific word in a specific document. We use the standard calculation adopted by the Python machine learning library Scikit-Learn[39].

$$\mathrm{IDF}(t,\mathcal{D}) = \ln\left(\frac{1+N_D}{1+n(t,\mathcal{D})}+1\right).$$

It avoids infinity at $n(t,\mathcal{D}) = 0$ by adding 1 to the denominator and the singularity of the log function if $N_D = 0$ by adding another 1. The TF-IDF is therefore given by

$$\mathrm{TF}-\mathrm{IDF}(t,d,\mathcal{D}) = \mathrm{TF}(t,d)\times\mathrm{IDF}(t,\mathcal{D}).$$

6.3 The BERT Transformer Model

BERT, the Bidirectional Encoder Representations from Transformers model in Devlin et al. (2018), is a supervised learning model which can be trained to perform specific tasks. It consists of a feedforward multi-layer neural network architecture that has been tailored to the task of textual analysis. The inputs to the BERT model consist of input nodes that receive a set of word tokens or embeddings where an embedding is a high-dimensional (typically 768 dimensions) vector representation of a specific word. These input words may represent one or even several sentences - special key words are used to indicate where a new sentence begins and ends. The output layer of BERT depends on the task being learned.

What makes BERT special is its internal Transformer architecture. The Transformer was introduced in Vaswani et al. (2017) and is designed to make it possible to learn about word meaning via its context. This was a significant conceptual leap forward in language modeling as it deviated from earlier approaches based on the use of Long Short Term Memory (LSTM) models which see language as an ordered sequence of words with only a knowledge of previous words. The architecture of the Transformer, which involves such components as 'self-attention heads' is beyond the scope of this summary. The simplest description is that it repeatedly calculates a set of dot products between all combinations of processed word embedding vectors in the input sentences. Through this mechanism, the Transformer can construct a numerical representation all of the words in a sentence (or set of sentences) at the same time. The order of the words is also provided, but only via a positional encoding vector that is blended with the actual word embeddings.

The BERT model is initially trained to learn two specific tasks. The process of training involves presenting each of the training examples of the task, and then updating the value of the weights and biases so that the final layer output more closely matches the target label of the task. This is done using some version of the back-propagation algorithm. In the first task, the input layer receives two sentences taken from a large corpus of documents, specifically English language Wikipedia (with 5,500M words) and the BooksCorpus (with 800M words). A word in each sentence is covered ('masked') at random and BERT is then trained to guess the covered words correctly. In the second task, BERT is presented with two sentences taken from the same corpus, which may or may not be successive. It is then

39 - See https://scikit-learn.org/stable/.

trained to decide whether the first sentence precedes the second one. In this way, BERT learns about language and context both within sentences and at a longer distance between sentences.

BERT models comes in various 'sizes', with the larger models having more trainable parameters. They also require greater computational and memory resources. In this paper we use the $BERT_{BASE}$ model which has 110 million trainable parameters, 12 encoder layers, 12 bi-directional self-attention heads and a word embedding dimension of 768.

References

- Apel, M., A. Betzer, and B. Scherer (2021). Real-Time Transition Risk. Working Paper.
- Ardia, D., K. Bluteau, K. Boudt, and K. Inghelbrecht (2022). Climate Change Concerns and the Performance of Green Versus Brown Stocks. Working Paper Available at SSRN: https://ssrn.com/abstract=3717722.
- Battiston, S., A. Mandel, I. Monasterolo, F. Schütze, and G. Visentin (2017). A climate stress-test of the financial system. *Nature Climate Change*.
- Blei, D., A. Ng, and M. Jordan (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3, 993–1022.
- Bollen, L., H. Mao, and X. Zeng (2011). Twitter mood predicts the stock market. *Journal of Computational Science* 2 (1), 1–8.
- Bua, G., D. Kapp, F. Ramella, and L. Rognone (2020). Transition Versus Physical Climate Risk Pricing In Euro Area Financial Markets: A Text-Based Approach. ECB Working Paper.
- Busch, T., M. Johnson, and T. Pioch (2022). Corporate carbon performance data: Quo vadis? *Journal of Industrial Ecology* 26 (1), 350–363.
- Carhart, M. M. (1997). On Persistence in Mutual Fund Performance. *Journal of Finance* 52 (1), 57–82.
- Chen, H., P. De, Y. J. Hu, and B. H. Hwang (2014). Wisdom of crowds: The value of stock opinions transmitted through social media. *The Review of Financial Studies* 27 (5), 1367–1403.
- Cody, E. M., A. J. Reagan, L. Mitchell, P. S. Dodds, and C. M. Danforth (2015). Climate Change Sentiment on Twitter: An Unsolicited Public Opinion Poll.
- Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova (2018). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv preprint arXiv:1810.04805.
- Dietz, S., C. Gollier, and L. Kessler (2018). The Climate Beta. *Journal of Environmental Economics and Management* 87, 258–274.
- Ducoulombier, F. (2021). Understanding the importance of scope 3 emissions and the implications of data limitations. *The Journal of Impact and ESG Investing* 1 (4), 63–71.
- Engle, R., S. Giglio, B. Kelly, H. Lee, and J. Stroebel (2020). Hedging Climate Change News. *The Review of Financial Studies* 33 (3), 1184–1216.
- Faccini, R., R. Matin, and G. S. Skiadopoulos (2021). Dissecting Climate Risks: Are they Reflected in Stock Prices? Available at SSRN 3795964.
- Fama, E. F. and K. R. French (2015). A Five-Factor Asset Pricing Model. *Journal of Financial Economics* 116 (1), 1–22.
- Gentzkow, M., B. Kelly, and M. Taddy (2019). Text as Data. *Journal of Economic Literature* 57 (3), 535–574.
- Gentzkow, M. and J. Shapiro (2010). What Drives Media Slant? Evidence from U.S. Daily Newspapers. *Econometrica* 78 (1), 35–71.
- Görgen, M., A. Jacob, M. Nerlinger, R. Riordan, M. Rohleder, and M. Wilkens (2020). Carbon Risk. Available at SSRN 2930897.
- Heston, S. and N. Sinha (2017). News vs. sentiment: Predicting stock returns from news stories. *Financial Analysts Journal* 73 (3), 67–83.

