

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 the economic and financial risks from climate change. Our *quantity-based* approach exploits information on how mutual fund managers trade in response to idiosyncratic changes in their climate risk beliefs. We exploit two types of idiosyncratic belief shocks: (i) instances when fund advisers experience local extreme heat events that are known to shift climate change beliefs, and (ii) instances when fund managers change the language in shareholder disclosures to express concerns about climate risks. We use the funds' observed portfolio changes around such idiosyncratic belief shocks to predict how investors will reallocate their capital in response to aggregate climate news shocks that shift the beliefs and asset demands of many investors and thus move equilibrium prices. We show that a portfolio that is long stocks that investors tend to buy after experiencing negative idiosyncratic climate belief shocks, and short stocks that investors tend to sell, appreciates in value in periods with negative aggregate climate news shocks. Our quantity-based portfolios have superior out-of-sample hedge performance compared to portfolios constructed using existing alternative methods. The key advantage of the quantity-based approach is that it learns from rich cross-sectional trading responses rather than time-series price information, which is particularly limited in the case of newly emerging risks such as those from climate change. We also demonstrate the versatility of the quantity-based approach by constructing successful hedge portfolios for aggregate unemployment and house price risk.

*This version: November 14, 2022. We thank Daron Acemoglu, Michael Barnett, Richard Berner, Lauren Cohen, Robert Engle, Nicola Gennaioli, Itay Goldstein, Wei Jiang, Annette Vissing-Jorgensen, Dexin Zhou, and seminar and conference participants at the NBER Summer Institute, SFS Cavalcade, WFA, NYU Stern, Columbia University, Cornell, MIT Sloan, New York Fed, Dallas Fed, San Francisco Fed, Federal Reserve Board, Central Bank of Chile, UCLA, UC Boulder, Fordham University, DIW, UNC-Duke, PRI, Hong Kong University, Bundesbank, UCLA, NYU Shanghai Volatility and Risk Institute conference, Johns Hopkins University, 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 and the Center for Sustainable Business at NYU Stern. We are grateful to Xuran Zeng for excellent research assistance.

[†]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 societal challenge. 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 & Stroebele 2021](#), [Stroebele & Wurgler 2021](#)). At present, only a few instruments are designed to directly hedge various climate risks, most prominently the relatively illiquid “catastrophe bonds” ([Tomunen 2021](#)). However, investors interested in hedging climate risks can still build hedge portfolios using other assets, such as stocks or bonds, that are exposed to climate risks. To do so, investors need to identify which assets would benefit and which would lose from the occurrence of a climate risk realization. A long-short portfolio that buys the former and sells the latter would increase in value when climate risk materializes, thus providing a valuable climate risk hedge (e.g., [Engle et al. 2020](#)).

The finance literature has proposed various approaches to building hedge portfolios for macro risks, the most prominent of which is the mimicking portfolio approach of [Lamont \(2001\)](#). The existing approaches typically rely on the availability of a long time series: the risk exposures of different assets—and thus the choice of which assets to buy and sell in the hedge portfolio—are inferred based on the historical comovement between asset prices and realizations of the hedge target. This makes existing approaches poorly suited in settings when the targeted risks are new or materialize infrequently, as in the case of climate risk.

In this paper we propose a new methodology to build portfolios to hedge newly emerging risks such as those from climate change. Our *quantity-based* approach uses cross-sectional information on investors’ trading activity to identify which stocks to hold in the hedge portfolio. The approach starts by identifying “idiosyncratic belief shocks”, which are shocks that shift individual investors’ beliefs about climate risk, but only affect a small group of individuals at a time. While such shocks do not move asset prices (they are, after all, idiosyncratic), they can still influence the affected investors’ trading activity. Based on this insight, the quantity-based approach systematically explores how investors *trade* in response to these idiosyncratic belief shocks, thus learning how investors’ demand for each asset is shifted by changes in perceived climate risks. For example, it identifies which stocks investors tend to buy and sell after a negative change in their climate beliefs. The hedge portfolio is built by going long the former and short the latter.

This portfolio is expected to rise in *price* when aggregate climate risk materializes. The reason is that, while idiosyncratic shocks only move quantities and not prices, the occurrence of an aggregate climate shock affects many investors. As long as investors respond to the aggregate climate risk shock similarly to how they respond to the idiosyncratic shock, the correlated shift in demand of many investors will move prices. Therefore, the hedge portfolio will rise in value when aggregate climate risks materialize.

We operationalize this quantity-based approach by building portfolios of U.S. stocks to hedge climate risk. To identify the positions in the hedge portfolio, we study mutual funds, an important group of investors that publicly report their portfolio holdings each quarter. We propose two ways to identify idiosyncratic shocks to the climate risk beliefs of mutual fund

managers. The first exploits geographically localized extreme heat events that have been shown in prior work to affect beliefs about aggregate climate risk (see, e.g., Egan & Mullin 2012, Deryugina 2013, Joireman et al. 2010, Li et al. 2011, Fownes & Allred 2019, Sisco et al. 2017). We consider three different measures of extreme local heat shocks, each affecting only a limited number of investors: the presence of fatalities and injuries due to extreme heat, the extent of crop indemnity payments due to extreme heat, and the occurrence of extreme temperatures relative to historical patterns in a location. We confirm that the occurrence of these shocks in an area leads to increased Google searches for the term “climate change”, providing evidence that these heat shocks indeed induce updates in climate beliefs.

The second approach to identifying idiosyncratic shocks to mutual fund managers’ climate risk beliefs is based on changes in the discussion of transition risk—the risk stemming from government responses to climate risks—in the mutual funds’ shareholder reports. This approach directly identifies a change in beliefs about climate risks at the individual fund level, without focusing on why these beliefs have changed.

