A quantity-based approach to constructing climate risk hedge portfolios*

Georgij Alekseev[†] Stefano Giglio[‡] Quinn Maingi[§]
Julia Selgrad[¶] Johannes Stroebel[‖]

Abstract

We propose a new methodology to build portfolios that hedge climate change risks. Our *quantity-based* approach explores how mutual funds holdings change when the fund adviser experiences a local extreme heat event that shifts beliefs about climate risks. We use the observed trading behavior to predict how investors will reallocate their capital when "global" climate news shocks occur, which shift the beliefs and asset demands of many investors simultaneously and thus move equilibrium prices. We show that a portfolio that holds stocks that investors tend to buy after experiencing a local heat shock appreciates in value in periods with aggregate climate news shocks. Our quantity-based approach yields superior out-of-sample hedging performance compared to traditional methods of identifying hedge portfolios. The key advantage of the quantity-based approach is that it learns from cross-sectional trading responses rather than time-series price information, which is limited in the case of climate risks. We also demonstrate the efficacy and versatility of the quantity-based approach by constructing successful hedge portfolios for aggregate unemployment and house price risk.

*This version: October 8, 2021. We thank Richard Berner, Robert Engle and seminar and conference participants at NYU Stern, the New York Fed, the Federal Reserve Board, the Central Bank of Chile, and IWH Halle for comments. We thankfully acknowledge financial support from the Norges Bank through a grant to the Volatility and Risk Institute at NYU Stern.

†NYU Stern. Email: galeksee@stern.nyu.edu.

‡Yale University, NBER, and CEPR. Email: stefano.giglio@yale.edu.

§NYU Stern. Email: rmaingi@stern.nyu.edu.

¶NYU Stern. Email: jselgrad@stern.nyu.edu.

‖NYU Stern, NBER, and CEPR. Email: johannes.stroebel@nyu.edu.

Climate change presents a major challenge to humanity: in addition to a wide range of social implications, both the physical effects of climate change and the regulatory efforts to slow carbon emissions have the potential to substantially disrupt economic activity. As investor awareness of the economic and financial risks of climate change has increased, there has been rising demand for financial instruments that hedge these risks (see Krueger et al. 2020, Giglio, Kelly & Stroebel 2021, Stroebel & Wurgler 2021). Given the scarcity of financial assets that directly transfer climate risks, investors and researchers have started to build hedging portfolios using existing assets such as equities, which change in value based on the realizations of climate risks (e.g., Engle et al. 2020). While existing procedures to construct such hedging portfolios—such as mimicking portfolio approaches along the lines of Lamont (2001)—are theoretically appealing, they are difficult to implement because of the scarcity of time-series data to learn the correlation of different assets' returns with the realizations of climate risks.

In this paper, we propose a novel methodology to build portfolios that hedge climate risks. Our *quantity-based* approach exploits the availability of *cross-sectional* information on investor trading responses to geographically localized shocks that change investor attention or beliefs about climate risks. Specifically, we study how mutual funds adjust their portfolios after a regional heat shock occurs in their adviser's location. We use this information to predict the aggregate response of investors when "global" climate news materialize—climate news that are large enough to affect many investors and therefore move equilibrium prices. We show that our quantity-based climate hedging portfolios—which go long assets that investors tend to buy after "local" heat shocks, and short those that they tend to sell—have superior hedging performance relative to existing approaches. While the main application of our methodology in this paper is to hedge climate risks, we also show that it can be applied to various other risks, like unemployment and housing price shocks, where investors' beliefs are affected by what they observe locally.

The central challenge when building portfolios to hedge a risk is to identify which assets are likely to do well or poorly upon the realization of the risk. The key idea behind our methodology is that we can identify "local" climate shocks that induce *some* investors (those exposed to the local shocks) to trade in a way similar to how *all* investors would trade when "global" climate shocks occur. The direction of the trades executed in response to local shocks would then predict the direction of price movements following global shocks. Potential mechanisms for why investors trade in response to local climate shocks are that these shocks may cause the investors to update their climate change beliefs or to pay more attention to global climate risks. For our empirical work, we identify a particular shock—local extreme heat events—based on a vast literature that has shown that such events are important drivers of climate change beliefs (see Egan & Mullin 2012, Deryugina 2013, Joireman et al. 2010, Li et al. 2011, Fownes & Allred 2019, Sisco et al. 2017). Based on these studies, we hypothesize (and validate empirically) that by studying how investors change their portfolio allocations in response to local heat shocks, we can predict how all investors will trade (and how prices will move) when global climate shocks occur.[1]

Our approach to building hedging portfolios consists of two steps. The first one is the cross-sectional study of mutual fund trading responses to local heat shocks. We consider different definitions of local heat shocks, based on measures of extreme temperatures as well as the presence of fatalities, injuries, or property damages from heat stress. These measures are only weakly correlated, allowing us to compare the performance of hedging portfolios based on different local shocks. Using these measures of local heat shocks, we compare how changes in mutual fund allocations across industries differ between funds that are exposed to the shocks and those that are not. Some industries enter the hedging portfolio in intuitive ways: for example, firms in the real estate sector, which has substantial physical risk exposures, are disproportionately sold by investors after a heat shock. The direction of trading in other industries is less immediately intuitive. For example, firms in the energy sector are disproportionately bought by mutual fund managers after localized heat shocks, perhaps consistent with traditional energy companies playing an important role in developing new (and cleaner) sources of energy.

In interpreting these results, it is important to keep in mind that what matters for building the hedging portfolio is whether the trading activity in response to the local shocks is mirrored in the response to the global shocks; there is no need for investor beliefs about firms' climate risk exposures to be correct, as long as they are consistent across investors, time, and between local and global shocks. We demonstrate this stability in various ways. First, we show that the industry-level trading activity in response to local shocks is similar across periods in our sample as well as across different investors. Second, we show that while our different measures of heat shocks have low correlations, their corresponding quantity-based hedging portfolios are significantly correlated, indicating that investors trade similarly in response to the different heat shocks. Finally, and perhaps most importantly, we show directly that this cross-sectional quantity information is indeed useful for learning about the pricing response to global shocks, by studying the out-of-sample hedging performance of these quantity-based portfolios for aggregate climate shocks.

The second step of our approach is the construction of the hedging portfolios for global shocks. These are long-short portfolios whose weights are determined by the response to the local shocks. We build hedging portfolios for our different heat shocks, and evaluate their performance against two alternative approaches for constructing hedging portfolios. The first one, the "narrative" approach, uses prior information about the determinants of climate risk exposures, together with firm characteristics such as ESG scores, to determine the hedging portfolio composition. For example, one possibility would be to build portfolios that are long high-ESG score firms and short low-ESG score firms, based on a prior that high-ESG score firms would do well when climate risks materialize (see Engle et al. 2020, Pástor et al. 2020, Hoepner et al. 2018). This approach has the advantage that, like ours, it does not require long time series to be implemented; however, it requires investors to have correct priors about how firm characteristics relate to performance following realizations of climate

[1]This approach relates to Choi et al. (2020), who find that carbon-intensive firms earn lower stock returns than other firms when the local stock exchange city experiences abnormally high temperatures that month.

risk. The second approach is the "mimicking portfolio" approach as in Lamont (2001), where climate risk series are projected onto a set of portfolio returns using time-series information. The mimicking portfolio approach relies strongly on time-series data: since it does not take an *a priori* view on which assets gain or lose when climate shocks occur, it needs to learn this from assets' performance during past climate risk realizations.

We assess the hedging performance of our quantity-based portfolios and these alternative hedging portfolios by computing the out-of-sample correlation between portfolio returns and various measures of global climate shocks in the 2015-2019 period. For the mimicking portfolio approach and the quantity-based approach, we construct the hedge portfolios using rolling 5-year windows of price and quantity data, respectively.[2] We construct a variety of "global" climate shocks to be used as the target of the hedge, drawing on a rapidly expanding literature that follows Engle et al. (2020) to construct different time series of news about physical and regulatory climate risks. Rather than choosing a preferred climate risk series, we evaluate how different approaches perform in hedging various series constructed by Engle et al. (2020), Faccini et al. (2021), and Kelly (2021), as well as national temperature shocks and attention to climate risk as measured through Google searches.

We find several interesting patterns. First, at a broad level, hedging climate risks is hard. Few approaches manage to achieve more than a 20% out-of-sample correlation with the climate shock series, confirming and extending the initial finding in Engle et al. (2020). Second, the vast majority of the hedging portfolio methodologies do not achieve consistent performance across the various climate shocks. That is, some approaches provide better hedging for some measures of global climate risks, and other approaches for other measures. Third, our quantity-based portfolios have the best average out-of-sample hedging performance relative to a wide range of implementations of the competing narrative or mimicking portfolio approaches. Specifically, it yields positive out-of-sample correlations between the hedging portfolio returns and every one of our aggregate climate shock series, with maximum correlations of above 30%. This suggests that the cross-sectional information on which the quantity portfolios are based is useful to hedge aggregate climate shocks.

In addition to highlighting the strengths of our quantity-based methodology, our empirical results uncover important downsides of the traditional approaches. The mimicking portfolio approach is very sensitive to the availability of time-series data, and suffers when the time series is particularly short. As an illustration, consider a mimicking portfolio that *only* uses the S&P 500. While this portfolio is composed of only one asset, historical data is still required to establish whether to take a long or short position: is the broader stock market likely to increase or decrease upon the realization of climate risks? This relationship turns out to be unstable over time: during 2010-2014, the S&P 500 comoved positively with climate risk realizations, while during 2015-2019, it comoved negatively, highlighting the challenges of mimicking portfolio approaches to constructing successful climate hedges. Narrative-based portfolios are immune to the short-sample issue, since historical data is not used to determine positions. However, deciding on positions in an a priori way is also hard,

[2] Prior to 2010, climate risks were hardly incorporated into market prices and likely did not affect investor behavior, making all of these approaches difficult to implement.

and somewhat arbitrary. For example, the energy sector includes many companies with substantial carbon emissions. As a result, one might intuitively think to short the energy sector to hedge against regulatory climate risks. Yet, during 2015-2019, the returns of the energy sector comoved positively with realizations of climate risk, perhaps due to a belief that green energy innovations will come in part from traditional energy firms.

The primary focus of our paper is to use our new quantity-based approach to construct portfolios that hedge realizations of climate risk. This is a natural application of our methodology: climate change is a first-order risk that has attracted investor attention only recently, and therefore has not yet built up enough time-series data to allow precise estimation of the climate risk exposures of different assets. However, our new approach can, in principle, be applied to hedging any macro risk series for which local events affect investors' beliefs about (or attention to) aggregate risks. For example, in recent work, Kuchler & Zafar (2019) show that locally experienced house price movements affect expectations about future U.S. house price changes; they also show that personally experienced unemployment affects beliefs about the future national unemployment rate. Consistent with our results on hedging climate risks, we show that the trading responses of mutual fund investors to local house price and unemployment changes allow us to construct portfolios that perform well at hedging innovations in the corresponding national series.

We structure our paper as follows. First, in Section 1, we describe a simple model of quantity-based information that justifies our approach to constructing hedge portfolios. We then describe the regional heat shocks that we include in our analysis in Section 2.1, and we list the climate news indices, which we use to validate the quantity-based hedge portfolios, in Section 3.2. In sections 2 and 3, we apply the novel quantity-based approach to the case of climate change realizations. We validate the quantity climate change hedge portfolios in Section 3.4 and further demonstrate the efficacy of our methodology by hedging unemployment and housing prices in Section 4. We conclude this paper by summarizing the results and listing ideas for future research.