• Hillert, A., H. Jacobs, and S. Müller (2014). Media makes momentum. *The Review of Financial Studies* 27 (12), 3467–3501.

• Huij, J., D. Laurs, P. A. Stork, and R. C. Zwinkels (2022). Carbon Beta: A Market-Based Measure of Climate Risk. Working Paper.

• Hutto, C. and E. Gilbert (2014). VADER: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text. Conference: *Proceedings of the Eighth International AAAI Conference on Weblogs and Social Media*, 35–71.

• Klaaßeen, L. and C. Stoll (2021). Harmonizing corporate carbon footprints. *Nature Communications* 12, 6149.

• Malo, P., A. Sinha, P. Korhonen, J. Wallenius, and P. Takala (2014). Good debt or bad debt: Detecting semantic orientations in economic texts. *Journal of the Association for Information Science and Technology* 65.

• McCombs, M. E. and D. L. Shaw (1972). The Agenda-Setting Function of Mass Media. *The Public Opinion Quarterly* 36 (2), 176–187.

• Meinarding, C., Y. Schöler, and P. Zhang (2020). Shocks to Transition Risk. Available at SSRN.

• Nordhaus, W. and P. Sztorc (2023). DICE 2013R: Introduction and User's Manual. Yale University and the National Bureau of Economic Research, USA.

• Pástor, L., R. F. Stambaugh, and L. A. Taylor (2021). Sustainable Investing in Equilibrium. *Journal of Financial Economics* 142 (2), 550–571.

• Peng, H., L. Xu, L. Bing, F. Huang, W. Lu, and L. Si (2020, Apr.). Knowing what, how and why: A near complete solution for aspect-based sentiment analysis. *Proceedings of the AAAI Conference on Artificial Intelligence* 34 (05), 8600–8607.

• Pörtner, H. O., D. C. Roberts, H. Adams, C. Adler, P. Aldunce, E. Ali, ..., and J. Birkmann (2022). Climate change 2022: Impacts, adaptation and vulnerability. *IPCC Sixth Assessment Report*.

• Roncalli, T., T. Le Guenedal, F. Lepetit, and T. Sekine (2020). Measuring and Managing Carbon Risk in Investment Portfolios. Available at SSRN 3681266.

• Sanh, V., L. Debut, J. Chaumond, and T. Wolf (2019). Distilbert, a distilled version of BERT: smaller, faster, cheaper and lighter. *CoRR abs/1910.01108*.

• Shiller, R. (2000). Irrational Exuberance. Princeton University Press.

• Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin (2017). Attention is all you need. *Advances in Neural Information Processing Systems 30*.

• Vaucher, B., V. Bouchet, and B. Herzog (2023). Look up! A Market-Measure of the Long-Term Transition Risks in Equity Portfolios. Working Paper. Scientific Portfolio.

• Webersinke, N., M. Kraus, J. Bingler, and M. Leippold (2021). Climatebert: A pretrained language model for climate-related text. arXiv preprint arXiv:2110.12010v3.

