Hedging Climate Change News

Robert F. Engle
NYU Stern and NBER

Stefano Giglio
Yale University, NBER, and CEPR

Bryan Kelly
Yale University, AQR Capital Management, and NBER

Heebum Lee
NYU Stern

Johannes Stroebel
NYU Stern, NBER, and CEPR

We propose and implement a procedure to dynamically hedge climate change risk. We extract innovations from climate news series that we construct through textual analysis of newspapers. We then use a mimicking portfolio approach to build climate change hedge portfolios. We discipline the exercise by using third-party ESG scores of firms to model their climate risk exposures. We show that this approach yields parsimonious and industry-balanced portfolios that perform well in hedging innovations in climate news both in sample and out of sample. We discuss multiple directions for future research on financial approaches to managing climate risk. (*JEL* G11, G18, Q54)

Received December 7, 2017; editorial decision December 12, 2018 by Editor Andrew Karolyi. Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

Earth's climate is changing, but uncertainty around the trajectory and the economic consequences of climate change is substantial. As a result, investors around the world desire products that allow them to hedge against the

This version: May 23 2019. We thank Harrison Hong, Andrew Karolyi, and Ross Valkanov; participants at the Climate Finance Workshop at Columbia University, the Climate Finance Conference at Imperial College, and the Volatility Institute Conference at NYU Stern; and a number of anonymous referees for helpful comments. The Norwegian Finance Initiative and the Global Risk Institute provided generous grant support. We thank Konhee Chang for outstanding research assistance. Supplementary data can be found on *The Review of Financial Studies* web site. Send correspondence to Robert Engle, NYU Stern, 44 West 4th Street, New York, NY 10012; telephone: 212-998-0710. E-mail: rengle@stern.nyu.edu. Disclaimer: The views and opinions expressed are those of the authors and do not necessarily reflect the views of AQR Capital Management, its affiliates, or its employees; do not constitute an offer, solicitation of an offer, or any advice or recommendation, to purchase any securities or other financial instruments, and may not be construed as such.

doi:10.1093/rfs/hhz072

realizations of climate risk. Because of the long run and nondiversifiable nature of climate risk, standard futures or insurance contracts in which one party promises to pay the other in the event of a climate disaster are difficult to implement. Indeed, no counterparty could credibly guarantee to pay claims during a climate disaster event that might materialize in many decades, in part because a bad outcome would mandate all contracts to be paid at the same time. Individual investors are therefore largely constrained to self-insure against climate risk.

In this paper, we propose an approach for constructing climate risk hedge portfolios using publicly traded assets. We follow a dynamic hedging approach similar to Black and Scholes (1973) and Merton (1973). In this approach, rather than buying a security that directly pays off in the event of a future climate disaster, we construct portfolios whose short-term returns hedge *news* about climate change over the holding period. By hedging, period by period, the innovations in news about long-run climate change, an investor can ultimately hedge her long-run exposure to climate risk. In the short run, such a portfolio differs from the Markowitz mean-variance efficient portfolio and will thus exhibit a lower Sharpe ratio; but, in the long run, the dynamic hedging approach will compensate investors for losses that arise from the realization of climate risk.

The primary objective of this paper is to provide a rigorous methodology for constructing portfolios that use relatively easy-to-trade assets (equities) to hedge against risks that are otherwise difficult to insure. We show that our approach, which uses tools from standard asset pricing theory, does indeed allow us to construct portfolios that can successfully hedge climate news out of sample. Having said that, we do not view our resultant hedge portfolios as the definitive best hedges against climate change risk, but instead as a starting point for further exploration. Along these lines, we will discuss many valuable directions for future research on using financial markets to hedge climate risk.

The first challenge to implementing a dynamic hedging strategy for climate risk is to construct a time series that captures news about long-run climate risk, and which can therefore help us to construct an appropriate hedge target. We start from the observation that when there are events that plausibly contain such information about changes in climate risk, this will likely lead to newspaper coverage of these events; indeed, newspapers may even be the direct source that investors use to update their subjective probabilities of climate risks. Our approach in this paper therefore is to extract a climate news series from textual analysis of news sources. A wide range of events covered in newspapers can carry potentially relevant information. Indeed, the list of topics that are often covered by newspapers in relation to discussions about climate risk includes extreme weather events (e.g., floods, hurricanes, droughts, wildfires, extreme temperatures), physical changes to the planet (e.g., sea level changes, glacial melting, ocean temperatures), regulatory discussions, technical progress in alternative fuel delivery, and the price of fossil fuels.

We construct two complementary indices that measure the extent to which climate change is discussed in the news media. The first index is calculated as the correlation between the text content of *The Wall Street Journal* (WSJ) each month and a fixed climate change vocabulary, which we construct from a list of authoritative texts published by various governmental and research organizations. The WSJ is among the most salient media outlets for market participants, and thus our index captures the intensity of climate change discourse that is accessible to the investment community at very low cost.

Our WSJ Climate Change News Index associates increased climate change reporting with news about elevated climate risk, based on the idea that climate change primarily rises to the media's attention when there is a cause for concern. An alternative approach is to directly differentiate between positive and negative news in our index construction. To this end, we construct a second news-based climate index that is designed to focus specifically on bad news about climate change. This index applies sentiment analysis to climate-related articles to measure the intensity of *negative* climate news in a given month.

In this paper, we do not try to distinguish between different *types* of climate change news. In particular, we do not distinguish between news about physical damages from climate change and news about regulatory risks that are related to climate change. These two risk measures might move independent from each other. For example, the Paris accord, which led to a pledge to reduce carbon emissions, might have represented an increase in regulatory risk and a decrease in physical risk. Separately measuring news series about physical and regulatory climate risk represents an interesting avenue for future research. Also, our focus in this paper is on *global* climate change news. Our indices ignore news about local climate events, which are not covered in the WSJ or in a large cross-section of newspapers.

The second step in implementing our dynamic hedging strategy is to construct portfolios that allow us to hedge innovations in these two news series. In particular, we seek to systematically explore which stocks rise in value and which stocks fall in value when (negative) news about climate change materializes. Then, by constructing a portfolio that overweights stocks that perform well on the arrival of such negative news, an investor will have a portfolio that is well-positioned to profit the next time when such news about climate change materializes. Continued updating of this portfolio based on new information about the relationship between climate news and stock returns will ultimately lead to a portfolio which is long the winners from climate change and short the losers.

Our econometric approach to forming such hedge portfolios follows standard methods in the asset pricing literature. If climate risk represents a risk factor for asset markets (i.e., if it is a factor that drives the comovement of different assets), it is possible to construct a well-diversified portfolio the return of which isolates the exposure to that risk factor. Investors can then hedge their climate risk exposure by trading this portfolio without changing their exposures to the other risk factors in their portfolios. Various approaches to construct such hedge portfolios have been proposed in the literature. The two main ones are cross-sectional regressions like Fama-MacBeth (in which the hedging portfolio is obtained through period-by-period cross-sectional regressions of asset returns onto exposures to the risk factors), and direct projections of the risk factors onto a set of asset returns (the so-called "mimicking portfolio approach").[1] Among the many prominent papers in this literature are Fama and MacBeth (1973), Chen, Roll, and Ross (1986), Huberman, Kandel, and Stambaugh (1987), Breeden, Gibbons, and Litzenberger (1989), Lamont (2001), Balduzzi and Robotti (2008), Lönn and Schotman (2017), and Roll and Srivastava (2018). Giglio and Xiu (2018) study the asymptotic properties of the different estimators in large cross-sections, and investigate their robustness to model specification errors. In this paper, we will apply the mimicking portfolio approach, as advocated by Lamont (2001).

The challenge with implementing this mimicking portfolio approach is that we only observe a limited number of months of climate news realizations, but have a large set of assets that we could use to form hedge portfolios. This leads to concerns about data mining, where we might end up constructing hedge portfolios that perform very well in sample but that are not stable going forward. To address this concern, we use characteristics that proxy for a firm's exposure to climate risk to parsimoniously parameterize the weights of the hedge portfolios. For example, one such characteristic might be the carbon footprint of each firm. In particular, it might be that when there is news about increasing climate risk, individuals will buy low-carbon-footprint stocks and sell high-carbon-footprint stocks. If this were the case, one could construct a portfolio that increases in value when there is (negative) news about climate risk using thousands of long and short positions based on just one parameter, the firms' carbon footprints.

We implement this characteristics-based approach by using firm-level environmental performance scores constructed by the ESG ("**E**nvironmental, **S**ocial, and **G**overnance") data providers MSCI and Sustainalytics to proxy for firms' climate risk exposure.[2] In particular, we use these scores as characteristics on which to sort individual stocks to form portfolios. We then construct the final hedge portfolios by projecting innovations in our climate news indices onto these ESG-characteristic-sorted portfolios, together with standard Fama-French factor-sorted portfolios (market, size, and value).

[1] The literature on cross-sectional regressions, like Fama-MacBeth, typically focuses on estimating the risk premiums of the factor, but risk premiums are simply the average excess returns of the corresponding hedge portfolios.

[2] Again, there is a question of what *type* of climate change risk exposure these measures capture. Specifically, they may more closely capture regulatory risks than physical risks, and other characteristics could be added to the analysis to capture different types of climate change exposures. For example, one could perhaps proxy for firms' physical climate risk by the distance of firms' headquarters or production facilities from the sea. Exploring different firm-level measures of climate risk exposure (both physical and regulatory) constitutes an interesting avenue for future research.

When we compare our hedge portfolios to alternative hedge portfolios that add simple industry bets (such as positions in the energy exchange-traded fund XLE) to the standard Fama-French factors, we find that our ESG-characteristic-based mimicking portfolios procedure produces hedge portfolios that perform better than the alternatives in hedging innovations in climate risk. In particular, our portfolios deliver higher in-sample and out-of-sample correlations with those innovations. For example, the return of the hedge portfolio based on the Sustainalytics E-Scores achieves out-of-sample correlations with the WSJ index innovations as high as 30%. Our hedge portfolios also do not resemble industry bets; rather, they identify, both within and across industries, those firms with the largest exposures to climate change risk, yielding a climate hedge portfolio that is relatively industry-balanced.

Our work contributes to a burgeoning literature that studies how climate change affects asset markets, and how asset markets in turn may affect the dynamics of climate change. Andersson, Bolton, and Samama (2016) propose a passive investment strategy tilted to low-carbon stock as a hedge against climate risk, while Choi, Gao, and Jiang (2018) explore how investors update their information about climate risk. Hong, Li, and Xu (2019) investigate whether international stock markets efficiently price drought risk, and Kumar, Shashwat, and Wermers (2018) explore whether fund managers misestimate the risk of climate disasters. Baldauf, Garlappi, and Yannelis (2018), Bakkensen and Barrage (2018), Bernstein, Gustafson, and Lewis (2019), Giglio et al. (2018), and Murfin and Spiegel (2018) explore the pricing of climate risk in real estate markets, while Giglio, Maggiori, and Stroebel (2015), Giglio et al. (2018) use real estate pricing data to back out very long-run discount rates that are appropriate for valuing projects aimed at mitigating climate change. Daniel, Litterman, and Wagner (2015) apply standard asset pricing theory to calibrate the social cost of carbon.