1 Quantity-Based Portfolios: a Simple Model

In this section, we describe a simple model to illustrate the mechanism behind our quantity-based approach to building climate change hedge portfolios.

Setup. Consider a unit continuum of investors $i \in [0,1]$ who choose to invest in either security A or B. Investor i's demand for security A is given by $q_A\big(p_A, \epsilon_A(i)\big)$, where p_A is the (relative) price of security A, and $\epsilon_A(i)$ gives investor i's beliefs on the (relative) future payoffs of security A. For simplicity, assume that $q_A\big(p_A, \epsilon_A(i)\big) = f(p_A) + g\big(\epsilon_A(i)\big)$, with f and g continuously differentiable, and $\frac{\partial f}{\partial p_A} < 0$. The market-clearing condition is:

$$\int_{i=0}^{i=1} q_A\big(p_A, \epsilon_A(i)\big) di = \bar{A},$$

where $\bar{A}$ is the supply of security A. The equilibrium is characterized by price p_A^* and asset allocations $q_A^*(i)$. We focus on the equilibrium in market A; market B clears by Walras' law.

An individual investor's beliefs can be decomposed into a common component ν_A and an investor-specific component $\omega_A(i)$, such that $\epsilon_A(i) = \nu_A + \omega_A(i)$. We interpret the common belief ν_A as driven by shocks or news that are observed by all investors. $\omega_A(i)$, instead, represents beliefs that are determined by "local" shocks that are only observed by investor i. We do not impose assumptions on the origins of the common and idiosyncratic component of beliefs. Also, note that in this simple model, investors simply "agree to disagree": there is no learning from prices about the beliefs or information of other investors.

Local Shocks. We first study changes in equilibrium prices and quantities in response to a "local" shock $\omega_A(i)$, for example because investor i—having experienced a local heat shock—now believes that stricter regulations to carbon emissions will reduce the future profitability of stock A. By the chain rule we have that $\frac{\partial q}{\partial \omega_A(i)} = \frac{\partial q}{\partial \epsilon_A(i)}$. Since each investor is "small" relative to the market,

$$\frac{\partial}{\partial \omega_A(i)} \int_{i=0}^{i=1} q_A\big(p_A, \epsilon_A(i)\big) di = 0.$$

Thus, $\frac{\partial p_A^*}{\partial \omega_A(i)} = 0$. However, since investor i's demand changes, $\frac{\partial q^*}{\partial \omega_{A(i)}} \neq 0$. In words, if investor i receives a "local" belief shock, her equilibrium allocation changes. However, since the shock only affects one (atomistic) investor, this does not affect equilibrium prices. Thus, investor i's change to her *equilibrium* allocation q^* is directly informative about her demand sensitivity to beliefs, $\frac{\partial q}{\partial \epsilon_A(i)}$.

From Local Shocks to Common Shocks. Suppose now there is global news about stock A, a change in ν_A. For example, all investors now believe that climate change regulation has become more likely, reducing their expected profitability of firm A. By the implicit function theorem and the chain rule, equilibrium price responses are given by:

$$\frac{\partial p_A^*}{\partial \nu_A} = -\frac{\int_{i=0}^{i=1} \frac{\partial q_A}{\partial \epsilon_A(i)} di}{\frac{\partial q_A}{\partial p_A}}.$$

In words, the sensitivity of prices to national news is directly proportional to average quantity sensitivities, $\int_{i=0}^{i=1} \frac{\partial q_A}{\partial \epsilon_A(i)} di$. Together with the earlier result, this shows how idiosyncratic quantity responses can be used to predict national price responses. Intuitively, by studying how investors react to local shocks that have no effect on the equilibrium price, we can predict how their demand shifts in response to news that affect all investors. Those news then move the demand function of many investors simultaneously, leading to price movements that can be predicted by the response to local news.

2 Portfolio Changes Following "Local" Shocks

We begin this section by outlining the construction of our "local" climate shocks based on measures of regional extreme heat events. Then, we describe the mutual fund portfolio and adviser location data, before studying how investors change their portfolios following heat shocks in their locations.

2.1 "Local" Climate Shocks: County-level Heat Shocks

Our research design requires a set of "local" shocks that fulfill four criteria. First, these shocks need to shift the asset demand of affected investors through influencing their attention to or beliefs about climate risks. Second, the local shocks can only affect a small group of investors, so that they move only those investors' asset demands but do not affect aggregate prices. Third, we need to be able to observe the trading behaviors of the affected investors. And fourth, the shifts in asset demand in response to local climate shocks need to correspond to the shifts in asset demands in response to global climate shocks.

Local heat shocks likely satisfy all four conditions. First, an extensive literature has shown that local heat shocks are important drivers of climate change attention and beliefs (Egan & Mullin 2012, Deryugina 2013, Joireman et al. 2010, Li et al. 2011, Fownes & Allred 2019, Sisco et al. 2017); we know from Giglio, Maggiori, Stroebel & Utkus (2021) and others that beliefs affect asset allocations. Second, these heat shocks are sufficiently concentrated geographically such that they only affect the beliefs of a small subset of investors. Third, we observe the trading activities and locations of an important set of investors. The fourth requirement—that investors react to local heat shocks in similar ways as to global climate risk shocks—cannot be directly verified. However, we validate it indirectly by confirming that the resulting portfolios can hedge aggregate climate risk.

We construct three different measures of extreme local heat events. Table 1 provides an overview of the constructed heat shocks, and the maps in Appendix Figures A.1 to A.3 visualize the geographic distributions of these events.

Fatalities or Injuries from Extreme Heat. Our first measure of extreme heat shocks captures whether there were any fatalities or injuries due to extreme heat in a given county. We construct this measure using monthly information from NOAA's National Center for Environmental Information, as collected in the *Spatial Hazard Events and Losses Database for the United States* (SHELDUS) database. Panel A of Table 1 shows that about 0.13% of all county-months in the U.S. between 2010 and 2019 had fatalities or injuries due to heat.

Crop Indemnity Payments due to Extreme Heat. We construct a second measure of local heat shocks from crop indemnity payments. The underlying data is collected by the U.S. Department of Agriculture, and we use a version maintained by SHELDUS.[3] We define

[3]Crop indemnity payments are insurance payments to farmers, which are paid when external disruptions lead to crop yields or revenues below the agreed amount in the insurance contract. The U.S. Department of

Table 1: Summary Statistics on Extreme Heat Measures

Panel A: Local Heat Shocks: Summary

Climate Shock	Event Description	Frequency Monthly	Frequency Sample
Heat: Fatalities/Injuries	Injuries or fatalities	0.13%	1.32%
Heat: High Indemnities	90th percentile indemnity payments	0.79%	0.54%
Heat: Record Temperature	10-year record of 3-month-average county temperature	0.91%	1.06%

Panel B: Local Heat Shocks: Monthly Jaccard Correlations

	Fatalities/Injuries	Indemnities	Record Temperature
Heat: Fatalities/Injuries	1.00		
Heat: Indemnities	0.01	1.00	
Heat: Record Temperature	0.02	0.07	1.00

Panel C: Local Heat Shocks: Sample Jaccard Correlations

	Fatalities/Injuries	Indemnities	Record Temperature
Heat: Fatalities/Injuries	1.00		
Heat: Indemnities	0.06	1.00	
Heat: Record Temperature	0.03	0.02	1.00

Note: Panel A provides an overview of the constructed heat measures. All shocks are coded as dummy variables indicating the existence of a severe heat event. The event description lists the dummy criteria. The "monthly" frequency shows the share of month-county observations across all U.S. from 2010 to 2019 that experience the heat event. The "sample" frequency shows the share of observations in our final sample that experience the heat event. The differences in "monthly" and "sample" frequency arise from mutual fund advisers generally being concentrated in high population density counties; injuries and fatalities from heat also disproportionately occur in areas with high population densities, while heat-related crop indemnity payments are more common in rural areas. Also, the "sample" frequency uses three-month windows instead of monthly windows, generally leading to a somewhat higher prevalence. Panel B shows the Jaccard correlation between the constructed heat measures across all county-months from 2010 to 2019, whereas Panel C shows the Jaccard correlation among our final sample. Intuitively, the Jaccard correlation measures the likelihood of observing both shocks conditional on observing one of them.

an extreme heat indemnity event when the monthly heat-related crop indemnity payments in a given county exceed the 90th percentile of non-zero payments across all U.S. county-months in the past 10 years; about 0.8% of county-months between 2010 and 2019 had such an event. Panel B of Table 1 highlights that the correlation of high crop indemnity heat events with heat-related fatalities and injuries is essentially zero. Crop indemnity payments are more frequent in low-density rural areas, whereas fatalities and injuries due to heat are more frequent in urban areas. Crop indemnity shocks therefore provide a source of variation for our analysis that is independent of the shock related to fatalities and injuries.

Extreme Temperatures. While SHELDUS extreme heat shocks capture the most devastating events—events that usually involve very high absolute levels of temperatures—they do not necessarily capture all instances when temperatures are high relative to normal temperatures in comparably colder regions. Therefore, we construct a third county-level heat

Agriculture reports these payments for several private insurance companies, covering more than 100 crops.

shock measure based using temperature data from the PRISM Climate Group. For each quarter, we compute the average daily temperature. We then define a heat shock as a period in which this average temperature exceeds the record quarterly average temperature for the given county over the past ten years.

2.2 Extreme Heat & Climate Change Attention—Google Searches

We next explore the ability of our local heat shock measures to affect local climate change attention, measured by Google searches for "climate change" (see Stephens-Davidowitz 2014, Choi et al. 2020, for similar approaches). Since Google search interests are not available at the county level and are often missing at the MSA level, we conduct this analysis at the state-month level, and aggregate our measure of heat shocks to the state level.[4]

The Google search series measure relative interest in a topic, e.g., the fraction of all Google searches in a region that search for "climate change." In every period, Google scales the relative search interest cross-sectionally to be between 1 and 100. This means that, in each period, the region with the most relative searches for a given term receives a score of 100. All other regions' scores represent their relative searches as a fraction of the relative searches of the highest-ranked region. For example, if region A is the region with the most relative searches and region B has half as many relative searches, then region B's score would be 50. Given this multiplicative scaling factor, we explore how local climate shocks affect the logarithm of the Google searches using the following specification:[5]

$$\log(\widetilde{G_{t,s}}) = \beta_S S_{t,s} + \delta_s + \gamma_t + \epsilon_{t,s}, \tag{1}$$

where $\widetilde{G_{t,s}}$ is the scaled Google search interest for climate change in state s at time t, and $S_{t,s}$ is the corresponding climate shock indicator. State and time fixed effects are captured by δ_s and γ_t. We cluster standard errors at the month and state level, and weight observations by the state's population.

Table 2 reports the coefficients β_S from running this regression separately for the different indicators of extreme local heat. All coefficients are positive and statistically significant. Intuitively, experiencing any fatalities or injuries from heat is associated with an increased relative interest in climate change by 5%. Reported crop indemnity payments increase the relative Google searches by 7%, whereas record temperatures are associated with an increase by 8%. These findings suggest that all heat measures can affect climate change awareness and are, therefore, suitable candidates to discipline the quantity hedge portfolio.

[4] In our baseline analysis, a state is recorded to experience a temperature shock if at least one of the counties experiences a temperature shock. Similarly, state-level fatality/injury shocks are defined by at least one fatality or injury occurring within the state during the month. The indemnity shocks are based on the sum of the indemnity payments within each state relative to the 90th percentile of non-zero payments across all states over the past 10 years. The findings are robust to alternative ways of aggregating county-level heat shocks to the state level, or to using more continuous measures, such as injuries or fatalities per captia.