We then study how U.S. active non-sector mutual funds change their portfolio allocation across industries when their managers experience one of these idiosyncratic climate belief shocks.¹ The four belief shock measures are at most weakly correlated with each other: different heat shocks tend to affect some areas more than others, and mutual funds’ disclosures are updated in response to many events other than heat shocks that also affect managers’ climate beliefs in idiosyncratic ways. Interestingly, however, we find that the industries that funds buy and sell in response to these four shocks are significantly correlated. In other words, fund managers that experience one type of idiosyncratic climate belief shock buy and sell similar industries as managers that experience a different type of shock. This finding suggests that the portfolio adjustments contain useful and consistent signals about the climate risk exposures of the different industries.

There are several interesting patterns that emerge by looking at which industries are bought or sold in response to climate belief shocks. For example, by the end of 2019, the auto industry was the industry with the strongest positive quantity response (that is, mutual funds tend to buy auto stocks after increasing concerns regarding climate risks). While this finding may appear surprising at first glance—automobiles are, after all, an important source of carbon emissions—the managers’ reaction may reflect their beliefs that the transition to electric vehicles provides substantial opportunities for incumbent car makers to sell more new vehicles over the coming years. We also find that investors tend to buy insurance companies in response to belief shocks, potentially because, in a world with heightened climate risks, insurance companies may face increased profits due to higher demand for insurance.

Of course, we cannot know with certainty what factors determine fund managers’ assessments of the various industries’ climate risk exposures. Indeed, if we had a good understanding of different industries’ exposures *ex ante*, the construction of hedge portfolios

¹We do not focus on trading in individual stocks, because the large number of stocks and the relatively sparse holding matrix would imply significant estimation error in each stock’s climate risk exposure. However, from a conceptual perspective, our approach expands to considering individual stocks as well as other asset classes, as long as holdings changes can be systematically observed.

would not require elaborate approaches. However, our approach does not require such a full understanding of the economic determinants behind each industries’ climate risk exposure. Instead, to build a hedge portfolio based on the observed quantity information, we rely on the consistency of fund managers’ behavior in response to idiosyncratic shocks and aggregate shocks. Said differently, we need investors to change their demands for different industries in the same way following idiosyncratic and aggregate climate belief shocks.

We validate this assumption in a variety of ways. First, we show that the industry-level trading activity in response to idiosyncratic shocks is similar across periods in our sample as well as across different investors. Second, we show that, while our different belief shocks are not correlated across investors, the corresponding quantity-based hedging portfolios are similar, indicating that investors trade in consistent ways in response to different climate belief shocks. Most importantly, we show that the cross-sectional quantity information is useful for predicting industry price responses to aggregate climate risk shocks. To do this, we use the trading of mutual funds in response to managers’ idiosyncratic belief shocks to build a long-short industry portfolio and study its out-of-sample hedging performance with respect to various measures of *aggregate* climate risk. We build separate hedge portfolios using our four idiosyncratic belief shocks (three based on heat and one based on transition risk disclosures), and we also build hedge portfolios based on an aggregation of the four belief shocks and based on an aggregation of the three heat-related shocks. Eventually, we evaluate their performance against alternative approaches for constructing hedge portfolios that have been proposed in the literature.

The first alternative approach—which we call the “narrative” approach—chooses long and short positions based on economic reasoning. For example, such an approach might suggest buying clean energy companies, selling coal companies, or buying companies with high ESG scores as in [Engle et al. \(2020\)](#), [Pástor et al. \(2021\)](#), and [Hoepner et al. \(2018\)](#). This approach will hedge climate risk if the underlying economic intuition is aligned with that of the average investor. This approach has the advantage that, like the quantity-based approach, it does not require long time series to be implemented; however, it requires having correct priors about investors’ perceptions of each industry’s exposures to climate risk.

The second alternative approach is the “mimicking portfolio” approach of [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’ return performance during past climate risk realizations.

We assess the hedging performance of our quantity-based portfolios and the various alternatives by computing the out-of-sample correlations between monthly portfolio returns and measures of aggregate climate shocks between 2015 and 2019. For the mimicking portfolio approach and the quantity-based approach, we construct the hedge portfolios using rolling five-year windows of price and quantity data, respectively.² To evaluate the hedging perfor-

²Prior 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. Consistent with this assessment, we find that none of the approaches can hedge climate news between 2000-2010.

mance with respect to aggregate risks, we explore a range of measures of aggregate climate shocks as hedge targets, 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 the portfolios perform in hedging various series constructed by [Engle et al. \(2020\)](#), [Faccini et al. \(2021\)](#), [Ardia et al. \(2020\)](#), and [Kelly \(2021\)](#), as well as national temperature shocks and attention to climate risk as measured through Google searches.