About Amundi

Amundi, a leading responsible investor engaged in the just transition

Amundi, the leading European asset manager and pioneer in responsible investing, ranks among the top 10 global players[1] and offers its 100 million clients savings and investment solutions in active and passive management, in traditional or real assets.

Already incorporating ESG criteria in 100% of its actively managed open-ended funds[2] and engaged in the energy transition and social cohesion, Amundi launched its ESG Ambitions 2025 plan, with 10 objectives to accelerate its ESG transformation and pave the way towards carbon neutrality by 2050. Amundi is also a member of the Net Zero Asset Managers initiative and a founder of the Investors for a Just Transition coalition.

With six international investment hubs[3], Amundi's clients benefit from the expertise of 5,300 employees in over 35 countries. A subsidiary of Crédit Agricole and listed on the stock exchange, Amundi currently manages €1904 trillion assets[4].

Amundi
ASSET MANAGEMENT

1 - Source: IPE "Top 500 Asset Managers" published in June 2022, based on assets under management as at 31/12/2021
2 - Since December 2021
3 - Boston, Dublin, London, Milan, Paris and Tokyo
4 - Amundi data as at 31/12/2022

About EDHEC-Risk Climate Impact Institute

Exploring double materiality – studying the impact of climate-change related risks on finance and the effects of finance on climate change mitigation and adaptation

Institutional Context

Established in France in 1906, EDHEC Business School now operates from campuses in Lille, Nice, Paris, London, and Singapore. With more than 110 nationalities represented in its student body, some 50,000 alumni in 130 countries, and learning partnerships with 290 institutions worldwide, it truly is international. The school has a reputation for excellence and is ranked in the top 10 of European business schools (Financial Times, 2021).

For more than 20 years, EDHEC Business School has been pursuing an ambitious research policy that combines academic excellence with practical relevance. Spearheaded by EDHEC-Risk Institute, its aim is to make EDHEC Business School a key academic institution of reference for decision makers in those areas where is excels in expertise and research results. This goal has been delivered by expanding academic research in these areas and highlighting their practical implications and applications to decision makers. This approach has been complemented by strategic partnerships and business ventures to accelerate the transfer of scientific innovation to the industry and generate financial benefits for the School and its constituencies.

In the Fall of 2022, EDHEC-Risk Institute became EDHEC-Risk Climate Impact Institute (EDHEC-Risk Climate). This transition reflects the importance assigned by the School to sustainability issues and builds on the foundations laid by EDHEC-Risk Institute research programmes exploring the relationships between climate change and finance.

Mission and Ambitions

EDHEC-Risk Climate's mission is to help private and public decision makers manage climate-related financial risks and make the best use of financial tools to support the transition to low-emission and climate-resilient economies.

Building upon the expertise and industry reputation developed by EDHEC-Risk Institute, EDHEC-Risk Climate's central ambition is to become the leading academic reference point helping long-term investors manage the risk and investment implications of climate change and adaptation and mitigation policies.

EDHEC-Risk Climate also aims to play a central role in helping financial supervisors and policy makers assess climate-related risks in the financial system and provide them with financial tools to mitigate those risks and optimise the contribution of finance to climate change mitigation and adaptation.

The delivery of these ambitions is centred around two long-term research programmes and a policy advocacy function.

The research programmes respectively look at the Implications of Climate Change on Asset Pricing and Investment Management and the Impact of Finance on Climate Change Mitigation and Adaptation.

The Institute also supports the integration of climate issues into the research agenda of the School's other financial research centres and into the product offering of the School's business ventures. In particular, it helps leading infrastructure research centre EDHECinfra build capacity on sectoral alignment and transition plans.

EDHEC-RISK
Climate Impact
INSTITUTE

EDHEC-Risk Climate Impact Institute Publications

2023

• Rebonato. R, Kainth, D. Melin, L, and D. O'Kane. Optimal Climate Policy with Negative Emissions. (March).

• Chini, E and M. Rubin.Time-varying Environmental Betas and Latent Green Factors (April).

• Maeso, J. M. and D. O'Kane. The Impact of Climate Change News on Low-minus-High Carbon Intensity Portfolios. (June).

For general enquiries about this publication, please contact: **research@climateimpactedhec.com**

For press or media-related enquiries, please contact: **maud.gauchon@climateimpactedhec.com**

EDHEC-Risk Climate Impact Institute - Nice
400 promenade des Anglais
BP 3116 - 06202 Nice Cedex 3
France
Tel. +33 (0)4 93 18 78 87

EDHEC-Risk Climate Impact Institute - London
10 Fleet Place, Ludgate
London EC4M 7RB
United Kingdom
Tel: + 44 207 332 5600

climateimpact.edhec.edu

EDHEC-RISK
Climate Impact
INSTITUTE