[5] Let $G_{t,s}$ be the unscaled Google search interest in climate change for month t and state s. We observe only $\widetilde{G_{t,s}} = G_{t,s}/\eta_t$, where η_t is the unobserved scaling factor for month t. By regressing $\log(\widetilde{G_{t,s}}) = \beta_S S_{t,s} + \delta_s + \gamma_t + \epsilon_{t,s}$, we ensure that the time fixed effect captures the scaling factor.

Table 2: Heat Shocks and Climate Attention

	Log(Google Search Volume)		
Heat: Fatalities/Injuries	0.05**		
	(0.03)		
Heat: High Indemnities		0.07**	
		(0.03)	
Heat: Record Temperature			0.08**
			(0.04)
R^2	0.77	0.77	0.77
State & Month FE	Y	Y	Y
N	5,506	5,506	5,506

Note: Table shows results from regression 1. Standard errors in parenthesis are clustered at the month and state level, and observations are weighted by the state's population size. Significance levels: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

2.3 Holding and Location Data

To explore how mutual fund managers trade following local extreme heat events, we need to observe their holdings and locations.

Holdings data. We use the Thomson Reuters Mutual Fund Holdings S12 database to obtain a time series of portfolio holdings of U.S. mutual funds. We combine the holdings with fund characteristics from CRSP. The mutual funds are linked from Thomson Reuters to CRSP using their Wharton Financial Institution Center Number (WFICN) as reported in WRDS MFLINKS. Most funds report their holdings in three-month intervals. Therefore, in the following analysis, we will focus on holding changes at three-month intervals.

We restrict holdings to assets with share codes 10, 11, 12, and 18, and exchange codes 1, 2, and 3,[6] which focuses our hedge assets on North American common stocks. Since we want to identify deliberate fund manager asset reallocations in response to local climate shocks, we restrict our analysis to actively managed funds. Therefore, we only keep funds that have Investment Objective Code 2 ("Aggressive Growth"), 3 ("Growth"), 4 ("Growth & Income"), or missing, and have CRSP Objective Code "Equity Domestic Non-Sector."[7]

We obtain stock characteristics from CRSP and Compustat. We assign holdings their end-of-month prices from CRSP. We obtain firm GICS industry codes from Compustat by merging the stocks on their CUSIP identifiers. The first four digits of the GICS code determine the stock's classification into the 24 "industry groups" that are the main focus of our analysis.[8]

[6]These share codes represent Ordinary Common Shares that are 'not further defined,' 'need not be further defined,' 'incorporated outside the U.S.,' or 'REITs (Real Estate Investment Trusts).' Exchange codes 1, 2, and 3 represent the NYSE, American Stock Exchange, and Nasdaq Stock Market, respectively.

[7]These restrictions are fairly standard (e.g., Song 2020), and we show that our results are robust to alternative choices.

[8]The Global Industry Classification Standard (GICS) is developed by MSCI and S&P based on earnings

Measuring Active Changes. In our main cross-sectional analysis, we want to explore how localized heat shocks induce changes in the share of value invested into industry I by fund f through active trading. We perform our analysis at the industry level, since higher granularity would result in potentially noisy estimates. For every fund f and time t, we define the active change in industry I holdings as:

$$ActiveChanges^{I}_{f,t} = \left[\left(\frac{\sum\limits_{j \in I_{t-1}} P_{j,t-1} S_{f,j,t}}{\sum\limits_{j} P_{j,t-1} S_{f,j,t}} \right) - \left(\frac{\sum\limits_{j \in I_{t-1}} P_{j,t-1} S_{f,j,t-1}}{\sum\limits_{j} P_{j,t-1} S_{f,j,t-1}} \right) \right] \frac{1}{(Share^{I}_{t})}, \quad (2)$$

where $P_{j,t-1}$ denotes the previous period price for stock j, $S_{f,j,t}$ denotes the number of shares of stock j held by fund f in period t, and $Share^{I}_{t}$ measures the market capitalization share of industry I as a fraction of the U.S. stock market. The term in square brackets therefore captures the active change in the portfolio share of a given industry. The reason for scaling by industry size is that a given increase in the portfolio share of a particular industry (i.e., shift of a given dollar amount invested) is likely to induce larger price movements for smaller industries. Since most funds report their holdings quarterly, we measure fund composition changes in three-month intervals, i.e., a time unit here represents a three-month interval.[9]

Location data. We also obtain data on mutual fund advisers' locations, since these fund advisers are primarily responsible for making asset-allocation decisions (see Chang 2019). U.S. mutual funds must publicly file with the SEC at regular intervals, and we use these filings to determine their fund advisers and their fund adviser locations. Specifically, we parse adviser locations from N-SAR filings until 2017 and parse them from N-CEN filings from 2018 onward, which replaced the N-SAR format. These SEC filings cannot be matched directly with Thomson Reuters or CRSP mutual fund data. Therefore, we apply a fuzzy string matching algorithm to match SEC filings with mutual funds. We only include almost perfect name matches, and successfully match 84.1% of fund-quarter observations. Overall, our sample that matches quarterly fund reports to location data includes 2,496 unique funds, making up 58,007 fund-quarter observations (an average of 23.2 observations per fund) between 2010 and 2019.[10]

and market perception in combination with revenues to classify companies.

[9] Alternatively, one could analyze a separate variable, $PassiveChanges^{I}_{f,t}$, where the first fraction uses $P_{j,t}$ instead of $P_{j,t-1}$, i.e., current period holdings are valued at current period prices. This alternative approach takes price changes into account, and would be a more suitable model if we assume that funds can freely and constantly rebalance their portfolio. We verified that, in practice, both approaches generate similar hedging portfolios.

[10] As we describe in more detail below a fund-quarter observation involves two consecutive holding reports spaced three months apart, allowing us to analyze the active trading of mutual funds over the quarter.

Figure 1: Locations of Mutual Fund Advisers

Panel A: Adviser Locations - Largest Counties

FIPS	County	State	% Funds	% Fund-Quarters
36061	New York	NY	22.4	21.1
25025	Suffolk (Boston)	MA	14.0	10.2
17031	Cook (Chicago)	IL	5.5	4.8
06075	San Francisco	CA	4.1	3.0
06037	Los Angeles	CA	3.3	3.6

Panel B: Adviser Locations - Largest States

State name	State	% Funds	% Fund-Quarters
New York	NY	25.8	24.9
Massachusetts	MA	14.3	10.3
California	CA	10.3	8.9
Illinois	IL	6.8	6.5
Pennsylvania	PA	5.7	5.7

Note: The map shows the distribution of the locations of mutual fund advisers in our final sample. The map only includes funds where all advisers reside in the same county and shows the number of funds for each county. Panel A shows the share of funds residing in the most represented counties in our sample, whereas Panel B shows this information for the most represented states. Both panels are based on the subset of funds whose advisers all reside in the same location.

For 67.6% of funds, all advisers reside in the same county. Whenever funds have multiple advisers who are not all located in the same county, we assign fund shock exposure as an average of fund adviser shock exposures. For example, if a fund has two advisers in county A and one adviser in county B, and county A is affected by a climate shock, we assume the fund is affected by a 2/3 climate shock. Figure 1 shows the geographic distribution of the fund advisers for the subset of funds where all advisers reside in the same county. While some areas see larger concentration of advisers, advisers are generally spread through the entire country. Panel A shows that about a quarter of advisers are located in New York (most of them in New York City), 14.3% of them are in Massachusetts (most of them in Boston), and about 10.3% of advisers are located in California (roughly equally split between San Francisco and Los Angeles). This gives us important geographic variation and therefore differential exposure to local heat shocks.

Summary Statistics. Panel A and B of Table 3 present summary statistics on the GICS industries and the portfolio holdings of the advisers in the final analysis sample. In the average quarter, the Energy sector (GICS code 1010) included 224 unique companies held by advisers in our sample, and made up about 7.2% of total market capitalization. The smallest industry by average market capitalization was "Auto & Components", comprising of an average of 43 firms with an average market capitalization of 0.9%. On average, mutual funds in our sample held 209 unique firms across 19.5 unique industries. At the 10th percentile, they held 33 firms across 14 industries.

Panel C of Table 3 shows summary statistics on the "active changes" variable. Intuitively, active changes of 0 imply that industry I's relative weight within the portfolio remained unchanged. On the other hand, active changes of 1 imply that I's relative weight increased by I's relative size. For example, if I makes up 10% of the market, and the fund increased I holdings from 5% to 15% of the portfolio, then the active changes would be 1. The average and the median active change in our sample is zero. The 1% quantile is -1.26, i.e., funds were decreasing their share invested in I by more than I's relative market size.

2.4 Estimating the Response to Local Climate Shocks

To understand how portfolio composition varies with exposure to local shocks, we estimate the following panel regression separately for each industry I:

$$ActiveChanges^I_{f,t} = \beta^I S_{loc(f),t} + \delta^I_t + \epsilon_{f,t}, \tag{3}$$

where $ActiveChanges^I_{f,t}$ is defined as in equation 2, $S_{loc(f),t}$ is a local, time-varying climate shock described in Section 2.1, and δ^I_t captures time fixed-effects. The main object of interest at this step is β^I: for each industry, this represents the differential change in fund holdings of that industry for funds affected by a "local" climate shock, relative to the change in holdings for funds that were not affected by a local climate shock. We refer to this coefficient as *industry-specific climate loading*.

Table 4 reports the estimated industry-specific climate loadings, for the three local heat measures, and averaged across the three, when we estimate this regression at the end of 2019, using 5 years of backward-looking data. Industries toward the top of the table have positive weights, indicating that funds affected by local heat shocks tend to increase their holdings in those industries; industries at the bottom of the table, with negative coefficients, are disproportionately sold by mutual funds whose advisers experience an extreme local heat event.

Table 3: Sample Summary Statistics

Panel A: Industry Summary Statistics

GICS	Industry	Number of Companies			Share of Market (%)		
		Avg.	Min	Max	Avg.	Min	Max
1010	Energy	224	197	248	7.2	4.1	10.9
1510	Materials	202	173	223	3.6	2.1	4.7
2010	Capital Goods	325	299	346	7.8	4.8	8.8
2020	Commercial & Prof. Serv.	131	121	145	1.5	1.3	1.7
2030	Transportation	70	59	85	2.6	1.8	3.2
2510	Auto & Components	43	40	46	0.9	0.6	1.2
2520	Consum. Durables & Apparel	119	109	137	2.0	1.3	2.5
2530	Consumer Services	135	117	150	2.4	2.1	3.0
2550	Retailing	154	144	162	6.0	3.4	7.0
3010	Food & Staples Retailing	27	22	33	1.4	1.1	1.6
3020	Food, Bev. & Tobacco	93	81	104	4.3	3.3	5.2
3030	Household & Pers. Prod.	37	34	45	1.6	1.2	1.9
3510	Health Care Equip. & Serv.	251	230	289	6.4	5.7	7.1
3520	Pharma., Biotech., & Life Sc.	364	261	510	7.8	6.3	9.6
4010	Banks	435	399	507	6.4	5.4	8.5
4020	Diversified Financials.	161	148	171	4.9	4.2	6.2
4030	Insurance	106	92	128	3.0	2.4	3.5
4510	Software & Services	279	259	304	9.4	7.8	13.3
4520	Tech. Hardw. & Equip.	220	172	275	5.4	2.1	7.4
4530	Semiconductors & Equip.	110	82	137	3.6	2.9	4.5
5010	Communication Services	42	31	53	1.7	1.3	2.5
5020	Media & Entertainment	108	84	135	5.2	2.2	12.1
5510	Utilities	90	76	103	2.7	2.3	3.2
6010	Real Estate	145	109	177	2.2	1.2	4.4

Panel B: Mutual Fund Summary Statistics

	Number of Companies			Number of Industries		
	Avg.	p10	p90	Avg.	p10	p90
Mutual Fund Holdings	209	33	467	19.5	14.0	24.0

Panel C: Active Changes Summary Statistics

	Mean	p1	p25	p50	p75	p99
Active Industry Change	-0.00	-1.26	-0.06	0.00	0.05	1.32

Note: Panel A shows among the universe of stocks held by the funds in our final analysis sample, the average, minimum, and maximum number of companies and market share for each industry at the monthly level between 2010 and 2019. The unit of observation is an industry-quarter and the sample size is 960. Similarly, Panel B shows the average and the 10th and 90th percentiles of companies and industries in our sample of eligible fund-quarters. The unit of observation is a fund-quarter (each report) and the sample size is 72,550 (note that active changes require two consecutive reports, which are not always available). Panel C shows summary statistics for the active industry changes as defined in Equation (2). The unit of observation is a fund-quarter-industry change and the sample size 1,156,344.