We document several patterns. First, at a broad level, hedging climate risks is hard, and few approaches manage to achieve more than a 20% out-of-sample correlation with the climate shock series, confirming and extending this finding from [Engle et al. \(2020\)](#). Second, both the mimicking portfolio approach and the narrative approach generally provide mixed results: they appear to provide decent hedges for some measures of aggregate climate risks, and bad hedges (often with negative out-of-sample correlations) for other measures. Third, our quantity-based portfolios have significantly better average out-of-sample hedging performance compared to the alternatives. Specifically, our quantity-based portfolios constructed using the aggregation of all four idiosyncratic shocks, the aggregation of three heat shocks, and the heat fatalities or injuries all yield positive out-of-sample correlations with *all* of our aggregate climate shock series, with maximum correlations of above 35%. The other quantity-based portfolios do almost as well, and almost all of them dominate the alternative approaches. This validates the idea that the cross-sectional information on which the quantity portfolios are based is useful for hedging aggregate climate news shocks.³

In addition to documenting the strengths of our quantity-based methodology, our empirical results highlight some 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 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: will the broader stock market 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 the mimicking portfolio approach to constructing successful climate hedges. Adding more base assets in the construction of the mimicking portfolio can help to better target the hedge, but requires estimating more parameters, again an issue in short samples. Narrative-based portfolios are immune to such short-sample issues, since historical data is not used to determine positions. However, deciding on positions in an a priori way is hard: as an example, for many industries, the different co-authors of this paper would have ex ante picked rather different holdings.

³Given data constraints, we cannot evaluate the *long-run* hedging performance of these portfolios for actual realizations of climate risks. For any of these approaches to work in the long run, the investor needs to ultimately identify the different assets' climate risk exposures correctly. The narrative approach requires the investor's economic reasoning to be correct, whereas the mimicking portfolio and quantity approach rely on the *average* investor being correct. That said, the latter approaches can still provide good hedging of negative climate news in the short run if investors are wrong in their assessments of assets' climate risk exposures but consistently so over time. We devote Section 4 to an in-depth discussion of this issue.

In the data, we find that seemingly plausible narrative portfolios have very different out-of-sample hedging properties. For example, buying clean energy stocks provides a solid hedge against the arrival of negative climate risk news, but shorting traditional energy companies provides a bad hedge, displaying negative correlations with climate risk news.

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, since climate risks have only recently attracted investors’ attention. As a result, there is not enough time-series data to allow researchers to precisely estimate the climate risk exposures of different assets based on price data alone. However, our approach can, in principle, be applied to hedging any macro risk series for which similar idiosyncratic shocks (e.g., stemming from local events or measurable through investor disclosures) affect investors’ beliefs about aggregate risks. For example, in recent work, [Kuchler & Zafar \(2019\)](#) show that locally experienced house price movements affect expectations about future U.S.-wide house price changes; they also show that personally experienced unemployment affects beliefs about the future national unemployment rate (also see [Bailey et al. 2018, 2019, 2020](#)). Consistent with our results on hedging climate risks, we show that the trading responses of mutual fund investors to local house price and unemployment shocks allow us to construct portfolios that perform well at hedging innovations in the corresponding national series.

Our work contributes to a growing literature that studies the interaction between climate change and asset markets (see [Giglio, Kelly & Stroebe1 2021](#), for a recent review). In equity markets, [Bolton & Kacperczyk \(2021a\)](#) and [Hsu et al. \(2022\)](#) document that high-pollution firms are valued at a discount. [Engle et al. \(2020\)](#) find that stocks of firms with lower exposure to regulatory climate risk experience higher returns when there is negative news about climate change, and [Barnett \(2020\)](#) shows that increases in the likelihood of future climate policy action lead to decreased equity prices for firms with high exposure to climate policy risk. [Choi et al. \(2020\)](#) document that the stocks of carbon-intensive firms underperform during periods of abnormally warm weather, where investors’ attention to climate risks is likely heightened. Climate risk is also priced in other asset classes such as real estate markets ([Baldauf et al. 2020](#), [Bakkensen & Barrage 2022](#), [Bernstein et al. 2019](#), [Giglio, Maggiori, Rao, Stroebe1 & Weber 2021](#), [Murfin & Spiegel 2020](#)) and municipal bond markets ([Painter 2020](#), [Goldsmith-Pinkham et al. 2021](#), [Acharya et al. 2022](#)).

Our quantity-based approach to forming hedge portfolios builds on prior work that studies how individuals form beliefs based on their personal experiences (e.g., [D’Acunto et al. 2022](#), [Kuchler & Zafar 2019](#), [Malmendier & Nagel 2011](#), [Alok et al. 2020](#), [Busse et al. 2015](#), [Chang et al. 2018](#)) and how such beliefs translate into actions ([Armona et al. 2019](#), [Armantier et al. 2015](#), [Bachmann et al. 2015](#), [Bailey et al. 2018, 2019](#), [Gennaioli et al. 2016](#), [Giglio, Maggiori, Stroebe1 & Utkus 2021a,b](#), [Roth & Wohlfart 2020](#)). Our approach also relates to a recent literature using quantity and holdings data in asset pricing (e.g., [Berk & van Binsbergen 2016](#), [Koijen & Yogo 2019](#)). We contribute to this literature by providing evidence that quantity information is useful for predicting price movements in response to aggregate shocks.

1 Quantity-Based Portfolios: A Simple Model

In this section, we describe a simple model that motivates our quantity-based approach to forming hedge portfolios.

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

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

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 idiosyncratic component $\omega_A(i)$, such that $\epsilon_A(i) = \nu_A + \omega_A(i)$. The common belief ν_A is driven by shocks or news that are observed by all investors, and that correspond to the types of news that investors might want to hedge (e.g., well-publicized news about accelerating global warming that shifts all investors' beliefs about physical climate risks). The idiosyncratic belief component $\omega_A(i)$ can instead be affected by "local" events that are only observed or experienced by investor i (e.g., a localized heat wave where investor i lives that impacts her views on climate risks). We do not impose assumptions on the origins of the common and idiosyncratic components of beliefs. There is no learning from prices about the beliefs or information of other investors; investors simply "agree to disagree".