1. Construction of the Hedge Portfolios: Theory

This section discusses our methodology to construct portfolios that hedge news about climate change. We denote by r_t an $n \times 1$ vector of excess returns over the risk-free rate of n assets at time t. We assume that these returns follow a linear factor model, in which asset returns are driven by innovations in climate news, which we denote by CC_t, as well as by p other (tradable or nontradable) risk factors v_t:

$$\underbrace{r_t}_{n\times 1} = (\underbrace{\beta_{CC}}_{n\times 1}\underbrace{\gamma_{CC}}_{1\times 1} + \underbrace{\beta_{CC}}_{n\times 1}\underbrace{(CC_t - E[CC_t])}_{1\times 1}) + (\underbrace{\beta}_{n\times p}\underbrace{\gamma}_{p\times 1} + \underbrace{\beta}_{n\times p}\underbrace{v_t}_{p\times 1}) + \underbrace{u_t}_{n\times 1}. \quad (1)$$

The vectors β_{CC} and β are risk exposures of the n assets to the climate news factor and the other p factors, respectively. Similarly, γ_{CC} and γ are the corresponding risk premiums for the climate news factor and the other risk factors. Finally, u_t is an idiosyncratic error term. In this basic setup, the risk exposures are constant; we relax this assumption below.

Our objective is to construct a hedge portfolio for CC_t. This is defined as a portfolio that has unit exposure (beta) to climate risk shocks CC_t, but no exposure to any of the other p factors v_t. This ensures that investors can change their exposure to climate risk by trading in this portfolio, without modifying their exposure to the other risk factors. The asset pricing literature has followed two main approaches to construct hedge portfolios: the Fama-MacBeth cross-sectional regression approach and the mimicking portfolio approach. Giglio and Xiu (2018) derive theoretical properties of the two estimators in large-dimensional settings.

In this paper, we follow the mimicking portfolio approach; for completeness, Appendix A.1 provides a review of the Fama-MacBeth procedure in our setting. In the mimicking portfolio approach, the climate risk factor CC_t is directly projected onto a set of excess returns of a set of portfolios, $\tilde{r}_t$:

$$CC_t = \xi + w'\tilde{r}_t + e_t. \tag{2}$$

The hedge portfolio for CC_t is constructed using the weights $\hat{w}$ estimated from this regression; its excess return is $h_t^{CC} = \hat{w}'\tilde{r}_t$. The vector e_t captures the measurement error in CC_t, so that this approach explicitly accounts for potential measurement error in the climate risk factor CC_t. A sufficient condition for this procedure to recover the desired hedge portfolio for climate news is that the returns of the portfolios used in the projection, $\tilde{r}$, span the same space as the true factors, (CC_t, v_t).[3]

1.1 Implementation and construction of the hedge portfolios

To build hedge portfolios using the mimicking portfolio approach, we choose a set of projection portfolios which are well diversified, so that idiosyncratic error is approximately eliminated, and which at the same time capture different dimensions of risk, so that their returns $\tilde{r}_t$ span the factor space. The portfolios used in the projection need to satisfy one further requirement. In particular, the setup described in Equation 1 includes the assumption that the risk exposures of the assets used in the estimation are constant over time. We therefore need to construct the portfolios $\tilde{r}$ in such a way that their exposures to the underlying risk factors are constant. A standard approach to achieve this is to form portfolios by sorting assets on characteristics. Indeed, to the extent that risk exposures of individual assets directly depend on these characteristics, sorting the assets by characteristics will ensure that the resultant portfolios have constant risk exposures. We follow this approach and choose a matrix

3 Formally, write the model in the following compact form by calling f the vector of all factors: $f_t \equiv (CC_t, v_t)$, with covariance matrix Σ_f and β_f the matrix of betas: $\beta_f = (\hat{\beta}_{CC}, \hat{\beta})$. Call η the $(p+1)\times 1$ vector with 1 as the first element and 0 everywhere else, so that $CC_t = \eta' f_t$. The population vector of weights w is $Var(\tilde{r}_t)^{-1} Cov(\tilde{r}_t, CC_t)$. If returns $\tilde{r}_t$ span the same space as the true factors, this means there exists an invertible matrix H such that $\tilde{r}_t = Hf_t$. We can then write $w = (H\Sigma_f H')^{-1} H\Sigma_f \eta = H'^{-1}\eta$. The return of this portfolio is $h_t^{CC} = w'\tilde{r}_t = w'Hf_t = \eta' H^{-1} Hf_t = \eta' f_t = CC_t$.

of firm-level characteristics Z_t, appropriately cross-sectionally normalized, to construct the portfolio returns as

$$\tilde{r}_t = Z'_{t-1} r_t,$$

where r_t are excess returns of individual stocks, and portfolio weights are equal to the normalized characteristics.[4] Substituting this expression into Equation 2, we write

$$CC_t = \xi + w' Z'_{t-1} r_t + e_t. \tag{3}$$

Equation 3 can be interpreted in two ways. It can either be thought of as a projection of the hedge target CC_t onto characteristic-sorted portfolios $Z'_{t-1} r_t$ that are assumed to have constant risk exposure and that span the entire factor space. Alternatively, it can be thought of as a constrained projection of CC_t on *all* individual asset returns r_t, but with time-varying weights $w' Z'_{t-1}$; the weights are modeled as a linear function of characteristics, so that any individual firm's weight depends on its risk exposure to the different factors. Equation 3 therefore performs a one-step dimension reduction that estimates the hedge portfolio, while modeling the time variation in risk exposures.

2. Hedging Climate Change News

In this section, we implement the mimicking portfolio approach to hedging climate risk that we described above. As we have highlighted in the Introduction, the relevant performance measure for the resultant hedge portfolios is how well they hedge innovations to climate news out of sample. However, given the relatively short time period for which we observe measures of both climate news and firm-level climate risk exposures, there are a limited number of out-of-sample test periods on which to evaluate the climate hedge portfolios.[5] As will become apparent below, there are many degrees of freedom in how to construct these hedge portfolios, including decisions about how to construct measures of firm-level climate risk exposures and about what other portfolios to include in regression 2. As a result, there is the danger of optimizing over these degrees of freedom to construct portfolios that provide optimal out-of-sample hedges to climate news over the short period we observe, but that may not be effective at hedging this news going forward.

To avoid such data mining concerns, we will clearly describe the various choices we encountered in the construction of the climate hedge portfolios.

[4] Note that we are exclusively working with excess returns, so there are no theoretical constraints on portfolio weights.

[5] In addition, even if we could easily extend our time series further into the past, it is unclear whether the additional sample periods would help us with constructing climate hedge portfolios today. In particular, it is plausible that climate risk has only started to be priced in stocks in recent years as investors' attention to this risk has increased. Indeed, some indirect evidence for such a suggestion comes from the fact that demand for ESG measures has substantially increased over the past few years. As a result, it is unclear whether firms with different climate risk exposures have had different excess returns in response to climate news that materialized in, say, the 1990s.

However, instead of optimizing over these degrees of freedom to find a portfolio that optimally hedges climate news over our short test sample, we make choices that appear reasonable to us, and that will hopefully lead to stable approaches to hedging climate news that is yet to occur. This discussion will highlight a number of important directions in which to further develop these climate hedge portfolios, and longer time series of measures of climate news and climate risk exposures will allow for more systematic ways of testing the true out-of-sample performance of different climate hedge portfolios.

2.1 Measuring climate change news

The first step in our analysis is to construct an index that measures innovations in news about climate risk. A variety of choices must be made when constructing this hedge target. How should we identify the news sources that reflect the information investors use in their climate risk-based investment decisions? Once we identify the appropriate news, how do we measure its relative intensity over time? How do we quantify the extent of good news versus bad news? And should one differentiate among subtypes of climate news (such as news about physical climate risks versus news about regulatory risks)?

Below, we follow two alternative approaches to building a climate news index. We believe they have the virtues of breadth and simplicity and offer scope for comparing trade-offs in some of our construction choices. At the same time, our indices have obvious imperfections and leave much room for other researchers to propose adjustments. Indeed, different investors might want to make different choices to ours in order to optimally align their hedge targets with the overall climate exposures of the rest of their portfolios. For example, investors with a strong coastal real estate portfolio might want to focus more on news about physical climate risk (because such real estate is strongly exposed to rising sea levels), while investors with a strong exposure to the coal industry might want to focus more on news about regulatory interventions in response to climate risk.[6]

2.1.1 *Wall Street Journal* climate change news index. The first index that we construct is based on climate news coverage in *The Wall Street Journal* (WSJ). Two considerations support our use of the WSJ. One is a desire to measure news that is relevant to and salient for investors concerned about climate risks, and the WSJ is among the most important media sources consumed by financial market participants. The second advantage is that we have access to the full text of WSJ articles since the early 1980s, which provides us with complete flexibility in choosing how to build the climate news index from raw news content.

[6] In addition, some researchers and investors may want to expand the list of publications they consider beyond our newspaper-based approach. Additional publications of interest could include coverage in scientific journals or social media posts.

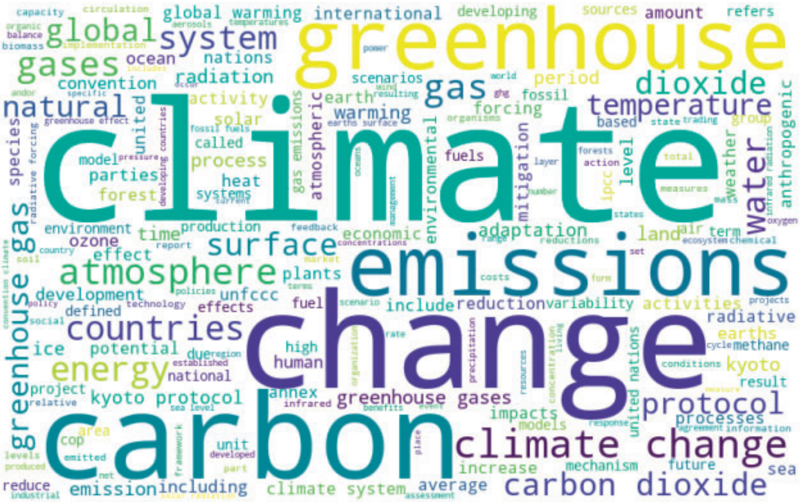

Figure 1
Climate change vocabulary
Word cloud summary of climate change vocabulary from a corpus of seventy-four authoritative climate change texts. Term sizes are proportional to their frequency in the corpus.

To quantify the intensity of climate news coverage in the WSJ, we compare the news content to a corpus of authoritative texts on the subject of climate change. In particular, we collect 19 climate change white papers from sources such as the Intergovernmental Panel on Climate Change (IPCC), the Environmental Protection Agency (EPA), and the U.S. Global Change Research Program. We complement these white papers with 55 climate change glossaries from sources such as the United Nations, NASA, the IPCC, the EPA, and others. Appendix A.2 presents the full list of these authoritative texts. We aggregate the seventy-four text documents into a "Climate Change Vocabulary (CCV)," which amounts to the list of unique terms (stemmed unigrams and bigrams) and the associated frequency with which each term appears in the aggregated corpus. Figure 1 provides an illustration of the CCV in the form of a word cloud, with term sizes proportional to their frequency. We form an analogous list of term counts for the WSJ. Each (daily) edition of WSJ is treated as a "document," and term counts are tallied separately for each document. Next, we convert WSJ term counts into "term frequency–inverse document frequency," or *tf-idf*, scores. Common terms that appear in most documents earn low scores because they are less informative about any individual document's content (they have low *idf*), as do terms that are rare in a given article (they have low *tf*). The *tf-idf* transformation defines the most representative terms in a given document to be those that appear infrequently overall, but frequently in that specific document (see Gentzkow, Kelly, and Taddy 2018).