Table 4: Industry Climate-β Coefficients

GICS	Description	Avg.	Fatalities/Injuries	Indemnities	Record Temp.
2510	Auto & Components	0.11	0.07	0.15	0.15
4520	Tech. Hardw. & Equip.	0.09	0.05	0.21	0.06
2030	Transportation	0.06	0.02	0.13	0.08
4530	Semiconductors & Equip.	0.05	0.05	−0.01	0.12
3010	Food & Staples Retailing	0.04	0.03	0.08	0.03
5010	Communication Services	0.03	0.04	0.02	0.00
1010	Energy	0.02	0.03	0.04	−0.01
3020	Food, Bev. & Tobacco	0.02	0.01	0.07	−0.01
4020	Diversified Financials.	0.02	0.01	0.01	0.04
5510	Utilities	0.02	0.01	0.03	0.02
4010	Banks	0.02	0.04	0.01	−0.03
2010	Capital Goods	0.02	0.01	0.06	0.00
4510	Software & Services	0.00	0.01	−0.04	0.03
4030	Insurance	−0.00	−0.03	0.06	0.00
3520	Pharma., Biotech., & Life Sc.	−0.01	0.01	−0.02	−0.02
6010	Real Estate	−0.01	−0.03	0.00	−0.00
5020	Media & Entertainment	−0.02	−0.03	0.05	−0.06
3030	Household & Pers. Prod.	−0.02	0.01	−0.07	−0.03
2530	Consumer Services	−0.02	−0.06	−0.02	0.05
1510	Materials	−0.03	−0.03	−0.01	−0.03
3510	Health Care Equip. & Serv.	−0.03	−0.02	−0.07	−0.01
2550	Retailing	−0.05	−0.07	0.01	−0.05
2520	Consum. Durables & Apparel	−0.06	0.02	−0.19	−0.08
2020	Commercial & Prof. Serv.	−0.12	−0.13	−0.28	0.09

Note: Industry climate beta coefficients as in equation (3). The coefficients are sorted by the average coefficient and are based on data from 2015 to 2019 inclusive. Therefore, these are the most current industry climate betas in our sample.

Two results in this table are interesting. First, the identities of industries that are bought/sold are not necessarily those expected ex ante. For example, while it may be expected that materials and real estate appear towards the bottom of the table, the automobile industry, tech and energy all appear at the top. One potential interpretation of this result is that these are industries that, while currently potentially producers of emissions, could be the source of innovation (e.g., electric vehicles, new energy sources) that actually might make them fare well in the face of stricter climate regulations.

Of course, some of these results could alternatively be due to estimation noise. Two considerations are then important. First, to the extent that these numbers reflect investors' beliefs, and those investor beliefs tend to be consistent between local and global shocks, this is sufficient to build a good hedging portfolio – there is no need for this purpose for the beliefs to be correct. Second, whether signal dominates noise in this estimation is an empirical question. Only a proper evaluation of the hedging ability of the quantity-based portfolio – examined in the next section – can reassure us that these estimates are meaningful.

A second interesting result from the table is that the ordering and the signs of the industries are correlated across different measures of heat shocks (columns). To explore this in greater detail, we report in Table 1 the raw correlations of the heat measures. As noted above, they are all close to zero. Table 5 reports the correlation and rank-correlation, respectively, of the industry-specific climate loadings, in the validation period of 2015-2019. Interestingly, all correlations of the actual industry ordering are high. This indicates that mutual funds change their portfolios in a consistent way in response to these different heat shocks.

Table 5: Across-Shock Correlation of Industry-Betas

Panel A: Pearson Industry Climate Beta Correlation

	Fatalities/Injuries	Indemnities	Record Temperature
Heat: Fatalities/Injuries	1.00		
Heat: Indemnities	0.56	1.00	
Heat: Record Temperature	0.20	0.25	1.00

Panel B: Spearman (Rank) Industry Climate Beta Correlation

	Fatalities/Injuries	Indemnities	Record Temperature
Heat: Fatalities/Injuries	1.00		
Heat: Indemnities	0.42	1.00	
Heat: Record Temperature	0.29	0.28	1.00

Note: Panel A shows the Pearson correlation among the industry climate beta coefficients determined from various local heat shocks as in equation (3). Similarly, Panel B shows the Spearman *rank* correlation among the industry climate beta coefficients. The coefficients are based on data from 2015 to 2019 inclusive. Therefore, these are the most current industry climate betas in our sample.

Another measure of estimation noise is the ability of the data to give consistent results across randomly selected subsamples. We tested this by splitting the sample into two mutually exclusive subsamples with varying split rules. For each different approach of subdividing the sample, we computed the climate beta coefficients for 100 independent iterations and, for each iteration, we computed the rank-correlation and correlation of the resulting coefficients between both subsamples. Panel A of Table 6 reports the average rank-correlation and correlation for stratified and fully random sampling. The stratified sampling ensures that each subsample receives approximately half of the observations of each period - location combination, whereas the fully random sampling imposes no restrictions on the selected observations for the two subsamples. Both approaches achieve relatively high coefficient correlations, ranging from 0.17 to 0.44, indicating that the sample consistently picks up a common signal and is not driven by major outliers.

Similarly, Panel B Table 6 reports the average correlations from splitting the sample either by funds, periods, or counties. For example, with the fund split, each fund fully belongs to exactly one of the two selected mutually exclusive subgroups. Again, the relatively high correlations indicate little sensitivity of our sample with regard to only using a subset of the data, randomly, stratified, or in terms of included funds, periods, or locations.

Finally, we investigate the shifts in industry climate coefficients over long time horizons. While Table 4 shows the resulting coefficients when using data from 2015 to 2019, i.e., the very end of our rolling 5-year sample, Appendix Table A.1 shows the resulting coefficients

Table 6: Across-Sample Split Correlation of Industry-Betas

Panel A: Random Split within Groups

	Stratified		Fully Random	
Climate Shock	Spearman	Pearson	Spearman	Pearson
Heat: Fatalities/Injuries	0.42	0.44	0.37	0.39
Heat: High Indemnities	0.24	0.29	0.22	0.29
Heat: Record Temperature	0.32	0.23	0.28	0.17

Panel B: Random Split between Groups

	Fund Split		Period Split		Location Split	
Climate Shock	Spearman	Pearson	Spearman	Pearson	Spearman	Pearson
Heat: Fatalities/Injuries	0.41	0.44	0.18	0.17	0.30	0.30
Heat: High Indemnities	0.22	0.29	0.21	0.22	0.07	0.04
Heat: Record Temperature	0.30	0.20	0.21	0.11	0.22	0.24

Note: This table shows the average Spearman (rank) and Pearson correlation of the quantity beta coefficients from 100 iterations of a sample split robustness test. Panel A shows the results of splits within groups. For each iteration, the stratified sample split randomly divides the sample into two mutually exclusive subsamples that are stratified by year-month and county. Similarly, for each iteration, the fully random split randomly divides the sample into two subsamples (without any restrictions on the resulting subsamples). Panel B shows the results of splits between groups, i.e., from splitting the sample either by funds, periods, or counties. For example, with the fund split, each fund fully belongs to exactly one of the two selected mutually exclusive subgroups.

when using data from 2010 to 2014. Notably, the automobile industry achieves the largest climate beta coefficients in both horizons, suggesting a strong market belief in the opportunities arising from the transition to green energy. Similarly, technological hardware and semiconductors are strongly favored after localized climate shocks during both periods. However, the industries with the most negative climate coefficients appear to have shifted more during the past ten years. While retailing, real estate, and insurance remain at the lower end of the ordering, many industries have shifted. Most notably, consumer durables and apparel was one of the most positive climate beta industries from 2010 to 2014, but is now one of the most negative industries.

The industry shifting occurs primarily for two reasons. First, both periods involve estimation noise which inevitably causes some shifts in the industry coefficients. And second, industries and government focuses are consistently changing. While an industry could have been an inefficient polluter or of negligible national interest in the past, it could have evolved to utilize greener operations or be viewed more favorably by politicians. Therefore, in our main analysis, we use a 5-year rolling window to determine the quantity betas to always reflect up-to-date information.

3 Quantity-Based Climate Hedging portfolios

We next describe how we use the climate-βs estimated above to build the quantity-based climate hedging portfolios. We also evaluate the out-of-sample hedging performance of our portfolios over the period of 2015-2019, and compare the hedging performance against other approaches in the literature.

3.1 Portfolio Construction and Description

For each month t, industry I, and local climate shocks S, we estimate $\beta^I_{S,t}$, as in equation 3, using the previous five years of data on mutual fund portfolio compositions and climate shocks. We construct the excess returns of the corresponding quantity-based hedging portfolio as:

$$QP_{S,t} = \sum_I \widehat{\beta^I_{S,t-1}}(R^I_t - R^f_t), \tag{4}$$

where R^I_t is the value-weighted industry return and R^f_t denotes the risk-free rate. Note that each component of the portfolio is an excess return, so we do not need to scale the $\beta^I_{S,t}$.

Panel A of Table 7 shows the monthly return correlation of the quantity-based hedging portfolios based on our three climate news shocks. Given that the three types of local climate shocks are practically uncorrelated (see Table 1), the high correlations in the return series provide strong evidence that our three shocks are picking up a common signal.

Table 7: Portfolio Return Correlations

Panel A: Pearson Portfolio Return Correlation			
	Fatalities/Injuries	Indemnities	Record Temperature
Heat: Fatalities/Injuries	1.00		
Heat: Indemnities	0.61	1.00	
Heat: Record Temperature	0.53	0.37	1.00
Panel B: Orthogonalized to Fama-French 3-factors			
	Fatalities/Injuries	Indemnities	Record Temperature
Heat: Fatalities/Injuries	1.00		
Heat: Indemnities	0.61	1.00	
Heat: Record Temperature	0.55	0.18	1.00

Note: Panel A shows the monthly return correlation among our three quantity portfolios for the validation period of 2015 to 2019. Panel B shows the corresponding monthly return correlation after orthogonalizing each portfolio with respect to the Fama-French market, size, and value factor.