Idiosyncratic Belief Shocks. We first study changes in equilibrium prices and quantities in response to an idiosyncratic shock $\omega_A(i)$, for example because investor i —having experienced a localized heat wave—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,

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

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 experiences an idiosyncratic change in her beliefs, her equilibrium allocation changes. However, since the shock only affects one (atomistic) investor, it 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 Quantities to Prices. Suppose now there is news about stock A that affects all investors’ beliefs (i.e., a change in ν_A). For example, an announcement by a senior politician might cause all investors to believe that climate change regulation has become more likely, reducing the expected profitability of heavily-emitting 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, $\int_{i=0}^{i=1} \frac{\partial q_A}{\partial \epsilon_A(i)} di$.⁴ Together with the earlier result, this shows how idiosyncratic quantity responses can help 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 affects all investors. Aggregate news then moves the demand function of many investors simultaneously, leading to price movements in response to aggregate shocks. Therefore, a hedging portfolio built using idiosyncratic quantity data can hedge aggregate climate shocks.

2 Idiosyncratic Belief Shocks & Portfolio Changes

Our quantity-based approach to constructing climate risk hedge portfolios requires identifying ‘idiosyncratic belief shocks’ that satisfy three criteria. First, the shocks should shift asset demands of affected investors by influencing their beliefs about climate risks (or their attention to these risks). Second, the shocks should only affect a small number of investors, so that they influence those investors’ portfolios without inducing a large price response. Third, changes in asset demand following the idiosyncratic belief shocks should predict changes in asset demand following aggregate news about climate risk, the events we are trying to hedge.

We have identified two types of idiosyncratic belief shocks that satisfy these criteria. The first type of shock builds on an extensive literature that identifies local extreme heat events as important drivers of climate change attention and beliefs in the affected populations (e.g., Egan & Mullin 2012, Deryugina 2013, Joireman et al. 2010, Li et al. 2011, Fownes & Allred 2019, Sisco et al. 2017). We construct several measures of extreme heat shocks at the county level and show that they indeed affect attention to climate risks as measured through Google searches. Our measures of extreme heat are sufficiently concentrated geographically to only affect the beliefs of a small subset of investors located in the affected counties.

The second type of shock is based on changes in the coverage of transition risks (i.e., climate risks that affect firms through regulation and other government responses) in mutual funds’ semi-annual shareholder reports. In contrast to the local heat shocks, which represent shifters of idiosyncratic climate risk beliefs, this disclosure-based approach attempts to

⁴The constant of proportionality can vary across securities. For example, the same quantity response may induce a larger price effect for stocks with smaller market capitalizations. To incorporate such effects in our empirical application, we estimate quantity responses relative to market capitalization; for a more structural approach, see Koijen et al. (2020).

directly measure the changes in the climate beliefs of different investors, without identifying why a particular investor may have changed her beliefs.

We next provide details on the construction of these two types of idiosyncratic belief shocks, before exploring investor trading responses to these shocks.

2.1 Idiosyncratic Belief Shocks: Extreme Heat Events

There are several plausible ways to identify extreme heat events as potential shifters of climate change attention and beliefs of the affected populations. In this paper we consider three shocks that we describe below: heat waves that involve fatalities or injuries; heat waves that induce large crop indemnity payments; and local temperature outliers. Table 1 provides an overview of the frequency of these shocks, and the maps in Appendix Figures A.1 to A.3 visualize the geographic distributions of the events.

Fatalities or Injuries from Extreme Heat. Our first extreme heat shock captures whether there were any fatalities or injuries due to extreme heat in a 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 0.1% of all county-months in the U.S. between 2010 and 2019 had fatalities or injuries due to extreme heat.

Crop Indemnity Payments due to Extreme Heat. We construct a second measure of extreme heat shocks from crop indemnity payments. The underlying data are collected by the U.S. Department of Agriculture, and we use a version maintained by SHELDUS.⁵ We normalize the crop indemnity payments by the number of acres reported as being planted adjusted by the insured’s share in the commodity. An extreme heat event is identified when the monthly normalized heat-related crop indemnity payments in a given county exceed the 99.9th percentile across all county-months in the preceding 10 years; about 0.21% 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 extremely low across our panel of county-months. 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 broadly independent of shocks due to fatalities and injuries.

Extreme Temperatures. While the extreme heat shocks described above capture the most devastating events—events that involve very high absolute levels of temperatures—they do not necessarily capture all instances where temperatures are high relative to normal patterns in colder regions. We thus construct a third county-level heat shock measure using

⁵Crop 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 Agriculture reports these payments for several private insurance companies, covering more than 100 crops.

Table 1: Summary Statistics on Extreme Heat Measures

<i>Panel A: Local Shocks: Summary</i>				
Climate Shock	Event Description	Frequency		
		Monthly	Sample	
Heat: Fatalities/Injuries	Injuries or fatalities	0.10%	1.32%	
Heat: High Indemnity Payments	99.9th percentile indemnity payments	0.21%	0.08%	
Heat: Extr. Temperature	Monthly max temp. 4°C above avg max	0.22%	1.21%	
Report: Transition Risk	Change in fund disclosures about transition risk	-	0.46%	

<i>Panel B: Local Shocks: Monthly Jaccard Correlations</i>				
	Fatalities/Injuries	High Indemnity Payments	Extr. Temperature	
Heat: Fatalities/Injuries	1.00			
Heat: High Indemnity Payments	0.00	1.00		
Heat: Extr. Temperature	0.01	0.00	1.00	

<i>Panel C: Local Shocks: Sample Jaccard Correlations</i>				
	Fat./Inj.	High Ind. Pay.	Extr. Temp.	Trans. Risk
Heat: Fatalities/Injuries	1.00			
Heat: High Indemnity Payments	0.01	1.00		
Heat: Extr. Temperature	0.04	0.01	1.00	
Report: Transition Risk	0.00	0.00	0.00	1.00