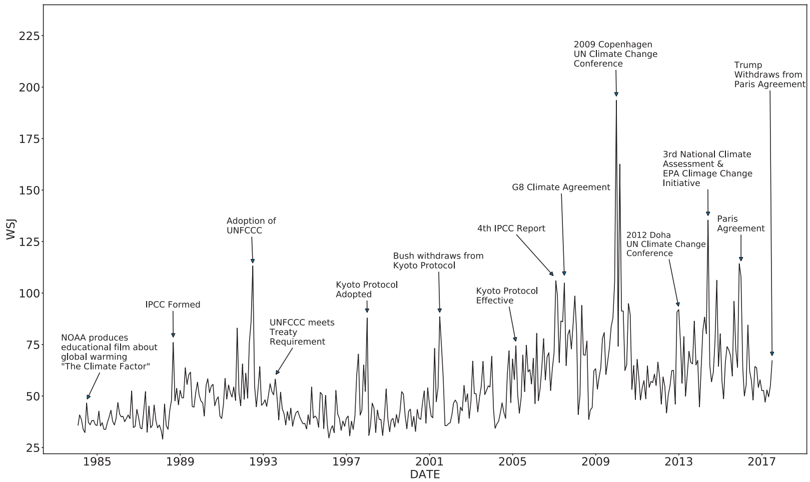

Figure 2
WSJ Climate Change News Index
This figure shows the WSJ Climate Change News Index from 1984 to 2017, annotated with climate-relevant news announcements.

The main choice going into our index construction is to treat the CCV as our definition of phraseology associated with climate change discourse. That is, our CCV takes a stand on the specific terms, and their relative usage intensity, to identify news about the topic of climate change. Like with the WSJ, we convert Climate Change Vocabulary term counts into *tf-idf*. We treat the aggregated CCV as a single document when calculating term frequencies, and apply the inverse document frequency calculation from the WSJ corpus.[7]

Finally, we construct our daily climate change index as the "cosine similarity" between the *tf-idf* scores for the CCV and each daily WSJ edition. Days in which the WSJ uses the same terms in the same proportion as the CCV earn an index value of one, while days in which the WSJ uses no words from the CCV earn an index value of zero. Approximately speaking, our raw WSJ Climate Change News Index describes the fraction of the WSJ dedicated to the topic of climate change each day, as defined by the texts that underlie the CCV. We scale this index by a factor of 10,000 to allow interpretation of the magnitudes of innovations in the index, which will represent our eventual hedge targets.

Figure 2 shows a time series of the WSJ Climate Change News Index since 1984. The figure shows that the intensity of climate news coverage has steadily increased since about the year 2000. In addition, the climate risk index spikes during salient climate events, such as the adoption of global climate treaties

[7] The choice to use the same *idf* for WSJ and CCV counts ensures that the document-frequency weights of CCV terms match the weights of WSJ terms. If we were to instead calculate *idf* based on the corpus of authoritative climate texts, we would down weight the most informative climate change terms and unduly distort the measurement of climate change discourse in the WSJ.

(e.g., the UNFCCC or the Kyoto protocol), or important global conferences to battle climate change (e.g., the 2009 UN Climate Change Conference in Copenhagen).

2.1.2 Crimson Hexagon's negative sentiment climate change news index. Implicit in our construction of the WSJ Climate Change News Index is the assumption that the number of climate change discussions increase when climate risk is elevated. In other words, the WSJ index embeds the view that, when it comes to climate change, no news is good news. While we view this as a plausible assumption, there is a risk of inaccurately capturing discussions of positive climate news (e.g., news about new mitigation technologies) as increases in climate risk. A separate potential shortcoming of the WSJ index is that, being based on a single source, it may be too narrow in its quantification of climate discourse among investors.

To address these possible concerns, we study a second news-based climate risk index that is designed to focus specifically on negative climate news, and that is drawn from a much more expansive collection of news articles. For this purpose, we use the services of the data analytics vendor Crimson Hexagon (CH). Starting in May 2008, Crimson Hexagon has collected a massive corpus of over one trillion news articles and social media posts. The underlying news sources cover over 1,000 outlets, including the WSJ, *The New York Times*, *The Washington Post*, Reuters, BBC, CNN, and Yahoo News. Coverage in terms of total articles available expands over time. Cross-sectionally, the distribution of article counts is fairly evenly distributed across news outlets, with the top-100 outlets accounting for approximately 14% of the total article count. For a given user-provided search term, CH applies a variety of proprietary natural language processing analytics, such as sentiment analysis and topic modeling, to construct time series of the sentiment of coverage of that term across the sources it collects.

We provide CH with the search phrase "climate change" and restrict our analysis to discussions in the news media (i.e., we exclude social media). Based on these choices for terms and content sources, CH provided us with an array of indices that summarize the total number of articles that include climate change news, as well as the fraction of those summarized to contain positive and negative climate change news. It also provided indices for further sentiment subcategories (e.g., fear, joy, anger), as well as a topic decomposition of climate-related articles. Thus, there are many potential degrees of freedom in using Crimson Hexagon data to construct a climate news series. For example, we could tune our choice of search terms, or optimize across each of the finer indices that CH supplies for any given set of search terms. As described above, given the brevity of our data sample, we need to guard against data mining, and we do so in this case by restricting ourselves to the most obvious search term ("climate change") and focusing on the most obvious category that resolves our desire for "signed" news, namely those that CH categorizes as basic "negative

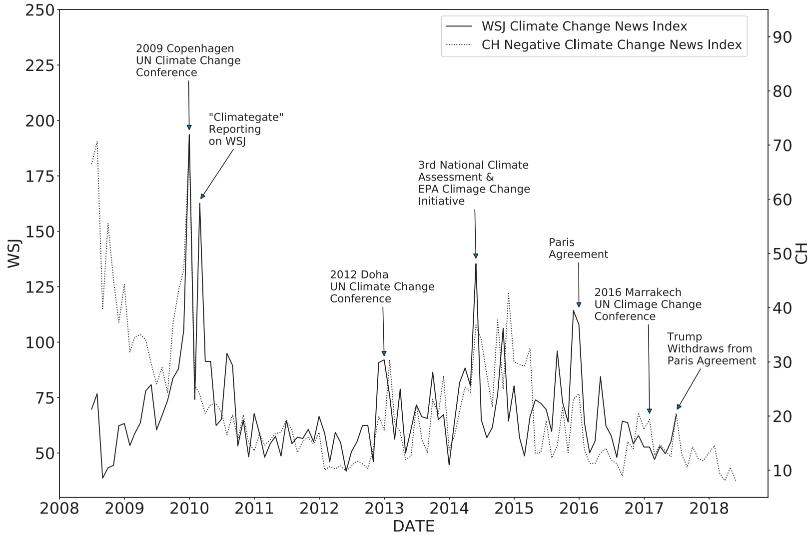

Figure 3
CH Negative Climate Change News Index
This figure shows the CH Negative Climate Change News Index from 2008 to 2017, overlaid against the WSJ Climate Change News Index, and annotated with climate-relevant news announcements.

sentiment." We calculate our CH Negative Climate Change News Index as the share of all news articles that are both about "climate change" and that have been assigned to the "negative sentiment" category; we multiply this measure by 10,000 in order to interpret the magnitudes of innovations in the index.

Figure 3 plots the time series of the CH Negative Climate Change News Index, in addition to that of the WSJ Climate Change News Index for comparison. Both indices regularly spike around salient climate events, such as climate conferences. The initial level of the CH index is somewhat higher than that of the WSJ index, though this is during a period for which Crimson Hexagon has relatively little data; this is also a period that will not be included in our final analysis (as we discuss below, our empirical analysis starts in September 2009, the first month for which we observe complete coverage of firm-level climate risk exposures). Interestingly, the WSJ index spikes in a number of instances in which the CH index does not. One of these was in early 2010, a period during which the WSJ extensively reported on the "Climategate" controversy.[8]

2.1.3 Constructing hedge targets. To measure innovations in climate news, we average the daily values for the WSJ Climate Change News Index and

[8] The Climategate controversy involved the publication of emails obtained through hacking a server at the Climatic Research Unit at the University of East Anglia. Several climate change "skeptics" alleged that these emails documented global warming to be a scientific conspiracy, with scientists manipulating data.

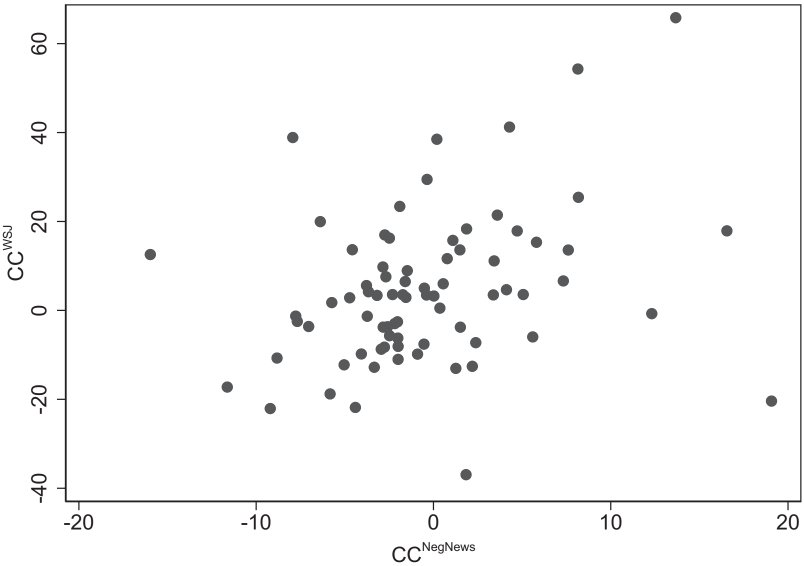

Figure 4
Correlation across CC_t measures
This figure shows a scatterplot highlighting the correlation across our two climate hedge targets, CC^{WSJ} and CC^{NegNes}. Each observation corresponds to 1 month between September 2009 to December 2016. The correlation coefficient is 0.30.

CH Negative Climate Change News Index to the monthly level, and then construct values of CC_t as residuals from an AR(1) model. This gives us our two monthly hedge targets: CC_t^{WSJ}, which captures innovations in the WSJ Climate Change News Index, and $CC_t^{NegNews}$, which captures innovations in the CH Negative Climate Change News Index. Figure 4 shows the correlation across these measures across the 88 months that will be included in our final analysis, September 2009 to December 2016. The correlation coefficient is 0.3, which suggests that, although both measures capture common elements of climate risk, they are by no means identical. As we have discussed above, which of the two series (or any one of the potential alternative series that we could have constructed) represents the ideal hedge target depends on the precise application; as a result, we view the construction of alternative hedge targets as an exciting area for further research.

2.2 Potential assets in hedge portfolios

After defining the hedge targets, the second step in implementing the mimicking portfolio hedge approach described in Section 1 is to determine the universe of assets used to build the hedge portfolio. In this project, we focus on constructing hedge portfolios using U.S. equities as the underlying assets. We obtain monthly individual U.S. stock return data from CRSP. We include only common equity securities (share codes 10 and 11) for firms traded on the NYSE, AMEX and

NASDAQ. Following Amihud (2002) and many others, we exclude penny stocks, defined as stocks with a price below $5 at the time of portfolio formation. This is to avoid including stocks whose returns are dominated by market microstructure issues. We also drop microcap stocks, defined as stocks with a market capitalization in the bottom 20% of the sample traded on the NYSE, following the observation in Fama and French (2008) that the returns of hedge portfolios obtained from long-short positions can be distorted by the inclusion of such microcaps (see also the discussion in Hou, Xue, and Zhang 2015).