We next investigate how much of the portfolio return correlations are driven by a potential common loading of the three climate-based hedging portfolios on some of the Fama-French factors. To identify the factor loadings of the quantity portfolios, we regress the portfolio returns on the returns of the market, size, and value factors:

$$QP_{S,t} = \alpha + \beta^M_c(R^M_t - R^f_t) + \beta^{SMB}_c SMB_t + \beta^{HML}_c HML_t + \epsilon_{t,S}. \tag{5}$$

Table 8 shows the regression results. All portfolios have a significant loading on the market factor, no loading on size, and, the "Fatalities/Injuries" portfolio has a significant positive loading on HML (note that the magnitude of these exposures is not meaningful given that the scale of the quantity portfolio is arbitrary). Overall, the time-series variation in the Fama-French factors captures 20-40% of the variation in the quantity portfolio.

In Panel B of Table 7, we show the return correlation of the three quantity-based hedging portfolios after orthogonalizing the returns to the three Fama-French factors (i.e., taking the correlations of the residuals from regression 5). When compared to Panel A, the correlation coefficients are very similar, suggesting that a common loading on the Fama French factors is not the main driver of a high return correlation across these portfolios.

Table 8: Factor Exposures of Hedging Portfolios

	Return of Quantity-Based Climate Hedging Portfolio		
	Fatalities/Injuries	High Indemnities	Record Temperature
$R^M - R^f$	0.06**	0.13***	0.19***
	(0.03)	(0.04)	(0.06)
SMB	0.02	−0.03	0.00
	(0.03)	(0.05)	(0.06)
HML	0.16***	0.05	−0.01
	(0.03)	(0.06)	(0.05)
Constant	−0.01	−0.07	0.06
	(0.09)	(0.14)	(0.14)
R^2	0.40	0.20	0.31
N	60	60	60

Note: Regression of monthly returns of the quantity-based climate hedging portfolios on the market, size, and value factors as in equation 5. The sample period is 2015-2019. Heteroskedasticity-robust standard errors in parentheses. Significance levels: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

3.2 Climate Hedge Targets

One challenge with designing portfolios that hedge climate risks is that there is no unique way of defining the hedge target. Climate change is a complex phenomenon and presents a variety of risks, including physical risks such as rising sea levels and transition risks such as the dangers to certain business models from climate regulation. Different risks may be relevant for different investors, and their realizations do not always co-move. In addition, climate change is a long-run threat, and we would ideally build a portfolio that hedges the long-run realizations of climate risk, something difficult to produce in practice.

To overcome these challenges, Engle et al. (2020) have shown that the objective of hedging long-run realizations of a given climate risk can be achieved by constructing a sequence of hedges against news (one-period innovations in expectations) about future realizations of climate risks. Following the initial work in Engle et al. (2020), researchers have developed a number of climate news series capturing a variety of different climate risks. In this paper, we do not take a stand on the right hedge target, but simply assess the ability of our approach to hedge different types of climate news shocks. To do so, we look at a broad range of measures proposed in the recent literature, which we describe in the following. Building on Engle et al. (2020), we consider innovations in these climate news indices from an AR(1) model as our hedge targets. Specifically, for a given climate news series c, we define these innovations—and hence our hedge targets—in month t as $CC_{c,t}$.

Engle et al. (2020). The Wall Street Journal (WSJ) and Crimson Hexagon Negative News (CHNEG) climate news indices created by Engle et al. (2020) are, to our knowledge, the first climate news series used as hedge targets. The first series captures the number of news articles in the WSJ dedicated to climate change (broadly assuming that "no news is good news"), the latter builds upon proprietary news aggregations from Crimson Hexagon combined with sentiment analysis that allows the separation of good news and bad news. Both indices capture a mix of physical and transitional risks. These news indices are monthly and capture the period of July 2008 to June 2017. Therefore, whenever we validate our portfolio with these climate news indices, the sample ends earlier in June 2017.

Ardia et al. (2021). Ardia et al. (2020) build on the WSJ index of Engle et al. (2020) by including several media outlets and differentiating between positive and negative news. Their index is called the Media Climate Change Concerns index (MCCC) and is available at daily frequency. The index covers the period of January 2003 to June 2018. We aggregate to monthly frequency by taking the average.

Faccini, Matin, and Skiadopoulos (2021). We include four of Faccini et al. (2021)'s climate news indices: news about international climate summits, global warming, natural disasters, and narrative indices. The international summits, global warming, and natural disasters indices measure news coverage of the respective topics, whereas Faccini et al. (2021) construct the narrative index by manual reading and classifying of 3,500 articles. Further, the international summits and narrative indices capture news about transitional risk, while the global warming and natural disasters indices are more likely to capture news about physical risk, though bad news about realizations of physical risks may also make subsequent regulation more likely. These news measures are available at daily frequency and capture the period of January 2000 to November 2019. We aggregate them to monthly frequency by taking the average.

Kelly (2021). Kelly (2021) creates three climate news series that reflect general, physical, and transitional risk, respectively. Each of these series is the product of the number of relevant Wall Street Journal articles in a month multiplied by their sentiment, such that higher levels correspond to more "bad news" about risk realizations in the respective category.

National Google searches. This climate news series is the national Google search interest in "climate change," capturing general attention to climate change and its associated risks. This index does not differentiate between positive and negative news, and it could be associated with any type of climate risk.

National Temperature Deviations. Just as local extreme temperatures increase climate change awareness, so do U.S.-wide extreme heat events. Barnett (2017) shows that monthly temperature innovations from a rolling one-sided Christiano-Fitzgerald bandpass filter induce significant stock market reactions. We replicate the approach and include such innovations as one of the climate news series.

3.3 Alternative Approaches to Build Hedging Portfolios

We want to compare the hedging performance of our quantity-based portfolios described in Section 3.1 against the hedging performance of two alternative approaches to constructing hedge portfolios: narrative-based approaches and mimicking portfolio-based approaches.

Narrative-based approaches. The first alternative approach we consider selects portfolio weights of different assets based on an *ex-ante* view of the exposure of those assets to climate risks. One example of such an approach would be to use environmental scores constructed by ESG data providers to build the portfolios, for example based on the prior view that high-ESG-score companies will fare better when climate risks materialize. An alternative approach would be to use specific groups of stocks (e.g., green energy stocks) under the prior view that those groups' exposures to certain types of climate risks are predictable ex-ante. We build several portfolios using such a narrative-based approach.

Our first narrative portfolio takes positions in all U.S.-listed stocks covered by Sustainalytics ESG scores: the portfolio's position in each stock is the stock's ESG score percentile in each period, minus 50. In other words, the portfolio takes a long position of 50 in the company with the highest ESG score and a short position of -50 in the company with the lowest score in each month. Stocks with the median ESG score are not held.

A second strategy within this narrative category uses groups of stocks to take a directional view. We build our portfolios using two ETFs: the Invesco Global Clean Energy ETF (Ticker: PBD), which invests in firms focused on the development of cleaner energy and conservation, and the Energy Select Sector SPDR Fund (Ticker: XLE), which tracks a market-cap-weighted index of U.S. energy companies in the S&P 500 index. This approach builds on the prior that realizations of climate change news should increase PBD's returns and decrease XLE's returns. Therefore, the hedging portfolio would go long PBD and short XLE.

Our third narrative-based portfolio is the stranded asset portfolio as in Jung et al. (2021) based on the XLE, VanEck Vectors Coal (KOL), and the SDPR S&P 500 (SPY) ETF, using the following weights: $0.3XLE + 0.7KOL - SPY$.

Mimicking portfolio-based approaches. Mimicking portfolio approaches build a portfolio of a pre-determined set of assets that is maximally correlated in-sample with the climate change shock. Using different assets (and imposing different constraints on the resulting mimicking portfolio) will produce different mimicking portfolios. To obtain the mimicking portfolios, we estimate the regression:

$$CC_{c,t} = wR_t + \epsilon_{c,t}$$

where $CC_{c,t}$ denotes the climate hedge target of type c in month t, w is a vector of N portfolio weights, and R_t is a vector of demeaned excess returns. The portfolio weights are estimated each month using a 5-year rolling window.

We consider different sets of excess returns to build mimicking portfolios. First, we use the three Fama-French factors (Market, SMB, and HML). Second, we use the two ETFs described above, PBD and XLE, in combination with the factors. Third, we add to the Fama-French factors the excess returns of the 24 GICS industries. For this industry-based portfolio, we estimate both the standard mimicking portfolio, and a regularized version via LASSO, choosing the tuning parameter by cross-validation in an attempt to minimize the dangers from in-sample overfitting.

3.4 Hedging Climate Shocks: Evaluation of Hedge Portfolios

In this section, we evaluate the hedging performance of the different portfolios. For the quantity-based and mimicking portfolio-based approaches, for every month in our testing period 2015-2019, we construct the portfolios as described above using a five-year rolling window of data; the narrative-based portfolios are unchanged over time. We focus on the post-2010 period to train our models, since, before 2010, investors paid very little attention to climate risks. As a result, we would not expect information on prices and quantities to be useful in building hedging portfolios today.

As a criterion to evaluate the various hedging approaches, we compare the out-of-sample correlations between the hedging portfolio returns and the AR(1) innovations to the various climate news series, $CC_{c,t}$.[11] Table 9 reports these out-of-sample correlations at the monthly frequency.[12] Each row in the table represents a hedging portfolio, whereas each column corresponds to a different climate news series. All climate news series are coded such that high numbers are indicative of negative climate news. Therefore, positive correlation coefficients show successful hedges. The same information is displayed in Figure 2. Each point in the dot plot is the correlation coefficient of a hedge portfolio return with one of the climate

[11] This evaluates the hedging ability of the portfolio up to a scaling parameter. Our methodology does not identify the scale of the positions of the hedging portfolio. Such a scale could also be estimated from a training sample, at the cost of having to rely on historical correlations between aggregate shocks and portfolio returns. We leave this analysis for future work.

[12] We validate the hedge portfolios at the monthly return frequency because, for many events, it is hard to pin down the occurrence to a specific day. For example, news coverage of heatwaves and similar natural disasters can stretch over weeks. Public announcement on policy changes can happen outside of market hours, such as when the EU introduces laws that affect U.S. companies. Moreover, sometimes, news coverage can predate policy changes by writing in anticipation of international summits.

news series. The different colors represent the different news series described above. The red rhombus shows the average among all correlations, and hedging portfolios are sorted top-to-botttom by this value.

Figure 2: Climate Hedge Performance of Various Portfolios

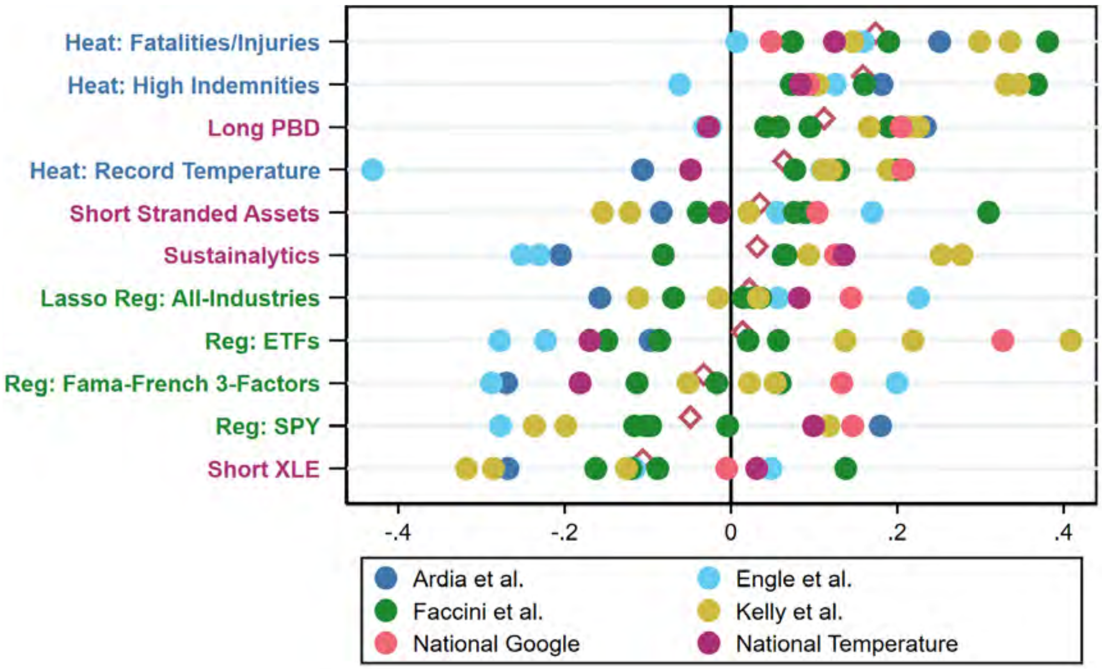

Note: Dot plot of monthly return correlations for various climate hedge portfolios with various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Different colors represent different groups to which the climate news series belong.