Note: Panel A provides an overview of the idiosyncratic belief shock measures. The “monthly” frequency shows the share of county-month observations in the U.S. from 2010 to 2019 that experience the event. The “sample” frequency shows the share of observations in our final sample (at the fund-quarter level) that experience the shock. Panel B shows the Jaccard correlations between the constructed heat measures across all county-months from 2010 to 2019, Panel C shows the Jaccard correlations among the shock measures across all fund-quarters in our final sample.

temperature data from the PRISM Climate Group. In particular, we flag county-months with a maximum temperature of at least 4 degrees Celsius (7.2 degrees Fahrenheit) above the county’s ten-year historical average maximum for the same month. We enforce that this maximum temperature is above 32 degrees Celsius (90 degrees Fahrenheit), which is the threshold for “extreme caution” by the U.S. National Weather Service, to maintain a sense of severity. About 0.22% of all county-months have such an extreme heat event. The “extreme temperature” events are more evenly spread across the United States; they are also not highly correlated with the previous two measures of extreme heat events.

Heat Shocks and Climate Change Attention and Beliefs. We next explore whether our local heat shock measures affect local climate change attention or beliefs, as measured by Google searches for the term “climate change” (see [Stephens-Davidowitz 2014](#), [Choi et al. 2020](#), for similar approaches). Since Google Trends data are not available at the county level and are often missing at the MSA level, we conduct this analysis at the state-month level, aggregating our measure of heat shocks to the state level.⁶

⁶A state is recorded to experience an “extreme 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

The Google search series measures relative interest in a topic, such as the fraction of all Google searches in a region for “climate change.” In every period, Google scales the relative search interest for a topic 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 search index using the following specification:⁷

$$\log(\widetilde{G}_{t,s}) = \beta_S S_{t,s} + \delta_s + \gamma_t + \epsilon_{t,s}, \quad (1)$$

where $\widetilde{G}_{t,s}$ is the observed (scaled) Google search interest for climate change in state s at time t , and $S_{t,s}$ is the corresponding indicator for a local extreme heat event. State and time fixed effects are captured by δ_s and γ_t .

Table 2: Heat Shocks and Climate Attention

	Log(Google Search Volume)		
Heat: Fatalities/Injuries	0.05** (0.03)		
Heat: High Indemnity Payments	0.05* (0.02)		
Extreme Temperature	0.05** (0.02)		
R^2	0.77	0.77	0.78
State & Month FE	Y	Y	Y
N	5,506	5,506	4,693

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

Table 2 reports the coefficients β_S from running this regression separately for the different indicators of extreme local heat corresponding to $S_{t,s}$. All coefficients are positive and statistically significant. Intuitively, experiencing any fatalities or injuries from heat, high reported crop indemnity payments and record temperatures are all associated with an increase in relative interest in climate change of 5%. These findings highlight that all of the extreme heat measures affect climate change awareness or beliefs.

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, and to using more continuous measures, such as injuries or fatalities per capita.

⁷Let $G_{t,s}$ 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.

2.2 Idiosyncratic Belief Shocks: Investor Disclosures

In addition to using extreme heat events as shifters of investors’ climate risk concerns, we also directly measure investor concerns about climate transition risks from mutual fund managers’ disclosures in their semi-annual shareholder reports (filed as N-CSR reports with the Securities and Exchange Commission). This disclosure-based measure allows us to capture shifts in climate change awareness or beliefs that could be driven by forces other than extreme heat events. In addition, this measure is not susceptible to possible concerns that local extreme heat events could affect investment decisions for reasons other than changing beliefs about climate risk (see, e.g., [Heyes & Saberian 2019](#)).

We first extract sentences from these reports that use terms such as climate change, carbon emission(s), greenhouse gas(es), and global warming. These terms were selected based on their cosine similarity to “climate change” in [Google’s word2vec implementation](#). We manually classify these sentences to determine whether they express concerns about or considerations driven by transition risk. For example, we identify the following passages:

Finally, we added two new holdings in the natural gas segment based on the growing concern, and ultimately legislated increased cost, associated with carbon emissions. [[Scout Funds, 2009Q4 N-CSR filing](#)]

Still, there were plenty of factors supporting investments in nuclear energy, especially as [...] concerns over global warming gathered steam [...]. [[Market Vectors ETF Trust, 2010Q3](#)]

Additionally, an international climate change deal signed in December may boost investment in clean energy sources such as solar. [[Nuveen Investments, 2015Q4](#)]

Overall, we identify 87 reports (corresponding to 803 fund-quarters, as a report can cover multiple funds) as expressing concerns about or considerations related to transition risk.

2.3 Holdings and Location Data

There are a number of reasons why mutual fund managers are a natural focus of our study of trading responses to idiosyncratic climate belief shocks. Mutual funds make up a substantial share of the investor universe, and their portfolio holdings are observable, at least at the quarterly frequency (see [Chen et al. 2010](#), [Frazzini & Lamont 2008](#), [Grinblatt & Titman 1989](#), [Wermers et al. 2012](#), for other uses of this data). In addition, as we describe below, we can observe the locations of the mutual funds’ advisers, which allows us to link the portfolio holdings to the occurrence of local heat shocks. Furthermore, mutual fund managers regularly report their strategic considerations to their shareholders, allowing us to construct our second measure of idiosyncratic belief shocks.