2.3 Measuring climate risk exposures

Having identified the set of possible assets to include in the hedge portfolio, the next empirical challenge is to systematically measure different firms' exposures to climate risk, that is, to identify the characteristics in Z_t that drive such exposures. Our approach in this paper is to build on measures of firms' environmental exposures produced by third-party ESG data providers. Indeed, there has been a growing interest in ESG investing among investors who are increasingly demanding assets that fulfill certain environmental ("E"), social ("S"), and governance ("G") criteria.[9] Given this trend, measuring the ESG characteristics of firms has become an important task for investors, and firm-level ESG scores are available from numerous providers that collect raw data gathered from sources such as firms' disclosures, SEC filings, and reports by governments or NGOs. These raw data are then translated into numerical ESG scores using proprietary algorithms.[10]

Our study uses information on firm-level ESG scores from two leading data providers, MSCI and Sustainalytics.[11] Both data providers construct various subscores that evaluate firms on different aspects of their ESG performance. From these subscores, we choose the broadest scores that plausibly proxy for firms' exposure to climate risk.

2.3.1 MSCI. We obtained from MSCI a data set of annual firm-level ESG scores between 1995 and 2016.[12] MSCI evaluates firms along several subcategories that capture either positive or negative environmental performance; Appendix A.3 presents the full list of subcategories. Each

[9] According to The U.S. SIF Foundation, the dollar value of ESG assets owned by institutional investors grew to $4.73 trillion in 2016, an increase of 11% a year since 2005.

[10] As noted in the Introduction, ESG scores may capture specific notions of climate change exposure; for example, they may better capture exposure to regulatory risks than exposure to physical damages from climate risks. The methodology in this paper could be easily applied using other firm characteristics that may capture different types of climate risk exposures.

[11] The number of ESG data providers, including firms such as Arabesque and TruValue Labs, is growing. Analyzing which of these E-Scores results in the optimal hedge portfolio would be an interesting avenue for further research, but in the absence of longer time series is likely subject to concerns of data mining.

[12] These scores were formerly known as KLD scores. In 2010, following MSCI's acquisition of RiskMetrics, KLD scores were retooled into what are now known as MSCI KLD scores.

subcategory is either scored as a "1" when the firm satisfies a certain condition, or a "0" if the firm does not satisfy the condition. For instance, a "1" in the positive *"Climate Change - Energy Efficiency"* subcategory means that the company operates in a relatively energy-efficient way. The thresholds for satisfying each condition are determined by MSCI and are not disclosed with the data. Following Hong and Kostovetsky (2012), we calculate an overall environmental score for each firm by subtracting the total scores in the negative environmental subcategories from the total scores in positive environmental subcategories. We call the resultant variable the "MSCI E-Score," where a higher score suggests a firm is more environmentally friendly. In principle, it would be possible to also construct E-Scores from only a selection of all "E" subcategories, perhaps by focusing on those subcategories that are particularly relevant for climate change. The out-of-sample performance of hedge portfolios constructed using different combinations of "E" subcategories could then be compared to select the one with the best performance. However, given the relatively short time series to evaluate the performance of the resultant hedge portfolios, even such an "out-of-sample" approach of finding the "best" E-Scores is naturally subject to data mining concerns. We hence decided to restrict ourselves to only analyzing the relatively broad overall E-Score, following prior approaches in the literature; we leave a more detailed exploration of the various subcategories to future research.

2.3.2 Sustainalytics. Sustainalytics provided us with monthly firm-level ESG scores beginning in September 2009. The broadest score in the data is the "Total ESG Score," which is the average of the "Total Environment Score," the "Total Social Score," and the "Total Governance Score." To determine each of the "E," "S," and "G" scores, Sustainalytics uses a number of subcategories and evaluates each firm's score by comparing it to peers in the same industry (Sustainalytics uses a nonstandard industry classification). For instance, the fifty-seven subcategories for the "Total Environment Score" include evaluations of a firm's efforts to reduce greenhouse gas emissions, increase renewable energy use, and reduce water use; Appendix A.3 presents the full list of subcategories. The scores in the subcategories are then aggregated by weighting them according to how exposed each industry is to each ESG risk, though this aggregation procedure is not well documented. Final scores are between 0 and 100. As before, a higher score suggests a firm is more environmentally friendly. We use the "Total Environment Score" in our empirical analysis.

2.3.3 Summary statistics. Our analysis of climate hedge portfolios focuses on the period between September 2009 and December 2016. This is a period for which we observe both measures of innovations of climate news, CC_t^{WSJ} and $CC_t^{NegNews}$, and both the Sustainalytics and MSCI E-Scores. We can therefore conduct a direct comparison of the performance of the various hedge portfolios for the two climate news series over this time horizon. For the MSCI E-Score,

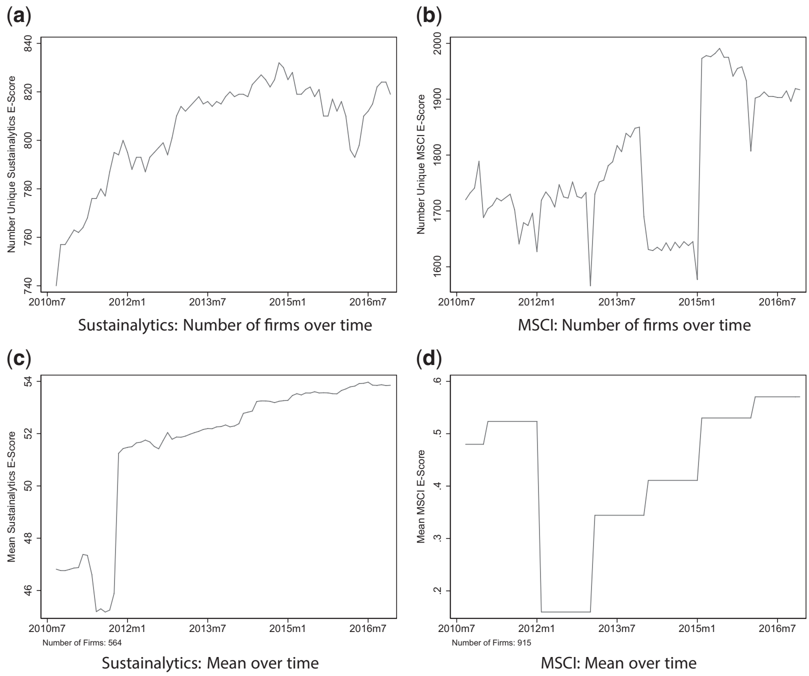

Figure 5
E-Scores: Summary statistics over time
This figure provides summary statistics for our two E-Scores. The top row shows the number of firms in our sample for which we observe E-Scores. The bottom row shows the average E-Score over time across those firms that we observe in every period in our sample. The left column shows these statistics for the Sustainalytics E-Score, and the right panel shows the statistics for the MSCI E-Score.

which is only reported annually, we assign the same score to all the months in the relevant year. Panels A and B of Figure 5 plot the number of firms in our pool of potential hedge assets for which we observe each E-Score over time. For Sustainalytics, we usually observe E-Scores for between 700 and 800 firms. MSCI E-Scores have broader coverage and are provided for between 1,700 and 1,900 firms.

Panels C and D of Figure 5 show the average values for each of the two E-Scores for a constant set of firms that we observe throughout the sample. The averages of each score contain a number of discontinuous breaks. For the MSCI E-Score, which is determined annually, these breaks could be either due to changes in firms' true ESG performance between years or due to changes in the modeling procedure. For Sustainalytics, which computes monthly scores, the discontinuous breaks are more likely due to changes in the modeling methodology over time, though we have been unable to obtain documentation on such changes that would allow us to verify this conjecture.[13] Such modeling changes would be problematic for building time-series models that perform well out of sample.

To minimize the complications from any modeling changes, we construct Z_t by cross-sectionally demeaning each E-Score in each month. However, this approach might still be problematic if changes to the model do not just shift the mean of the E-Scores over time, but also the cross-sectional dispersion. In that case, the meaning of absolute differences in the demeaned E-Score would change over time. As a second way to construct measures of Z_t, we therefore rank the E-Scores of all firms at each point in time, and then demean and rescale the ranked measure such that it ranges from -0.5 to +0.5. This approach preserves the ordinal content of the E-Scores but discards any information contained by the absolute differences between scores. Ranking-based approaches come with a number of issues. In particular, panels A and B of Figure 5 highlight that the number of firms for which E-Scores are available changes throughout the sample period. Firms added later in the sample are plausibly systematically different from those added earlier; for example, they might be less exposed to climate risk. The cross-sectional ranking of the same firm might therefore change over time without the true climate exposure of that firm changing. As a result, neither the demeaned absolute value nor the demeaned and rescaled ranked value of E-Scores are ex ante superior methods to construct climate exposures in Z_t. We will therefore present hedge portfolios using both approaches to constructing exposure measures and compare their relative performance.[14]

An interesting question is what firm characteristics are captured by the two E-Scores. A first hypothesis is that they primarily pick up industry-membership, whereby firms in "clean" industries, such as wind and solar energy, are assigned high E-Scores, and firms in "dirty" industries such as coal mining are assigned low E-Scores. To explore the extent to which the scores are primarily capturing a firm's industry, we begin by taking the firm-level E-scores in December 2016 (the last period in our data) and regressing them onto industry fixed effects. When regressing the absolute value of the Sustainalytics E-Score on 2-digit SIC code fixed effects, the adjusted R-squared of the regression is .103; it is .184 when regressing on fixed effects for 4-digit SIC codes. The measures of R-squared were similar when using the ranked measure of the Sustainalytics E-Score. When regressing the absolute value of the MSCI E-Score on 2-digit SIC codes (4-digit SIC codes), the adjusted R-squared of the regression is .099

[13] Most uses of ESG scores by the financial services sector build on the cross-section of ESG scores at a given point in time, for example, by forming portfolios that have a relatively higher performance on these measures. Such use cases often do not require a stable meaning of the same numerical score over time.

[14] The climate exposure measures in Z_t can be constructed from the various raw E-Scores in other ways. For example, one could cross-sectionally standardize each absolute measure to have a constant standard deviation over time. Alternatively, one could rank firms' E-Scores within industry rather than across all firms. However, in the absence of longer time series, a systematic analysis of which of these approaches obtains the best out-of-sample fit during our sample period is subject to the data mining concerns described earlier. As a result, we did not pursue these alternative approaches in this project.