Together, Table 9 and Figure 2 summarize the hedging ability of all different methodologies. The first three rows of Table 9 report the hedging performance of the quantity portfolios. These portfolios tend to produce relatively high out-of-sample correlations for a large variety of measures. Figure 2 (blue rows) summarizes these results: the heat indexes appear at the very top of the figure, with almost all correlations in the positive domain. In fact, the "Heat: Fatalities/Injuries" portfolio correlates positively with *all* climate news innovations. "Heat: High Indemnities" can hedge all but the CHNEG series, whereas "Heat: Record Temperature" only fails for the national temperature deviations, WSJ, and MCCC. Given the low correlation among most of the temperature shocks (see Table 1), the high consistency in the results is notable. All heat quantity portfolios provide excellent hedges for Faccini et al. (2021)'s international summits and global warming indices, as well as Kelly (2021)'s general and physical risk indices. This suggests that mutual fund adviser's climate change awareness following an extreme heat event increases both in terms of physical and transitional risk.

Next (rows 4-7 of Table 9), we can study the performance of the narrative approach portfolios. The main advantage of these portfolios is that they do not require estimating

Table 9: Climate Hedge Performance of Various Portfolios

	Faccini, Matin, Skiadopoulos				Kelly [et al.]			Engle et al.		Ardia et al.	Google	Temp.
	IntSummit	GlobWarm	NatDis	Narrative	General	Transitional	Physical	WSJ	CHNEG	MCCC	National	National
Heat: Fatalities/Injuries	**.39**	**.18**	**.05**	**.06**	**.31**	**.14**	**.35**	**.06**	**.13**	**.25**	**.07**	**.11**
Heat: High Indemnities	**.37**	**.16**	**.08**	**.06**	**.34**	**.10**	**.36**	**.16**	−.08	**.18**	**.11**	**.07**
Heat: Record Temperature	**.21**	**.13**	**.18**	**.07**	**.13**	**.19**	**.12**	−.40	**.09**	−.11	**.22**	−.05
Long PBD	**.06**	**.09**	**.19**	**.04**	**.21**	**.17**	**.23**	−.02	−.03	**.23**	**.20**	−.03
Short XLE	−.09	−.16	−.12	**.14**	−.32	−.13	−.29	−.12	**.05**	−.27	−.01	**.03**
Short Stranded Assets	−.04	**.08**	**.31**	**.09**	−.12	**.02**	−.15	**.06**	**.17**	−.08	**.10**	−.01
Sustainalytics	**.13**	−.08	**.06**	**.07**	**.25**	**.28**	**.09**	−.25	−.23	−.20	**.13**	**.14**
Benchmark5	−.02	**.06**	**.06**	−.11	**.05**	−.05	**.02**	**.20**	−.29	−.27	**.13**	−.18
ETF5	−.09	−.15	**.06**	**.02**	**.41**	**.22**	**.14**	−.28	−.22	−.10	**.33**	−.17
allInd5	−.12	**.08**	−.12	**.03**	−.04	−.08	**.09**	**.09**	**.11**	**.08**	**.15**	**.13**
LassoInd5	**.03**	**.01**	−.07	**.04**	**.03**	−.11	−.02	**.06**	**.23**	−.16	**.14**	**.08**

Note: Monthly correlations for various climate hedge portfolios' returns with various climate news series AR(1) innovations. Each row represents a hedge portfolio, whereas each column corresponds to the innovations of a climate news series. Positive correlation coefficients are highlighted in bold. Also, all climate news series are coded such that high numbers indicate negative climate news. Therefore, positive correlation coefficients show successful hedges. While the long, short, Sustainalytics, and quantity portfolios stay constant along the rows, the regression and lasso regression portfolios show in each cell the portfolio that was specifically trained on the respective climate news series.

the portfolio weights from historical data, since the direction of the trade, and the relative weights, are determined on ex-ante information. For example, the first portfolio of the group features a long position in the ETF PBD – the long position is motivated by economic reasoning, in that PBD is a clean energy fund where the natural hedging position is long (with respect to physical risks but most importantly transition risks). The second row is a short position in XLE, motivated by the fact that XLE is dominated by polluting companies; the same reasoning pins down the sign of the trading strategy of the other portfolios in this group. The Table (and the red rows of Figure 2) shows that the performance of these portfolios is mixed, with the worst results given by the short XLE position, and the best results given by the long PBD position. The Sustainalytics portfolio is strong in hedging the Kelly (2021) climate news series but fails in hedging the indices of Engle et al. (2020) and Ardia et al. (2020).

The last four rows of Table 9 report the hedging performance of mimicking portfolios based on aggregate time-series information (see also the green group of rows in Figure 2). The performance of these portfolios varies substantially across climate news series. They all produce reasonable out-of-sample correlations with the Google search innovations, but also display substantial variation with the remaining news series. For example, the portfolio built using the three Fama-French factors has a relatively high correlation of 0.2 with the WSJ index from Engle et al. (2020), in addition to a 0.13 correlation with the Google index. But it also displays a relatively high *negative* correlation with the CHNEG index of -0.29 from Engle et al. (2020), and similarly negative correlations with national temperatures and with MCCC from Ardia et al. (2020). All of the other correlations are close to zero. Note also that mimicking portfolios are estimated separately for each target (that is, there is one mimicking portfolio for the Google search index, one for the national temperatures, and so on). This gives mimicking portfolios additional flexibility compared the other methodologies (which instead do not use information in the target to build the hedging portfolio). Yet, the performance of the approach is mixed.

The results based on the traditional approaches highlight a few important points about the construction of hedging portfolios for climate risks. First, hedging climate risk is hard. We have little historical data available, which makes specifically the mimicking portfolio approach particularly noisy in practice. Second, using climate characteristics to build a hedging portfolio can give encouraging results (PBD is able to hedge out of sample all but 3 news series), especially because it does not require estimating portfolio weights using historical data. However, as the table shows, there is an inherent difficulty in choosing the right climate characteristics, or even the direction of the trade, based only on prior information. Beyond PBD, the other three portfolios in this group do not perform as well, nor consistently across measures. In fact, the short XLE trade – a very natural one ex ante – has the worst performance of all hedging portfolios. Finally, the results vary significantly across climate news series (the targets of the hedge). Many of the hedging portfolios considered so far perform well in hedging the Google searches and the narrative index, whereas for most of the other climate news targets, the performance is mixed and inconsistent, with some approaches better hedging some targets and not others.

Overall, the results show that the quantity approach, which does not rely on ex-ante information about climate exposures, and estimates it based on cross-sectional (rather than time-series) information, delivers superior out-of-sample hedging performance compared to the alternative methods.

3.5 Robustness

We now consider the robustness of our results with respect to a variety of different choices made in building the quantity portfolio. Each row of Figure 3 shows the out-of-sample correlations as in Figure 2, but in this case, for variations of the "Heat:Fatalities/Injuries" index (the best performing of the quantity portfolios).

Figure 3: Climate Hedge Performance - Robustness Tests

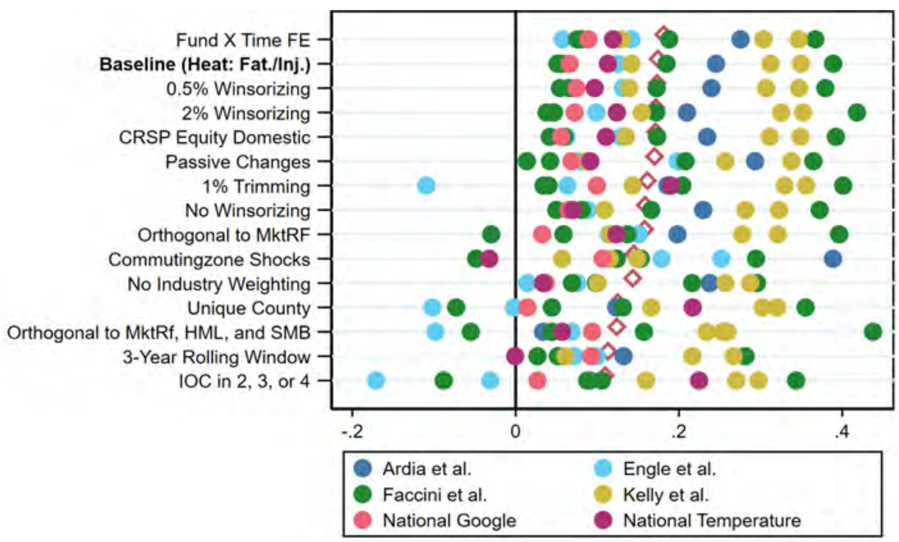

Note: Dot plot of monthly return correlations for various climate hedge portfolios with various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Different colors represent different groups to which the climate news series belong. For more details, refer to Table 9. The climate hedge portfolios are variations of "Heat: Fatalities/Injuries" and, therefore, show numerous robustness tests.

We consider the following variations: add the interaction of fund and time fixed effects to regression 3; define industry exposure changes in terms of current prices (previously defined as *PassiveChanges*); changes in the winsorization (the baseline winsorizes at 1%, we report here the results using trimming instead of winsorizing at 1%, and alternatively winsorizing at 0.5% or 2%, and without winsorization); define the relevant universe of funds with only the CRSP objective code or only the Thomson Reuters IOC; do not weight industry changes by their relative market size; aggregate climate shocks from counties to commuting zones; and lastly, only keep funds where all advisers reside in the same county. As the Figure shows, most of these changes have minimal effects on the performance of the measure. Some of them improve the performance of the portfolio, others deteriorate it, but, overall, the results appear quite robust to these changes.

4 Hedging Macro Factors

While the main focus of this paper is on hedging climate risks, quantity-based portfolios can be built to hedge any other macroeconomic risk. In this section we explore two other applications, to hedge national unemployment and housing shocks. As in the case of climate risks, key to applying our methodology is the availability of "local" shocks that affect investors' beliefs in a way that mirrors the response to the aggregate shocks.