For our approach to work, mutual fund managers must adjust their portfolios in response to perceived changes in climate risk. Such portfolio adjustments could occur for various reasons. Managers may believe that equilibrium valuations do not fully account for climate risks and that they can thus earn alpha by reducing their holdings of more exposed stocks

(Krueger et al. 2020). Alternatively, managers may view climate change as an additional risk to hedge, either because of their investors’ preferences (Ceccarelli et al. 2021) or to manage flows in response to national climate events (Dou et al. 2022). Notably, these motivations all generate similar predictions for investors’ responses to idiosyncratic belief shocks; when managers become more concerned about climate risk, they sell stocks which are more exposed to climate risk, and buy those that are less exposed.

Portfolio holdings data. We use the Thomson Reuters Mutual Fund Holdings S12 database to obtain a panel of portfolio holdings of U.S. mutual funds. We combine the holdings data with fund characteristics from CRSP.⁸ Most funds report their holdings every three months, and our analysis 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 3,⁹ which focuses the universe of assets for our hedging portfolio on North American common stocks. Since we wish to identify deliberate fund manager asset reallocations in response to idiosyncratic belief shocks, we restrict our analysis to actively managed funds, i.e., those with an Investment Objective Code of 2 (“Aggressive Growth”), 3 (“Growth”), 4 (“Growth & Income”), or missing, and that have “Equity Domestic Non-Sector” as their CRSP Objective Code (see Song 2020).

We obtain stock-level characteristics from CRSP and Compustat and assign end-of-month prices from CRSP to the holdings. We obtain firm GICS industry codes from Compustat by merging the stocks on their CUSIP identifiers. The first four GICS digits determine the stock’s classification into the 24 “industry groups” that are the main focus of our analysis.¹⁰

Measuring Active Portfolio Changes. In our main analysis, we explore how idiosyncratic climate belief shocks induce changes in the portfolio share of fund f in industry I *through active trading* between consecutive holdings reports. Since holdings are usually reported at three-month intervals (often, though not always, at the end of a quarter), we measure fund composition changes over such intervals. We perform our analysis at the industry level, since the sparsity of the stock-level holding matrix would lead to very noisy estimates of stock-level exposures. For every fund f and month t with a holdings report, we define the active change in industry I holdings as:

$$ActiveChanges_{f,t}^I = \left[\left(\frac{\sum_{j \in I} P_{j,t-3} S_{f,j,t}}{\sum_j P_{j,t-3} S_{f,j,t}} \right) - \left(\frac{\sum_{j \in I} P_{j,t-3} S_{f,j,t-3}}{\sum_j P_{j,t-3} S_{f,j,t-3}} \right) \right] \frac{1}{(Share_t^I)}. \quad (2)$$

⁸We link mutual funds across Thomson Reuters and CRSP using their Wharton Financial Institution Center Number (WFICN) as reported in WRDS MFLINKS.

⁹These 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.

¹⁰The Global Industry Classification Standard (GICS) is developed by MSCI and S&P based on [earnings and market perception in combination with revenues](#) to classify companies.

$P_{j,t-3}$ is the price for stock j at the end of month $t-3$, the time of the prior report. $S_{f,j,t}$ is the number of shares of stock j held by fund f at the end of month t , and $Share_t^I$ is the market capitalization of industry I as a share of the U.S. stock market. The term in square brackets thus captures the active three-month change of the share of industry I in fund f 's portfolio.¹¹ The reason for scaling by industry size is that a given change 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. We winsorize the active changes measure at the 1% level to mitigate the effect of outliers due to, for example, fund mandate changes.

Investor location data. We also obtain data on the location of mutual fund advisers, which are primarily responsible for making asset-allocation decisions (see [Chang 2019](#)). Specifically, we parse adviser locations from funds' SEC filings (N-SAR filings until 2017, N-CEN filings from 2018 onward). Since SEC filings cannot be matched directly with Thomson Reuters or CRSP mutual fund data, we apply a fuzzy string matching algorithm to match SEC filings with mutual funds. We focus on near-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 2019.¹²

For 67.6% of funds, all advisers reside in the same county. For extreme heat shocks, whenever funds have multiple advisers who are not all located in the same county, we assign fund-level climate 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 local extreme heat shock, we assume the fund is affected by $2/3$ of a local extreme heat shock; our results are robust to alternative aggregation choices.