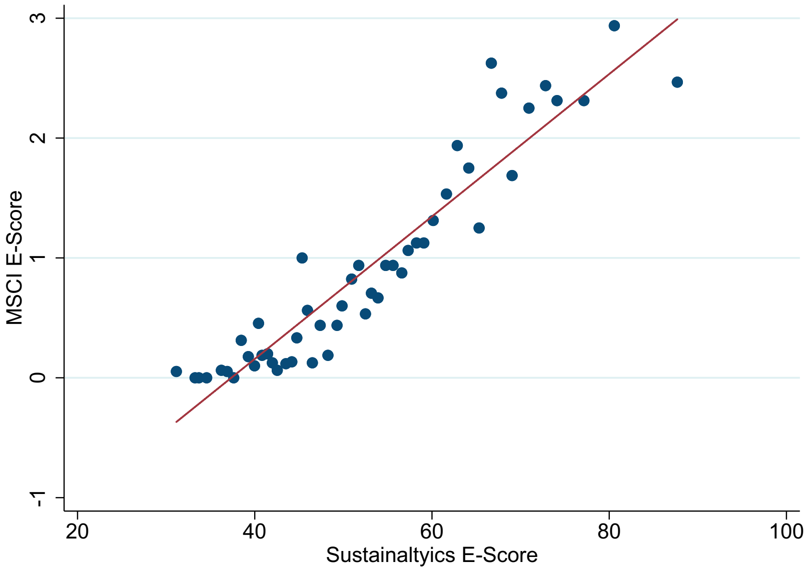

Figure 6
Correlation across E-Scores, December 2016
This figure shows a binned scatterplot that highlights the correlation across the Sustainalytics and MSCI E-Scores for all 796 firms in our sample that have both scores in December 2016. The correlation coefficient is 0.65.

(.203). These numbers show that, although there is some industry effect in determining E-Scores, most of the variation occurs within relatively narrow industries, rather than across industries.

Indeed, the three 2-digit SIC industries with the lowest Sustainalytics E-Scores are Personal Services (SIC code 72), Water Transportation (SIC code 44), and Motion Pictures (SIC code 78), probably not the first industries that come to mind when thinking of "dirty" industries. Similarly, the 2-digit SIC industries with the highest Sustainalytics E-Scores are Building Materials & Gardening Supplies (SIC code 52), Textile Mill Products (SIC code 22), and Furniture & Homefurnishings Stores (SIC code 57). When ranking by MSCI E-Scores, we similarly find that low-scoring firms are not necessarily those one would expect ex ante, such as those operating in the oil and gas sector.

A second question is the extent to which the MSCI and Sustainalytics E-Scores capture the same object. Figure 6 shows the correlation across the raw Sustainalytics and MSCI E-Scores in December 2016. They have a positive correlation of about 0.65, suggesting that they are both measuring aspects of the same object. However, enough independent variation occurs across the two measures to suggest that their usefulness in constructing climate hedge portfolios might vary. Indeed, we show below that the performance of the hedge portfolios varies noticeably when these hedge portfolios are constructed using the different E-Scores.

2.4 Forming hedge portfolios

In this section, we construct hedge portfolios for innovations in climate news, CC_t, using the mimicking portfolio approach described in Section 1.1. As discussed above, we use two different approaches to transform the raw E-Scores into the characteristic vector Z_t:

(1) Using firms' cross-sectionally demeaned absolute value of the E-Score ("absolute scores", e.g., $Z_t^{SUS\_A}$)

(2) Ranking the firms cross-sectionally by their E-Score, and then standardizing these rankings to range between -0.5 and +0.5 ("ranked scores", e.g., $Z_t^{SUS\_R}$).

Recall that one of the conditions for the mimicking portfolio approach to isolate climate change risk (and to avoid picking up other potentially correlated risks in the economy) is that the projection portfolios have to span all the risk factors driving returns. In addition to portfolios sorted on the climate characteristics, we therefore also include in regression 2 three additional factors that might be correlated with climate risk and that are known to be important in explaining the cross-section of returns: size (using cross-sectionally standardized market value to create Z_t, so that half the firms, sorted by market value, have positive weight, and half have negative weight; note that this portfolio will be long large firms and short small firms), value (using cross-sectionally standardized values of book-to-market to create Z_t), and the market (setting Z_t to equal the share of total market value).[15] For example, when we use the absolute Sustainalytics E-Score to measure firms' climate risk exposures, regression 3 becomes

$$CC_t = \xi + w_{SUS} Z_{t-1}^{SUS\_A'} r_t + w_{SIZE} Z_{t-1}^{SIZE'} r_t + w_{HML} Z_{t-1}^{HML'} r_t \\ + w_{MKT} Z_{t-1}^{MKT'} r_t + e_t, \tag{4}$$

where $w_{SUS}, w_{SIZE}, w_{HML}$ and w_{MKT} are scalars that capture the weight of the corresponding portfolios in the mimicking (hedge) portfolio for CC_t. For comparability, we also analyze the performance of hedge portfolios constructed using returns of the exchange-traded funds (ETFs) XLE and PBD instead of the returns of portfolios of stocks sorted by their E-Scores. XLE is the ticker of the Energy Select Sector SPDR ETF, which represents the energy sector of the S&P 500. PBD is the ticker of the Invesco Global Clean Energy ETF, which is based on the WilderHill New Energy Global Innovation Index and comprises companies that focus on greener and renewable sources of energy and technologies facilitating cleaner energy. Constructing hedge portfolios based on those ETFs allows us to (a) analyze the extent to which our E-Score-based hedge portfolios simply represent a market tilt away from "brown energy"

15 To maximize the number of stocks used to construct the hedge portfolios, we include stocks even if some of the characteristics Z_t are missing for that stock. To do so, we set all missing characteristics equal to zero.

Table 1
Full-sample regression: WSJ Climate Change News Index

	(1)	(2)	(3)	(4)	(5)
$Z_{t-1}^{SUS\_A'} r_t$	1.416***				
	(0.436)				
$Z_{t-1}^{SUS\_R'} r_t$		67.789***			
		(17.834)			
$Z_{t-1}^{MSCI\_A'} r_t$			12.658*		
			(6.849)		
$Z_{t-1}^{MSCI\_R'} r_t$				53.743*	
				(27.401)	
r_t^{XLE}					0.085
					(0.810)
r_t^{PBD}					0.208
					(0.630)
$Z_{t-1}^{HML'} r_t$	1.221	2.309	−5.862	−5.941	−6.772
	(7.019)	(6.873)	(6.878)	(6.858)	(8.093)
$Z_{t-1}^{SIZE'} r_t$	−5.680**	−6.034**	−5.511*	−5.459**	−2.765
	(2.350)	(2.289)	(2.773)	(2.696)	(2.474)
$Z_{t-1}^{MKT'} r_t$	0.783	0.789	0.841	0.789	0.091
	(0.642)	(0.628)	(0.692)	(0.680)	(1.285)
Constant	2.894	2.673	4.659*	4.891*	5.959**
	(2.681)	(2.613)	(2.700)	(2.669)	(2.897)
R-squared	.153	.187	.083	.088	.047
N	88	88	88	88	88

This table shows results from regression 4. The dependent variable captures innovations for the WSJ-Based Climate News measure. The unit of observation is a month, and the sample runs between September 2009 and December 2016. Standard errors are presented in parentheses. $^*p<.1$; $^{**}p<.05$; $^{***}p<.01$.

and toward "green energy" and (b) explore whether hedge portfolios based on XLE and PBD would have performed better than our E-Score-based hedge portfolios.[16]

2.5 In-sample fit results

We begin by exploring the in-sample fit of various versions of regression 4 over the full sample period. Table 1 shows regressions when hedging innovations to the WSJ Climate Change News Index, CC_t^{WSJ}, described in Section 2.1. Columns 1 and 2 show that portfolios based on Sustainalytics E-Scores have a positive and significant relationship with CC_t^{WSJ}; in periods with more innovations in negative climate news, a portfolio that goes long firms with higher (more "green") E-Scores has relatively larger excess returns. The R-squared measures of these regressions show that the portfolios based on the Sustainaltyics E-Scores can hedge 15%–19% of the in-sample variation in

16 As before, there are many degrees of freedom for how to compute hedge portfolios based on ETFs, and we do not want to suggest that portfolios constructed using XLE and PBD constitute the "best" ETF-based portfolios for hedging climate risk. Indeed, we view the analysis of which ETFs and other funds are most helpful in hedging climate risk to be an exciting area for future research.

CC_t. Columns 3 and 4 show that portfolios based on the MSCI E-Scores also have higher excess returns during periods with innovations in negative climate news; the R-squared measures of the regressions are lower than those in Columns 1 and 2. Portfolios based on ranked versions of both E-Scores have a slightly higher in-sample fit than portfolios based on absolute demeaned values. In addition to the ESG scores, size appears to correlate with climate change exposure: larger firms appear more exposed than smaller firms to climate change news, in the sense that they perform worse when the amount of news coverage of climate change in the WSJ increases. Column 5 includes the returns of XLE and PBD instead of the return of a characteristic-sorted portfolio. The in-sample fit of this regression is lower than that of any of the regressions in Columns 1–4, even though we have fewer explanatory variables in those regressions. This suggests that the characteristic-weighted portfolios might have some advantages over a hedge approach that creates industry tilts using energy-related ETFs.[17] It also shows that most of the R-squared in Columns 1–4 is the result of the characteristics-weighted portfolios, and not of the other portfolios, which are also included in Column 5.

Table 2 presents the same set of regressions as Table 1, but hedges innovations in the CH Negative Climate Change News Index, $CC_t^{NegNews}$. As before, the in-sample fits of the hedge portfolios based on Sustainalytics E-Scores are higher than the fits of the hedge portfolios based on MSCI E-Scores; similarly, the in-sample fits of the portfolios constructed using ranked E-Scores are marginally higher than those of the portfolios constructed using the absolute (demeaned) E-Score. Finally, the in-sample fits of all four portfolios based on E-Scores are somewhat higher than that of the portfolio based on XLE and PBD.[18] Overall, the relative performance of the various hedge portfolios is similar whether we are trying to hedge the WSJ Climate Change News Index or the CH Negative Climate Change News Index.

How would the hedge portfolios implied by these regressions look? To determine each firm i's weight in the hedge portfolio, we construct the following sum, where $Z_{i,t}$ values are taken as of December 2016: $\hat{w}_{SUS\_A} Z_{i,Dec16}^{SUS\_A'} + \hat{w}_{SIZE} Z_{i,Dec16}^{SIZE'} + \hat{w}_{HML} Z_{i,Dec16}^{HML'} + \hat{w}_{MKT} Z_{i,Dec16}^{MKT'}$, and where the various $\hat{w}$-terms represent the estimated coefficients from regression 4. This means that a firm's weight in the hedge portfolio is determined by its E-Score as well as its book-to-market ratio and its size. The resultant portfolio is the portfolio that an investor would form in December 2016 to hedge climate news in January 2017. Table 3 presents the average portfolio positions by 2-digit SIC code classification for

[17] The inclusion of the other factors in regression 4 make the resultant hedge portfolios in Column 5 of Table 2 different from a simple industry-tilt away from the market. Indeed, the resultant hedge portfolio will have a beta of 1 with CC_t, and a beta of zero with the other factors. Factor neutrality, not industry neutrality, is a desirable property of hedge portfolios.

[18] It is interesting to note that when hedging negative climate change news, the value-growth dimension seems to be aligned with the risk exposure. In particular, the table shows that value firms appear more exposed to climate news than growth firms.