4.1 Macro shocks

We start by building national and local series for unemployment and housing. We obtain national and county-level monthly unemployment figures from the U.S. Bureau of Labor Statistics. We then define unemployment shocks as AR(1) innovations at the quarterly frequency for both our local and global shocks (we include month fixed effects to remove seasonality):

$$Unemp_{t,c} = \theta Unemp_{t-1,c} + \delta_m + \epsilon_{t,c}. \tag{6}$$

Our housing index is the Zillow Home Value Index (ZHVI), which is available at different geographic levels. We define local housing shocks as the AR(1) innovations of the growth rate of the ZHPI price series, at the county level and quarterly frequency. Similarly, global housing price shocks follow the same definition but use the national ZHPI.

$$\Delta Log(ZHPI_{t,c}) = \theta \Delta Log(ZHPI_{t-1,c}) + \epsilon_{t,c}. \tag{7}$$

Intuitively, by using these shocks, our methodology captures the response of mutual funds to unexpected rises in local unemployment or housing prices. We obtain the innovations by applying the regression from 2010 to 2019. To align with the climate sample, we validate the performance from 2015 to 2019 inclusive.

4.2 Validation of the macro hedge portfolios

Applying Regression (3) with the local macro shocks instead of local climate change shocks gives us industry-specific macro betas. We then construct the unemployment and housing price quantity hedge portfolios as in Equation (4). Moreover, for comparison, we construct mimicking portfolios as in Section 3.3.

Table 10 shows the out-of-sample correlation of our constructed macro hedge portfolios with AR(1) innovations of the national unemployment series and the growth rate of the housing price index. Both the Fama-French factor projection and the all-industry lasso regression portfolio achieve favorable correlations with shocks in the national indices. In other words, their returns are high when either unemployment or house prices grow unexpectedly. In comparison, the quantity portfolio achieves the second-best result for the ZHPI index and the best result for the unemployment index, highlighting the great hedging ability of the quantity-based approach.

The table also shows that the quantity portfolio based on local unemployment shocks hedges national unemployment shocks but not housing shocks; and, vice versa, the portfolio based on local housing shocks hedges national housing shocks but not unemployment shocks. This is consistent with the fact that local unemployment shocks move investors' beliefs (and determine their trading) in a way that resembles what happens after a global unemployment shock (but not a housing shock); and similarly for housing shocks.

Table 10: Macro Hedge Performance

	ZHPI	Unemployment
Reg: Fama-French 3-Factors	**.08**	**.17**
Reg: SPY	**.07**	**.12**
Reg: All-Industries	**.15**	−.03
Lasso Reg: All-Industries	**.07**	**.19**
Quantity: ZHPI	**.13**	**.01**
Quantity: Unemployment	**.04**	**.21**

Note: Monthly correlations for various unemployment and housing price hedge portfolios' returns with AR(1) innovations of the national indices. Each row represents a hedge portfolio, whereas each column corresponds to the innovations of either the log ZHPI or the national unemployment index. Positive correlation coefficients are highlighted in bold. While the quantity portfolios stay constant along the rows, the regression and lasso regression portfolios show in each cell the portfolio that was specifically trained on the respective macro series.

As previously argued, we believe one of the main advantages of the quantity-information approach is that fact that it does not rely on having a long time-series of historical data. To further test this claim, we repeat the analysis but train all models, i.e., mimicking and quantity portfolios, on a three-year rolling window (instead of five). A three-year window allows the strategies to adjust faster to structural changes, which could be due to regulatory or technological developments. Table 11 reports the results. Notably, all but the unemployment quantity hedge portfolio perform worse when disciplined only with the previous three years, making the unemployment quantity hedge portfolio the clear leader. While the ability to hedge the ZHPI index in this case is lower, the out of sample correlation is still comparable with the one obtained by the other hedge strategies and, in fact, it is the second-best among the considered options. Overall, the two tables confirm that the quantity-based approach could provide a new tool for building hedging portfolios for a variety of economic risks.

Table 11: Macro Hedge Performance

	ZHPI	Unemployment
Reg: Fama-French 3-Factors (3Y)	**.02**	**.13**
Reg: SPY (3Y)	**.03**	**.05**
Reg: All-Industries (3Y)	**.09**	**.04**
Lasso Reg: All-Industries (3Y)	−.05	−.15
Quantity: ZHPI (3Y)	**.03**	−.01
Quantity: Unemployment (3Y)	−.01	**.22**

Note: A variation of Table 10 where all hedge strategies are based on a three-year rolling window (instead of five).

5 Conclusion and Directions for Future Research

In this paper, we introduce the quantity-based approach to hedging climate change news and macro series innovations. In both cases, the quantity hedge portfolios perform well and often better than more traditional hedging approaches. Not only does this insight open up a new methodology of constructing hedge portfolios for all kinds of risks, but it also reveals three facts on mutual fund investors. First, the systematic reactions to local climate shocks show that mutual funds believe in climate change risk and actively try to mitigate it. Second, because investors situated in afflicted areas behave differently from others, investors rely on their immediate environment to predict nationwide patterns. And third, on average, mutual fund investors are successful at hedging climate change risk and national macro series.

We believe that the quantity-based methodology opens up a promising field of research. Future literature can study the hedging capabilities of other indices or the responses from a different class of investors. In particular, it would be interesting to evaluate the effectiveness of retail investors to hedge various risks. Also, researchers need to find out how to combine the quantity-based methodology most efficiently with traditional approaches.

References

Ardia, D., Bluteau, K., Boudt, K. & Inghelbrecht, K. (2020), 'Climate change concerns and the performance of green versus brown stocks', *National Bank of Belgium, Working Paper Research* (395).

Barnett, M. (2017), Climate change and uncertainty: An asset pricing perspective, Technical report, Working Paper.

Chang, S. (2019), 'Local industry bias in investor behavior: Evidence from mutual funds', *Working Paper* .

Choi, D., Gao, Z. & Jiang, W. (2020), 'Attention to Global Warming', *The Review of Financial Studies* **33**(3), 1112–1145.

Deryugina, T. (2013), 'How do people update? the effects of local weather fluctuations on beliefs about global warming', *Climatic Change* **118**(2), 397–416.

Egan, P. J. & Mullin, M. (2012), 'Turning personal experience into political attitudes: The effect of local weather on americans? perceptions about global warming', *The Journal of Politics* **74**(3), 796–809.

Engle, R. F., Giglio, S., Kelly, B., Lee, H. & Stroebel, J. (2020), 'Hedging climate change news', *The Review of Financial Studies* **33**(3), 1184–1216.

Faccini, R., Matin, R. & Skiadopoulos, G. (2021), 'Are climate change risks priced in the us stock market?', *Working Paper* .

Fownes, J. & Allred, S. (2019), 'Testing the Influence of Recent Weather on Perceptions of Personal Experience with Climate Change and Extreme Weather in New York State', *Weather, Climate, and Society* **11**(1), 143–157.

Giglio, S., Kelly, B. & Stroebel, J. (2021), 'Climate finance'.

Giglio, S., Maggiori, M., Stroebel, J. & Utkus, S. (2021), 'Five facts about beliefs and portfolios', *American Economic Review* **111**(5), 1481–1522.

Hoepner, A. G., Oikonomou, I., Sautner, Z., Starks, L. T. & Zhou, X. (2018), 'Esg shareholder engagement and downside risk'.

Joireman, J., Truelove, H. B. & Duell, B. (2010), 'Effect of outdoor temperature, heat primes and anchoring on belief in global warming', *Journal of Environmental Psychology* **30**(4), 358–367.

Jung, H., Engle, R. & Berner, R. (2021), Climate stress testing, Technical report, Working Paper.

Kelly, B. (2021), ‘Tbd’.

Krueger, P., Sautner, Z. & Starks, L. T. (2020), ‘The importance of climate risks for institutional investors’, *The Review of Financial Studies* **33**(3), 1067–1111.

Kuchler, T. & Zafar, B. (2019), ‘Personal experiences and expectations about aggregate outcomes’, *The Journal of Finance* **74**(5), 2491–2542.

Lamont, O. A. (2001), ‘Economic tracking portfolios’, *Journal of Econometrics* **105**(1), 161–184.

Li, Y., Johnson, E. & Zaval, L. (2011), ‘Local warming: daily temperature change influences belief in global warming’, *Psychological Science* **22**(4), 454–459.

Pástor, L., Stambaugh, R. F. & Taylor, L. A. (2020), ‘Sustainable investing in equilibrium’, *Journal of Financial Economics* .

Sisco, M. R., Bosetti, V. & Weber, E. U. (2017), ‘When do extreme weather events generate attention to climate change?’, *Climatic Change* **143**(1), 227–241.

Song, Y. (2020), ‘The mismatch between mutual fund scale and skill’, *The Journal of Finance* **75**(5), 2555–2589.

Stephens-Davidowitz, S. (2014), ‘The cost of racial animus on a black candidate: Evidence using google search data’, *Journal of Public Economics* **118**, 26–40.

Stroebel, J. & Wurgler, J. (2021), ‘What do you think about climate finance?’.

A Appendix

A.1 Tables

Table A.1: Historical Industry Climate-β Coefficients (2010-2015)

GICS	Description	Avg.	Fatalities/Injuries	Indemnities	Record Temp.
2510	Auto & Components	0.06	0.01	0.13	0.08
2520	Consum. Durables & Apparel	0.05	0.03	0.04	0.13
4520	Tech. Hardw. & Equip.	0.04	0.03	0.07	0.05
4530	Semiconductors & Equip.	0.03	0.05	0.00	0.01
4510	Software & Services	0.02	0.01	0.06	0.01
2530	Consumer Services	0.01	0.03	−0.03	0.02
4010	Banks	0.01	0.01	−0.01	0.02
1510	Materials	0.01	0.03	−0.00	−0.01
1010	Energy	0.00	0.01	0.00	−0.00
2030	Transportation	0.00	0.00	−0.02	0.01
5020	Media & Entertainment	0.00	−0.01	0.02	0.01
2010	Capital Goods	0.00	0.02	−0.03	−0.00
2020	Commercial & Prof. Serv.	−0.00	−0.00	0.00	0.01
3510	Health Care Equip. & Serv.	−0.00	−0.00	−0.03	0.04
4020	Diversified Financials.	−0.00	−0.00	0.02	−0.03
4030	Insurance	−0.01	−0.02	−0.02	0.02
3030	Household & Pers. Prod.	−0.01	−0.01	−0.02	0.01
3520	Pharma., Biotech., & Life Sc.	−0.01	−0.04	0.01	0.02
3020	Food, Bev. & Tobacco	−0.01	−0.03	−0.04	0.06
3010	Food & Staples Retailing	−0.01	−0.00	−0.07	0.02
2550	Retailing	−0.02	−0.05	−0.00	0.02
6010	Real Estate	−0.02	−0.04	−0.06	0.05
5010	Communication Services	−0.02	−0.03	−0.01	−0.01
5510	Utilities	−0.03	−0.04	−0.06	0.00

Note: Industry climate beta coefficients as in equation (3). The coefficients are sorted by the average coefficient and are based on data from 2010 to 2014 inclusive.