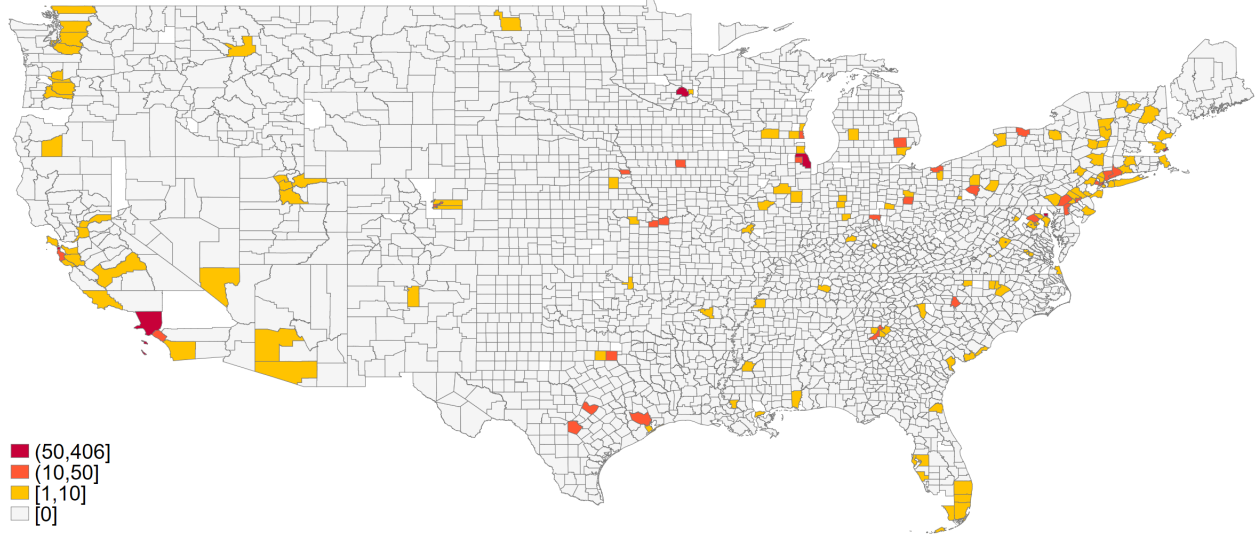
Figure 1 shows the geographic distribution of fund advisers for the subset of funds where all advisers reside in the same county. While some areas have a larger concentration of advisers, advisers are generally spread throughout the entire country. The table shows that about a quarter of advisers are located in New York (most of them in New York City), 14.3% are located in Massachusetts (most of them in Boston), and 10.3% are located in California (roughly equally split between San Francisco and Los Angeles). This gives us important geographical variation and therefore differential exposure to local extreme heat shocks.

Summary Statistics. Panels A and B of Table 3 present summary statistics on the GICS industries and the portfolio holdings of the mutual funds in our final sample. In an average quarter, funds in our sample held 224 unique companies in the Energy sector (GICS code 1010). In the average sample-month, the energy sector accounted for 7.2% of fund holdings. The smallest industry by average market capitalization was "Auto & Components", comprising of an average of 43 firms with an average market capitalization share of 0.9%. On

¹¹ We verify below that our approach is robust to considering a separate variable, $PassiveChanges_{f,t}^I$, where the first fraction uses $P_{j,t}$ instead of $P_{j,t-3}$, 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 constantly rebalanced their portfolios.

¹²As we describe in more detail below, a fund-quarter observation requires two consecutive holding reports spaced three months apart, allowing us to analyze the active trading of mutual funds over the period.

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. Panel A of the table 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 the map and the two panels are based on the subset of funds whose advisers all reside in the same location.

Table 3: Sample Summary Statistics

<i>Panel A: Industry Summary Statistics</i>							
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

<i>Panel B: Mutual Fund Summary Statistics</i>						
	Number of Companies			Number of Industries		
	Avg.	p10	p90	Avg.	p10	p90
Mutual Fund Holdings	209	33	467	19.5	14.0	24.0

<i>Panel C: Active Changes Summary Statistics</i>						
	Avg.	p1	p25	p50	p75	p99
Active Industry Change	-0.00	-1.26	-0.06	0.00	0.05	1.32

Note: Panel A shows, for 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 holdings 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 measure as defined in Equation (2). The unit of observation is a fund-quarter-industry change and the sample size is 1,156,344.

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 $ActiveChanges_{f,t}^I$. Intuitively, active changes of 0 imply that there were no active changes in industry I 's relative weight within fund f 's portfolio, while active changes of 1 imply that I 's weight in the portfolio increased by a percentage equal to the industry's market share. For example, if industry I makes up 10% of the market, and the fund increased I holdings from 5% to 15% of the portfolio, then $ActiveChanges^I$ would be 1. The average and median active changes in our sample are zero. The first percentile is -1.26 corresponding to decreases in an industry's portfolio share that are larger in percentage terms than that industry's market share.

2.4 Estimating the Response to Idiosyncratic Climate Shocks

To understand how portfolio composition changes with idiosyncratic climate belief shocks, we estimate the following panel regression separately for each industry I :

$$ActiveChanges_{f,t}^I = \beta_t^{I,S} S_{f,t} + \delta_t^I + \epsilon_{f,t}^I, \quad (3)$$

where $ActiveChanges_{f,t}^I$ is defined as in Equation 2 and δ_t^I captures month fixed-effects. $S_{f,t}$ is one of the four idiosyncratic belief shocks described in Section 2.1 or one of the two composite shocks described below.

For the three local extreme heat shocks, $S_{f,t}$ is non-zero whenever there was an extreme heat shock in the fund advisers' locations in at least one of the three months during which we measure the active portfolio changes (i.e., months t , $t-1$, and $t-2$).

For the disclosure-based approach, we assign reports expressing transition risk concerns a score of "1". In our baseline analysis, we use the *change* of this measure between consecutive reports to measure idiosyncratic changes in beliefs about climate transition risks. We assign $S_{f,t}$ a value of "+1" if the fund's concern about climate risks increase between reports (e.g., when transition risk is mentioned for the first time), a value of "0" if the concern is unchanged (e.g., it is not mentioned or keeps being mentioned over time); and a value of "-1" the fund's concern about climate risks decreases between reports. Our approach is robust to using alternative ways of aggregating this text-based information.¹³ To align the semi-annual N-CSR reports and the quarterly holding changes, we set $S_{f,t}$ equal to "+1" / "0" / "-1" whenever at least one of the months during which we measure the active portfolio changes lies between subsequent N-CSR reports.

We also consider two composite shocks. The "Pooled: Heat" shock is constructed as the maximum value across the three heat shocks. The "Pooled: All" shock is constructed as the sum of the "Pooled: Heat" shock and the disclosure-based shock.