Table 2
Full-sample regression: CH Negative Climate Change News Index

	(1)	(2)	(3)	(4)	(5)
$Z_{t-1}^{SUS\_A'} r_t$	0.266*				
	(0.141)				
$Z_{t-1}^{SUS\_R'} r_t$		12.286**			
		(5.864)			
$Z_{t-1}^{MSCI\_A'} r_t$			1.089		
			(2.173)		
$Z_{t-1}^{MSCI\_R'} r_t$				6.641	
				(8.696)	
r_t^{XLE}					−0.092
					(0.252)
r_t^{PBD}					0.036
					(0.196)
$Z_{t-1}^{HML'} r_t$	−4.536**	−4.390*	−5.934***	−5.919***	−5.520**
	(2.272)	(2.260)	(2.182)	(2.177)	(2.519)
$Z_{t-1}^{SIZE'} r_t$	−0.137	−0.179	0.210	0.100	0.501
	(0.761)	(0.753)	(0.880)	(0.856)	(0.770)
$Z_{t-1}^{MKT'} r_t$	0.315	0.314	0.287	0.295	0.297
	(0.208)	(0.206)	(0.219)	(0.216)	(0.400)
Constant	−0.115	−0.137	0.313	0.306	0.376
	(0.868)	(0.859)	(0.857)	(0.847)	(0.902)
R-squared	.125	.133	.090	.094	.089
N	88	88	88	88	88

This table shows results from regression 4. The dependent variable captures innovations for the *Newspaper-based negative climate news* measure. The unit of observation is a month, and the sample runs between September 2009 and December 2016. Standard errors are presented in parentheses. $^{*}p<.1$; $^{**}p<.05$; $^{***}p<.01$.

the industries with the six largest negative average portfolio weights and the industries with the six largest positive average portfolio weights. We only present the portfolio positions based on the absolute E-Scores, because they look very similar to the positions in the hedge portfolio constructed using the ranked E-Scores. For the portfolio constructed using Sustainlytics E-Scores to hedge innovations in the CH Negative Climate Change News Index, for example, the largest short position is “General Building Contractors,” followed by “Water Transportation.” The largest long positions are “Building Materials & Gardening Supplies” and “Tobacco Products.” This analysis highlights that the resultant hedge portfolios will not necessarily conform with common priors that the optimal way to hedge climate change news involves primarily going long green energy stocks and short oil companies; this is consistent with our observation that industry membership can only explain a small amount of the cross-sectional variation in firm-level E-Scores.

2.6 Out-of-sample fit results

The most important test of the hedge portfolios is their ability to hedge out-of-sample innovations to climate news, that is, to hedge innovations in months that were not included in the estimation of the portfolio weights. To construct a first

Table 3
Largest average short and long positions (by 2-digit SIC code)

A. WSJ Climate Change News Index

Sustainalytics E-Score (absolute)		**MSCI E-Score (absolute)**	
Top negative portfolio weights	***SIC2***	***Top negative portfolio weights***	***SIC2***
Coal mining	12	Water transportation	44
Water transportation	44	Petroleum & coal products	29
Insurance agents, brokers, & service	64	Motion pictures	78
Mining nonmetallic minerals, except fuels	14	Communications	48
Transportation services	47	Security & commodity brokers	62
Security & commodity brokers	62	Oil & gas extraction	13
Top positive portfolio weights	***SIC2***	***Top positive portfolio weights***	***SIC2***
Building materials & gardening supplies	52	Pipelines, except natural gas	46
Tabacco products	21	Tabacco products	21
Food & kindred products	20	Miscellaneous manufacturing industries	39
Paper & allied products	26	Lumber & wood products	24
Textile mill products	22	Paper & allied products	26
Furniture & homefurnishings stores	57	Textile mill products	22

B. CH Negative Climate Change News Index

Top negative portfolio weights	***SIC2***	***Top negative portfolio weights***	***SIC2***
General building contractors	15	General building contractors	15
Water transportation	44	Nondepository institutions	61
Coal mining	12	Auto repair, services, & parking	75
Insurance agents, brokers, & service	64	Communications	48
Holding and other investment offices	67	Water transportation	44
Insurance carriers	63	Insurance carriers	63
Top positive portfolio weights	***SIC2***	***Top positive portfolio weights***	***SIC2***
Railroad transportation	40	Chemical & allied products	28
Transportation by air	45	Textile mill products	22
Furniture & homefurnishings stores	57	General merchandise stores	53
Textile mill products	22	Lumber & wood products	24
Building materials & gardening supplies	52	Building materials & gardening supplies	52
Tobacco products	21	Tobacco products	21

This table shows the industries (2-digit SIC code) with the largest average short and long positions in the estimated hedge portfolios resulting from regressions presented in Tables 1 and 2. Panel A explores hedge portfolios based on regression 4 using innovations in the WSJ Climate Change News Index as CC_t, and panel B explores hedge portfolios based using innovations in the CH Negative Climate Change News Index as CC_t. All portfolios are constructed using the absolute demeaned value of the E-Scores. Within each portfolio, industries are arranged in descending order of the absolute values of the portfolio weights.

measure of the out-of-sample performance of the hedge portfolios, for every period t we run regression 4 using data between periods t_{min} and $t-1$, where t_{min} corresponds to the first month for which we observe all climate exposures and CC_t series (September 2009). We then form the hedge portfolio based on these estimates and explore the correlation of the returns of that hedge portfolio in period t with CC_t. This corresponds to the approach one would have taken to hedge climate news in real time. Because we require a certain amount of data to estimate regression 4, we only compare the out-of-sample performance of the hedge portfolios starting in period $t_{min}+30$ (March 2012).[19]

[19] Further reducing the number of portfolios onto which to project CC_t may improve the out-of-sample performance of the hedging portfolio. Given the short sample size available, in this paper we decided to not optimize the hedge portfolio further along this dimension.

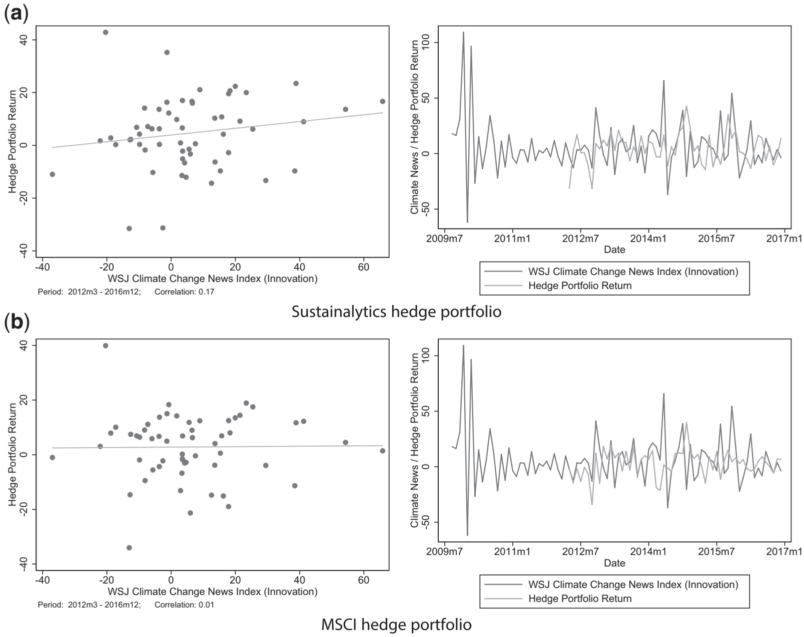

Figure 7
Out-of-sample fit: WSJ Climate Change News Index
This figure explores the out-of-sample performance of hedge portfolios constructed to hedge the *WSJ-Based Climate News Measure*. The top panel presents hedge portfolios built on the absolute values of the Sustainalytics E-Score, and the bottom panel presents portfolios built on the absolute values of the MSCI E-Score.

Figure 7 presents the out-of-sample performance of portfolios constructed to hedge innovations in the WSJ Climate Change News Index. The top panels show portfolios constructed using absolute values of the Sustainalytics E-Score, and the bottom panels show portfolios that build on the absolute values of the MSCI E-Score. The left columns present scatterplots of the out-of-sample returns of the hedge portfolios together with the realizations of the innovation of climate news. The right panels plot the time series of the climate news series and the return series of the hedge portfolios. There is a clear, positive out-of-sample correlation with CC_t of 0.17 for the Sustainalytics hedge portfolio. In other words, the hedge portfolios indeed have higher returns during periods with positive innovations to climate news. Portfolios based on MSCI E-Scores or ETFs, on the other hand, have very little ability to hedge innovations in the WSJ Climate Change News Index, with an out-of-sample correlation of just 0.01.

Panel A of Table 4 provides additional information about the out-of-sample performance of the various portfolios designed to hedge innovations in the WSJ Climate Change News Index. The first column is the most important one, showing the correlation between the realizations of CC_t^{WSJ} and the

Table 4
Cross-correlations: WSJ Climate Change News Index

A. Out-of-sample fit

	CC^{WSJ}	$H_{OOS}^{SUS\_A}$	$H_{OOS}^{SUS\_R}$	$H_{OOS}^{MSCI\_A}$	$H_{OOS}^{MSCI\_R}$	H_{OOS}^{ETF}	r_t^{XLE}	r_t^{PBD}
CC^{WSJ}	1.000							
$H_{OOS}^{SUS\_A}$	0.174	1.000						
$H_{OOS}^{SUS\_R}$	0.206	0.973	1.000					
$H_{OOS}^{MSCI\_A}$	0.013	0.688	0.621	1.000				
$H_{OOS}^{MSCI\_R}$	0.019	0.677	0.624	0.988	1.000			
H_{OOS}^{ETF}	−0.005	0.427	0.349	0.861	0.852	1.000		
r_t^{XLE}	0.068	−0.138	0.004	−0.097	−0.039	−0.141	1.000	
r_t^{PBD}	0.111	0.185	0.272	0.294	0.350	0.190	0.656	1.000

B. Cross-validation fit

	CC^{WSJ}	$H_{Cross}^{SUS\_A}$	$H_{Cross}^{SUS\_R}$	$H_{Cross}^{MSCI\_A}$	$H_{Cross}^{MSCI\_R}$	H_{Cross}^{ETF}	r_t^{XLE}	r_t^{PBD}
CC^{WSJ}	1.000							
$H_{Cross}^{SUS\_A}$	0.244	1.000						
$H_{Cross}^{SUS\_R}$	0.300	0.976	1.000					
$H_{Cross}^{MSCI\_A}$	0.039	0.742	0.671	1.000				
$H_{Cross}^{MSCI\_R}$	0.067	0.733	0.676	0.982	1.000			
H_{Cross}^{ETF}	−0.069	0.454	0.390	0.678	0.651	1.000		
r_t^{XLE}	0.068	0.041	0.072	−0.009	−0.034	0.297	1.000	
r_t^{PBD}	0.111	0.272	0.266	0.310	0.298	0.470	0.656	1.000

This table shows cross-correlations of different portfolios and innovations in the WSJ Climate Change News Index. Panel A focuses on the performance of hedge portfolios from our out-of-sample approach, and panel B focuses on the performance of hedge portfolios from our cross-validation approach.

returns of the various hedge portfolios (e.g., $R_{OOS}^{SUS\_A}$ corresponds to the out-of-sample returns of a hedge portfolio constructed using absolute values of the Sustainalytics E-Score). The hedge portfolios based on Sustainalytics E-Scores substantially outperform the hedge portfolios based on the MSCI E-Scores. In addition, hedge portfolios based on ranked E-Scores marginally outperform those based on absolute E-Scores, though the returns of portfolios based on absolute and ranked E-Scores from the same data provider are highly correlated. Finally, the out-of-sample performance of the Sustainalytics E-Score-based hedge portfolios is substantially better than that of portfolios based on ETFs. The returns of most hedge portfolios are *negatively* correlated with the returns to XLE, suggesting that these hedge portfolios are likely to hold short positions in the energy firms that constitute XLE. Similarly, we observe a positive correlation between the returns of all climate hedge portfolios and the returns of PBD, suggesting that the hedge portfolios likely hold long positions in many of the green energy firms that constitute PBD.