A.2 Figures

Figure A.1: Distribution of "Heat: Fatalities or Injuries"

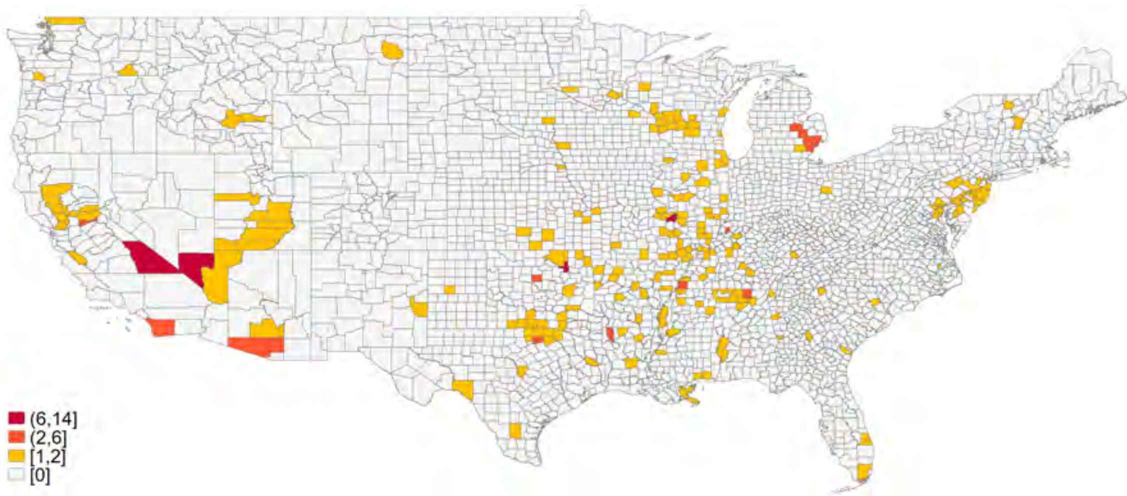

Note: Distribution of the "Heat: Fatalities or Injuries" climate shock from 2010 to 2019. The color-coding shows the number of county-months that experienced the "local" climate shock during the interval.

Figure A.2: Distribution of "Heat: High Indemnities"

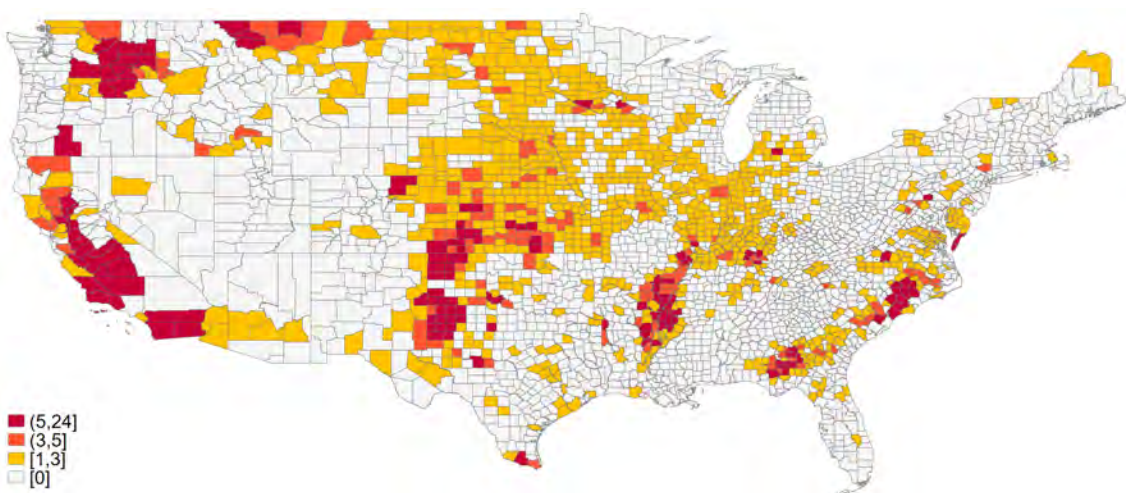

Note: Distribution of the "Heat: High Indemnities" climate shock from 2010 to 2019. The color-coding shows the number of county-months that experienced the "local" climate shock during the interval.

Figure A.3: Distribution of "Heat: Record Temperature"

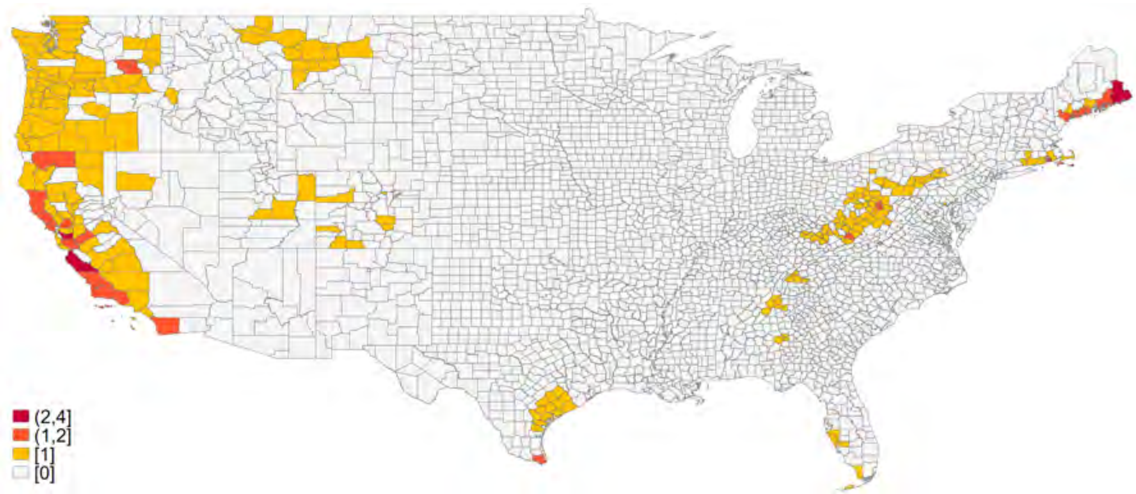

Note: Distribution of the "Heat: Record Temperature" climate shock from 2010 to 2019. The color-coding shows the number of county-months that experienced the "local" climate shock during the interval.

Figure A.4: Climate Hedge Performance - 3-Year Rolling Windows

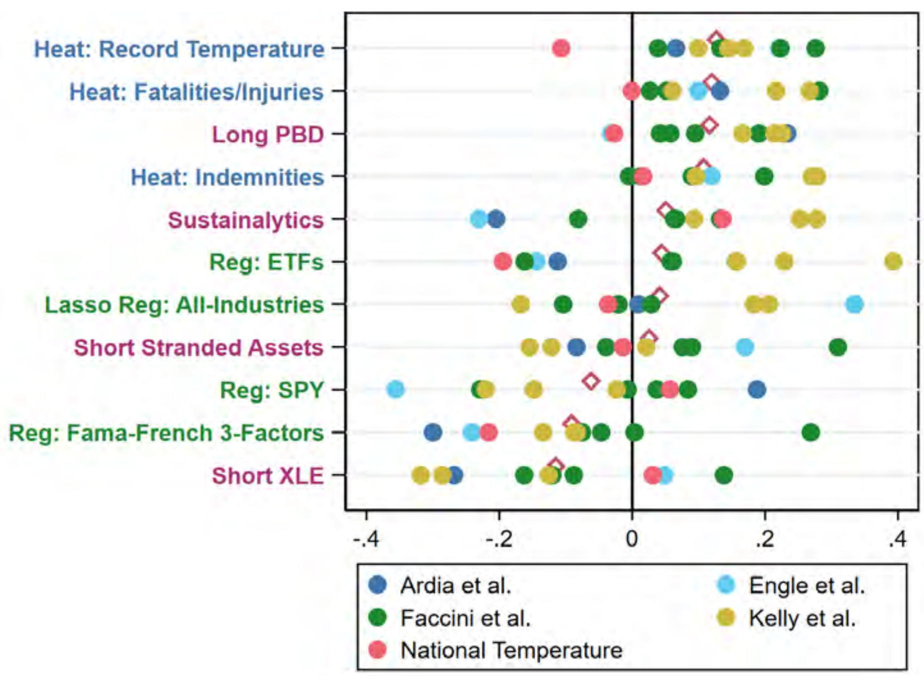

Note: Dot plot of monthly return correlations using 3-year rolling windows; see Table 9.

Figure A.5: Climate Hedge Performance - "Heat: Indemnities" Robustness

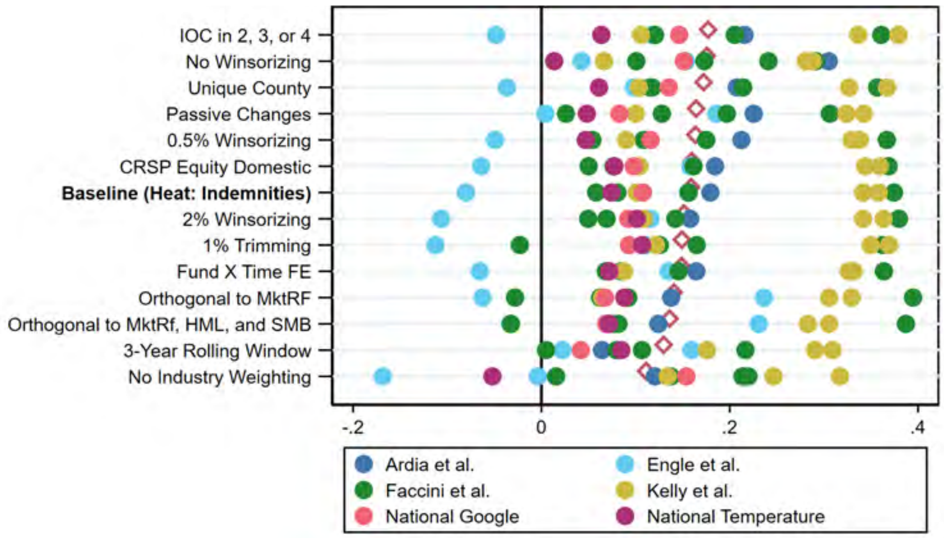

Note: Dot plot of monthly return correlations for various climate hedge portfolios with various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Different colors represent different groups to which the climate news series belong. For more details, refer to Table 9. The climate hedge portfolios are variations of "Heat: Indemnities" and, therefore, show numerous robustness tests.

Figure A.6: Climate Hedge Performance - "Heat: Record Temperature" Robustness

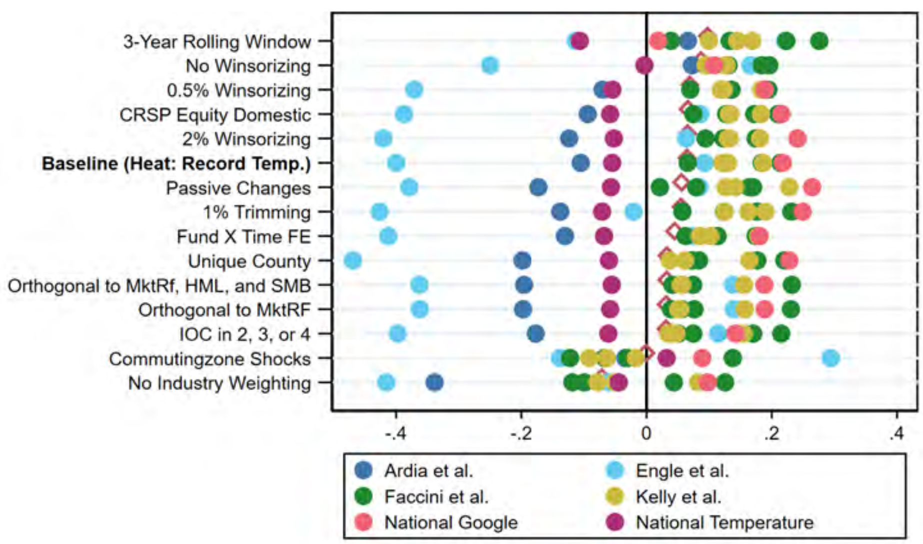

Note: Dot plot of monthly return correlations for various climate hedge portfolios with various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Different colors represent different groups to which the climate news series belong. For more details, refer to Table 9. The climate hedge portfolios are variations of "Heat: Record Temperature" and, therefore, show numerous robustness tests.