We also conduct a second test for the performance of the hedge portfolios based on a cross-validation approach. In particular, for every period t' we run regression 4 for all periods $t \neq t'$, and then use the resultant estimates to construct a hedge portfolio in a similar way as described above. The return of that hedge portfolio in period t' is then compared to $CC_{t'}$. Panel B of Table 5 explores the cross-validation performance of the various hedge portfolios. The hedge portfolios based on Sustainalytics E-Scores continue to outperform those based on MSCI E-Scores or ETFs substantially.

In sum, the hedge portfolios built using the Sustainalytics E-Score perform out of sample substantially better than any other hedge portfolio we have considered. The worse hedging performance of portfolios based on MSCI E-Scores highlights the importance of choosing characteristics that properly capture cross-sectional variation in exposure to climate change risks.

Figure 8 and Table 5 present results similar to those in Figure 7 and Table 4, but analyze the performance of portfolios designed to hedge innovations in the CH Negative Climate Change News Index. Portfolios based on Sustainalytics E-Scores have a similar ability to hedge this second climate news series as they had in hedging the CH Negative Climate Change News Index, both in the out-of-sample evaluation and in the cross-validation evaluation. The hedging ability of the MSCI indexes is in this case much higher than for the WSJ measure of climate change risks, suggesting that the MSCI E-Scores are more suited to capture negative climate change news as opposed to general coverage of climate change by the WSJ. Overall, the out-of-sample correlation between realization of climate change news and the hedge portfolios are 0.22 when using Sustainalytics E-Scores and 0.18 when using MSCI E-Scores.

3. Conclusion and Directions for Future Research

We demonstrate how a mimicking portfolio approach can be successful in hedging innovations in climate change news across a number of out-of-sample performance tests. Across our two indices for climate news, the hedge portfolios based on Sustainalytics E-Scores have the best in-sample fit as well as the best out-of-sample and cross-validation performance. Portfolios based on MSCI E-Scores and ETFs have a lower (but still positive) ability to hedge innovations in climate news. There are no systematic differences in the relative performance of hedge portfolios based on absolute or ranked versions of the raw E-Scores. In general, however, the differences between the out-of-sample and cross-validation performance of some of the portfolios highlight that the portfolios we construct are somewhat sensitive to the exact time series on which our models are trained. This is likely the result of only having a relatively few data points in each of our estimations. As we observe longer time series of E-Scores and climate news measures, our proposed method should deliver ever-better portfolios to hedge climate change news. Similarly, moving from hedging climate news that materializes over a monthly level to hedging on a daily level should allow researchers to substantially expand their training data, and thereby improve the out-of-sample performance of the hedge portfolios.

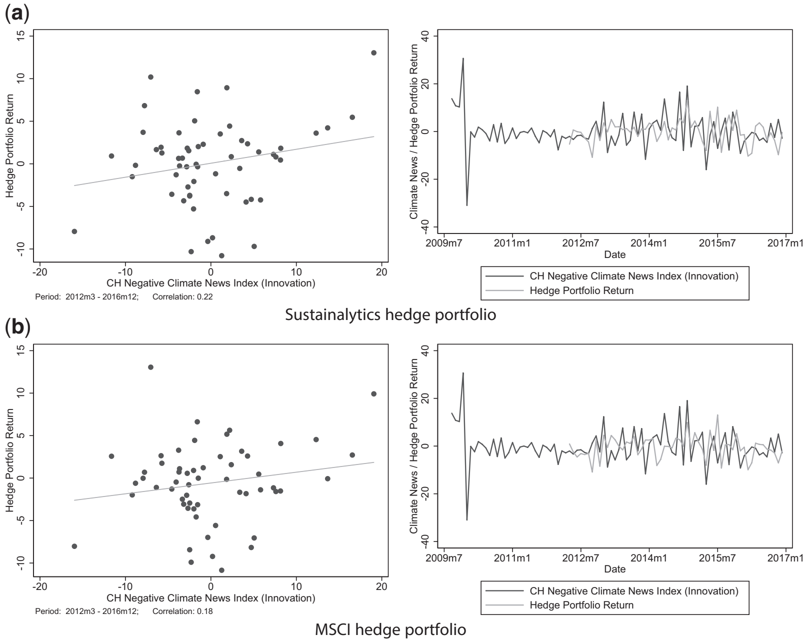

Figure 8
Out-of-sample fit: CH Negative Climate Change News Index
This figure explores the out-of-sample performance of hedge portfolios constructed to hedge CH-based negative climate news measure. The top panel presents hedge portfolios built on the absolute values of the Sustainalytics E-Score, and the bottom panel presents portfolios built on the absolute values of the MSCI E-Score.

More generally, we view this article as providing a rigorous methodology for constructing portfolios that hedge against risks that are otherwise difficult to insure. We do not view our resultant hedge portfolios as the definitive best hedges against climate change risk, but instead as a starting point for further exploration. Indeed, future research could consider many valuable directions for climate finance, and we discussed a number of the dimensions that should be explored further, including the addition of more assets to the hedge portfolios (such as international stocks) and the formation of hedge portfolios based on both characteristic-sorted portfolios and ETFs.

One additional important direction for future work is to integrate more and better data to measure firm-level climate risk exposures. These data could come from commercial data providers or could be constructed by researchers themselves, for example, by including information such as geographical proximity to potential climate disasters (e.g., rising sea levels or hurricane-prone regions). Indeed, articles in this volume, such as Choi, Gao, and Jiang (2018) and Kumar, Shashwat, and Wermers (2018) make valuable progress toward developing new ways to quantify climate risk exposures.

Table 5
Cross-correlations: CH Negative Climate Change News Index

A. Out-of-sample fit

	$CC^{NegNews}$	$H_{OOS}^{SUS\_A}$	$H_{OOS}^{SUS\_R}$	$H_{OOS}^{MSCI\_A}$	$H_{OOS}^{MSCI\_R}$	H_{OOS}^{ETF}	r_t^{XLE}	r_t^{PBD}
$CC^{NegNews}$	1.000							
$H_{OOS}^{SUS\_A}$	0.217	1.000						
$H_{OOS}^{SUS\_R}$	0.183	0.992	1.000					
$H_{OOS}^{MSCI\_A}$	0.179	0.869	0.852	1.000				
$H_{OOS}^{MSCI\_R}$	0.175	0.865	0.850	0.998	1.000			
H_{OOS}^{ETF}	0.157	0.780	0.767	0.961	0.960	1.000		
r_t^{XLE}	−0.066	−0.412	−0.353	−0.387	−0.367	−0.410	1.000	
r_t^{PBD}	0.063	0.061	0.112	0.096	0.127	0.119	0.656	1.000

B. Cross-validation fit

	$CC^{NegNews}$	$H_{Cross}^{SUS\_A}$	$H_{Cross}^{SUS\_R}$	$H_{Cross}^{MSCI\_A}$	$H_{Cross}^{MSCI\_R}$	H_{Cross}^{ETF}	r_t^{XLE}	r_t^{PBD}
$CC^{NegNews}$	1.000							
$H_{Cross}^{SUS\_A}$	0.148	1.000						
$H_{Cross}^{SUS\_R}$	0.154	0.991	1.000					
$H_{Cross}^{MSCI\_A}$	0.024	0.864	0.836	1.000				
$H_{Cross}^{MSCI\_R}$	0.048	0.885	0.861	0.993	1.000			
H_{Cross}^{ETF}	0.053	0.829	0.799	0.973	0.968	1.000		
r_t^{XLE}	−0.066	−0.208	−0.183	−0.205	−0.237	−0.223	1.000	
r_t^{PBD}	0.063	0.169	0.171	0.158	0.157	0.185	0.656	1.000

This table shows cross-correlations of different portfolios and innovations in the CH Negative Climate Change News Index. Panel A focuses on the performance of hedge portfolios from our out-of-sample approach, and panel B focuses on the performance of hedge portfolios from our cross-validation approach.

Another direction for follow-on work is to develop alternative definitions of the climate change risks. One interesting question is whether it is important to differentiate between physical and policy-oriented climate risks. For example, a tax on greenhouse gas emissions, if comprehensively applied at an appropriate level, would reduce the demand for climate hedge portfolios and consequently the cost of insuring against climate change. Thus, good regulation will mean less need for climate hedges. But regulation itself creates winners and losers from regulatory risk, and one might therefore want to construct regulatory hedge portfolios. The stability of such regulatory hedge portfolios may well be sensitive to the prevailing political environment.

A related question pertains to the expected returns of the various hedge portfolios. Indeed, an increasing use of climate hedge portfolios by investors will increase the price (and thus reduce the expected returns) of those firms whose stock provides the most effective hedge against innovations in climate change news. This lower expected return corresponds to the insurance premium paid for the climate hedge portfolio. An interesting avenue for future work will be to quantify the cost of the climate hedge portfolios by looking at the associated risk premiums.[20] It is also interesting to study the general equilibrium effects resulting from the fact that a lower cost of capital for firms with high E-Scores might actually have a direct effect on the climate trajectory. For example, to the extent that green energy firms see a reduction in their cost of capital, this might allow them to achieve efficient scale faster, and thereby affect the path of greenhouse gas emissions. The design of structural asset pricing models that feature such general equilibrium feedback loops seems a promising direction for research.

Appendix

A.1 Review of the Fama-MacBeth Approach

In this section, we review the Fama-MacBeth estimator for hedge portfolios in the context of our model. To apply the Fama-MacBeth procedure, the econometrician needs to take a stand on all the factors in the model: CC_t and v_t. Once the factors in the model are determined, the procedure follows two steps. In the first step, the risk exposures β_{CC} and β are estimated via time-series regressions of returns onto the factors, CC_t and v_t. In particular, for each asset i, $(\hat{\beta}^i_{CC}, \hat{\beta}^i)$ are estimated from the time-series regression:

$$r^i_t = \alpha^i + \beta^i_{CC} CC_t + \beta^i v_t + u_t.$$

In the second step, in each period t, hedge portfolios for all factors are obtained via cross-sectional regressions of returns r_t onto the estimated betas $(\hat{\beta}_{CC}, \hat{\beta})$:

$$r_t = h^{CC}_t \hat{\beta}_{CC} + h_t \hat{\beta} + e_t,$$

where $\hat{\beta}_{CC}$ and $\hat{\beta}$ are the betas estimated in the first step. The slopes of this regression in each period t are precisely the returns of the hedge portfolio in period t: h^{CC}_t (that hedges CC_t) and h_t (that hedges the remaining factors v_t). The hedge portfolios h^{CC}_t and h_t have, by construction, a beta of one with respect to the corresponding factors and zero with respect to all other factors. Their time-series means (the expected excess returns of the hedge portfolios) recover the risk premiums of the factors: $E[h^{CC}_t] = \gamma_{CC}$ and $E[h_t] = \gamma$. The Fama-MacBeth procedure for constructing hedge portfolios has two potential drawbacks. First, it requires knowing all the factors in the model, CC_t and v_t. Second, the procedure is not robust to measurement error in the factor of interest, CC_t, which is a natural concern in many settings, including in ours (see the further discussion of omitted factors and measurement error in Giglio and Xiu 2018).

A.2 Source of Climate Change Vocabulary

To create the Climate Change Vocabulary (CCV), we collect twelve climate change white papers from various sources including the Intergovernmental Panel on Climate Change (IPCC), Environmental Protection Agency (EPA), and the U.S. Global Change Research Program. We complement this with fifty-nine climate change glossaries from sources such as United Nations, NASA, IPCC, and EPA.

A.2.1 Twelve climate change white papers Table A1 reports the institution, title and published year of climate change white papers that we use to construct the CCV.

[20] Note that this requires substantial time-series data, because realizations of negative climate news in sample might actually lead the hedge portfolios to outperform over any given period.

Table A1
List of climate change white papers

Source	Title	Year
IPCC	IPCC Synthesis Report	1990, 1995, 2001, 2007, 2014
IPCC	IPCC Special Report: The Regional Impacts of Climate Change: an assessment of vulnerability	1997
IPCC	IPCC Special Report: Aviation and the Global Atmosphere	1999
IPCC	IPCC Special Report: Methodological and Technological Issues in Technology Transfer	2000
IPCC	IPCC Special Report: Safeguarding the Ozone Layer and the Global Climate System: Issues Related to Hydrofluorocarbons and Perfluorocarbons	2005
IPCC	IPCC Special Report: Carbon Dioxide Capture and Storage	2005
IPCC	IPCC Special Report: Renewable Energy Sources and Climate Change Mitigation	2011
IPCC	IPCC Special Report: Managing the Risks of Extreme Events and Disasters to Advance Climate Change Adaptation	2012
American Association for the Advancement of Science	What We Know: The Reality, Risks, and Response to Climate Change	2014
UC Berkley	American Climate Prospectus	2015
U.S. EPA	Climate Change Indicators in the United States (4th edition)	2016
Science	Social and Economic Impacts of Climate	2016
IMF	The Effects of Weather Shocks on Economic Activity	2017
U.S. Global Change Research Program	Our Change Planet: The U.S. Global Change Research Program for Fiscal Year 2017	2017
U.S. Global Change Research Program	Climate Science Special Report (4th National Climate Assessment, Vol. I)	2017

IPCC reports scientific and technical assessments of the current state of climate change. Generally, these reports comprise three volumes: one for each of the Working Groups of the IPCC. In addition to the main reports, Summaries for Policymakers and Synthesis Reports are provided. A Synthesis Report integrates materials covered by Assessment Reports and Special Reports. It is a nontechnical report targeting policy makers and addressing a broad range of policy-relevant but policy-neutral questions. Summary for Policymakers is an abridged version of the full Synthesis Report. In addition, IPCC Special Reports provide an assessment of a specific issue relating to climate change. They are generally structured similar to a volume of an Assessment Report. IPCC, Intergovernmental Panel on Climate Change; EPA, Environmental Protection Agency; IMF, International Monetary Fund.

A.2.2 Fifty-nine climate change glossaries

We collect climate change glossaries, both words and their definition, from U.S. Environmental Protection Agency (EPA), BBC, United Nations(UN), Center for Climate and Energy Solutions Glossary of Key Terms, Intergovernmental Panel on Climate Change (IPCC), World Health Organization (WHO), European Climate Adaptation Platform, International Petroleum Industry Environmental Conservation Association(IPIECA), Lenntech, Wikipedia, Met Office, Integrated Regional Information Networks(IRIN), Climate Change in Australia, Guardian, International Rivers, Mekong River Commission, Exploratorium, *New York Times*, U.S. Forest Service, U.S. Department of Transportation, Durham Region, Classroom of the Future, Government of Canada, International Food Policy Research Institute (IFPRI), New Zealand Government, University of Miami, German Climate Finance, California Government, South West Climate Change Impacts Partnership (SWCCIP), Scent of Pine, Natural Climate Change, UN Climate Change Conference, Center for Strategic and International Studies(CSIS), Watts Up With That?, U.K. Climate Impacts Programme (UKCIP), Climate Change Zambia, Canadian Broadcasting Corporation(CBC), Auburn University, Global Warming Solved, REDD+, Climate Resilience Toolkit(CRT), What's Your Impact, The Nitric Acid Climate Action Group (NACAG), Garnaut Climate Change Review, Climate Policy Information Hub, Explaining Climate Change, Four Degrees Preparation, The European Initiative for Upscaling Energy Efficiency in the Music Event Industry (EE MUSIC), Regional Education and Information Centre (REIC), Ecology, Climate Reality Project, National Geographic, Agricultural Marketing Resource Center (AgMRC), Global Greenhouse Warming, and Conservation in a Changing Climate.

A.3 Subcategories for "E" scores

A.3.1 MSCI

Positive indicators are Environmental Opportunities - Clean Tech, Waste Management - Toxic Emissions and Waste, Waste Management - Packaging Materials & Waste, Climate Change - Carbon Emissions, Property/Plant/Equipment, Environmental Management Systems, Natural Resource Use - Water Stress, Natural Resource Use - Biodiversity & Land Use, Natural Resource Use - Raw Material Sourcing, Natural Resource Use - Financing Environmental, Environmental Opportunities - Green Buildings, Environmental Opportunities in Renewable Energy, Waste Management - Electronic Waste, Climate Change - Energy Efficiency, Climate Change - Product Carbon Footprint, Climate Change - Insuring Climate Change Risk, Environment - Other Strengths. Negative indicators are Regulatory Compliance, Toxic Emissions and Waste, Energy & Climate Change, Impact of Products and Services, Biodiversity & Land Use, Operational Waste, Water Stress, Environment - Other Concerns.

A.3.2 Sustainalytics

Subcategories are Formal Environmental Policy, Environmental Management System, External Certification of Environmental Management Systems (EMS), Environmental Fines and Non-monetary Sanctions, Participation in Carbon Disclosure Project, Scope of Corporate Reporting on GHG emissions, Programmes and Targets to Reduce GHG Emissions from Own Operations, Programmes and Targets to Increase Renewable Energy Use, Carbon Intensity, Carbon Intensity Trend, % of Primary Energy Use from Renewables, Operations Related Controversies or Incidents, Reporting Quality Non-Carbon Environmental Data, Programmes and Targets to Protect Biodiversity, Guidelines and Reporting on Closure and Rehabilitation of Sites, Environmental and Social Impact Assessments, Oil Spill Reporting and Performance, Waste Intensity, Water Intensity, Percentage of Certified Forests Under Own Management, Programmes & Targets to Reduce Air Emissions, Programmes & Targets to Reduce Air Emissions, Programmes & Targets to Reduce Water Use, Other Programmes to Reduce Key Environmental Impacts, GHG Reduction Programme, Programmes and Targets to Improve the Environmental Performance of Own Logistics and Vehicle Fleets, Programmes and Targets to Phase out CFCs and HCFCs[21]

[21] CFCs refers to chlorofluorocarbons, and HCFCs refers to Hydrochloroflourocarbons.

in Refrigeration Equipment, Formal Policy or Programme on Green Procurement, Environmental Supply Chain Incidents, Programmes to Improve the Environmental Performance of Suppliers, External Environmental Certification Suppliers, Programmes and Targets to Stimulate Sustainable Agriculture, Programmes and Targets to Stimulate Sustainable Aquaculture/Fisheries, Food Beverage & Tabacco Industry Initiatives, Programmes and Targets to Reduce GHG Emissions from Outsourced Logistics Services, Data on Percentage of Recycled/Reused Raw Material Used, Data on Percentage of Forest Stewardship Council (FSC) Certified Wood/Pulp as Raw Material, Programmes and Targets to Promote Sustainable Food Products, Food Retail Initiatives, Products & Services Related to Controversies or Incidents, Sustainability Related Products & Services, Revenue from Clean Technology or Climate Friendly Products, Automobile Fleet Average CO2 Emissions, Trend Automobile Fleet Average Fleet Efficiency, Products to Improve Sustainability of Transport Vehicles, Systematic Integration of Environmental Considerations at R&D Stage, Programmes and Targets for End-of-Life Product Management, Organic Products, Policy on Use of Genetically Modified Organisms (GMO) in Products, Environmental & Social Standards in Credit and Loan Business, Responsible Asset Management, Use of Life-Cycle Analysis(LCA) for New Real Estate Projects, Programmes and Targets to Increase Investment in Sustainable Buildings, Share of Property Portfolio Invested in Sustainable Buildings, Sustainability Related Financial Services, Products with Important Environmental/Human Health Concerns, Carbon Intensity of Energy Mix, Mineral Waste Management, Emergency Response Programme.

References

Amihud, Y. 2002. Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets* 5:31–56.

Andersson, M., P. Bolton, and F. Samama. 2016. Hedging climate risk. *Financial Analysts Journal* 72:13–32.

Bakkensen, L., and L. Barrage. 2018. Flood risk belief heterogeneity and coastal home price dynamics: Going under water? Working Paper.

Baldauf, M., L. Garlappi, and C. Yannelis. 2018. Does climate change affect real estate prices? only if you believe in it. Working Paper.

Balduzzi, P., and C. Robotti. 2008. Mimicking portfolios, economic risk premia, and tests of multi-beta models. *Journal of Business and Economic Statistics* 26:354–68.

Bernstein, A., M. Gustafson, and R. Lewis. 2019. Disaster on the horizon: The price effect of sea level rise. *Journal of Financial Economics* Advance Access published March 23, 2019, 10.1016/j.jfineco.2019.03.013.

Black, F., and M. Scholes. 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81:637–54.

Breeden, D., M. Gibbons, and R. Litzenberger. 1989. Empirical tests of the consumption-oriented capm. *Journal of Finance* 44:231–62.

Chen, N.-F., R. Roll, and S. Ross. 1986. Economic forces and the stock market. *Journal of Business* 59:383–403.

Choi, D., Z. Gao, and W. Jiang. 2018. Attention to global warming. Working Paper.

Daniel, K., R. Litterman, and G. Wagner. 2015. Applying asset pricing theory to calibrate the price of climate risk. Working Paper, Columbia University.

Fama, E., and K. French. 2008. Dissecting anomalies. *Journal of Finance* 63:1653–78.

Fama, F., and J. MacBeth. 1973. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy* 81:607–36.

Gentzkow, M., B. T. Kelly, and M. Taddy. 2018. Text as data. Working Paper.

Giglio, S., M. Maggiori, and J. Stroebel. 2015. Very long-run discount rates. *Quarterly Journal of Economics* 130:1–53.

Giglio, S., M. Maggiori, J. Stroebel, and A. Weber. 2018. Climate change and long-run discount rates: Evidence from real estate. Working Paper.

Giglio, S., and D. Xiu. 2018. Asset pricing with omitted factors. Working Paper.

Hong, H., and L. Kostovetsky. 2012. Red and blue investing: Values and finance. *Journal of Financial Economics* 103:1–19.

Hong, H., F. Li, and J. Xu. 2019. Climate risks and market efficiency. *Journal of Econometrics* 208:265–81.

Hou, K., C. Xue, and L. Zhang. 2015. Dissecting anomalies: An investment approach. *Review of Financial Studies* 28:650–705.

Huberman, G., S. Kandel, and R. Stambaugh. 1987. Mimicking portfolios and exact arbitrage pricing. *Journal of Finance* 42:1–9.

Kumar, N., A. Shashwat, and R. Wermers. 2018. Do fund managers misestimate climatic disaster risk? Working Paper.

Lamont, O. 2001. Economic tracking portfolio. *Journal of Econometrics* 105:161–84.

Lönn, R., and P. Schotman. 2017. Empirical asset pricing with many assets and short time series. Working Paper, Maastricht University.

Merton, R. 1973. An intertemporal capital asset pricing model. *Econometrica* 41:867–87.

Murfin, J., and M. Spiegel. 2018. Is the risk of sea level rise capitalized in residential real estate. Working Paper.

Roll, R., and A. Srivastava. 2018. Mimicking portfolios. Working Paper, California Institute of Technology.