The main objects of interest are the climate-shock-specific estimates of $\beta^{I,S}$: for each industry I , this represents the differential active change in fund holdings of that industry for

¹³In our baseline analysis below, we also do not exploit changes in this measure for mutual funds that are specifically sustainability-focused; while managers from these funds frequently discuss climate risks, those risks are likely a key component in portfolio construction throughout our sample.

funds affected by idiosyncratic belief shock S , relative to the change in holdings for funds that were not affected by such shocks in the same quarter. We refer to this coefficient as the *industry-specific climate quantity beta*. Estimates of $\beta^{I,S}$ can vary over time with the sample over which regression 3 is estimated.

Table 4 reports the estimated $\beta^{I,S}$ coefficients, for each of the six shocks. The table reports the estimates obtained at the end of 2019 using five years of backward-looking data. Within each shock, we scale the coefficients such that they range between -1 and +1, allowing us to also present average values across four idiosyncratic belief shocks. Industries towards the top of the table are those that mutual funds disproportionately buy after receiving idiosyncratic climate belief shocks, while industries towards the bottom of the table are those that investors disproportionately sell (relative to the funds that do not receive the shock in that quarter). Our interpretation of the estimates in Table 4 is that they indicate the relative climate risk exposures of various industries as perceived by mutual fund investors.

Interpreting these estimates is not trivial. Each industry’s exposure to climate risk is determined by a variety of economic mechanisms, which are difficult to anticipate. Relative to the narrative approach (that takes a strong stand on what those mechanisms and exposures are), our quantity-based approach aggregates information from market participants (specifically, mutual funds) to help identify these industry exposures.

Ultimately, a crucial question is whether this procedure ends up yielding successful hedge portfolios. Before documenting that the resulting portfolios do indeed provide comparatively successful climate hedges—suggesting that the local quantity responses are successful at predicting price responses to aggregate climate news shocks—we first discuss possible economic mechanisms behind some of the estimated quantity betas.

The largest positive average climate quantity responses are found in the “Auto & Components” sector. Since automobiles produce a large share of current carbon emissions (Ritchie et al. 2020), one might have thought that tighter limits on those emissions would pose a risk to automotive firms and that their valuations would thus decline upon news of increased physical or transition risks. On the other hand, the auto sector is at the forefront of the technological transition towards a green economy, with electric vehicles playing an important role in decreasing carbon emissions. A spike in demand for those vehicles from households and firms hoping to reduce their carbon footprints, combined with substantial global subsidies for electric vehicle purchases, have the potential to cause a faster-than-usual replacement of existing vehicles with electric vehicles. As a result, it is plausible that a faster transition will increase global automotive sales (and ultimately profits) for a sustained period of time. This would benefit not only the focused electric vehicle makers such as Tesla, but also traditional incumbents such as Ford and General Motors as well as their suppliers. This sentiment is reflected in numerous equity analyst reports that we reviewed, and is reflected in headlines such as “General Motors is a buy as its transition to electric vehicles gains steam, Berenberg says”, “General Motors’ EV Plans Present ‘Golden Opportunity’: Wedbush Analyst Says”, and “Ford Stock Set to Benefit From Electric Vehicle Push”.

Similarly, one might have thought the insurance sector to be negatively exposed to news about realizations of physical climate risk, due to higher insurance claims following natu-

Table 4: Industry-Specific Climate Quantity Betas

GICS	Description	Avg.	Fat./Inj.	Ind.	Extr.	Tran.	Pool. All	Pool. Heat
2510	Auto & Components	0.54	1.00	0.30	0.69	0.18	1.00	1.00
4520	Tech. Hardw. & Equip.	0.51	0.74	0.71	1.00	-0.39	0.51	0.87
4530	Semiconductors & Equip.	0.30	0.71	0.27	0.37	-0.14	0.73	0.70
3030	Household & Pers. Prod.	0.29	0.38	0.10	0.45	0.21	0.54	0.42
5010	Communication Services	0.24	0.67	-0.01	0.01	0.31	0.39	0.58
2010	Capital Goods	0.24	0.27	0.02	0.21	0.45	0.55	0.30
3020	Food, Bev. & Tobacco	0.21	0.34	0.01	0.23	0.25	0.42	0.33
5510	Utilities	0.21	0.35	0.15	0.51	-0.17	0.46	0.42
4510	Software & Services	0.18	0.38	0.23	0.23	-0.14	0.34	0.40
4020	Diversified Financials.	0.15	0.47	0.28	0.02	-0.16	0.42	0.44
4010	Banks	0.12	0.65	0.13	0.01	-0.30	0.55	0.59
1010	Energy	0.12	0.49	0.30	-0.11	-0.21	0.44	0.46
3010	Food & Staples Retailing	0.10	0.58	-0.46	0.33	-0.05	0.71	0.48
4030	Insurance	0.08	-0.07	0.12	0.18	0.11	0.26	0.01
3520	Pharma., Biotech., & Life	0.07	0.34	-0.12	0.16	-0.08	0.31	0.30
2520	Consu. Durables & Appa	0.03	0.50	-0.76	-0.64	1.00	0.37	-0.03
5020	Media & Entertainment	-0.00	-0.13	0.05	0.40	-0.33	-0.13	-0.01
2530	Consumer Services	-0.06	-0.65	0.11	0.30	-0.01	-0.41	-0.45
3510	Health Care Equip. & Serv.	-0.06	0.03	1.00	-0.28	-1.00	-0.11	0.14
2030	Transportation	-0.07	0.49	0.50	-0.63	-0.64	0.14	0.42
1510	Materials	-0.13	-0.00	-1.00	0.46	-0.00	0.23	0.13
2550	Retailing	-0.20	-0.44	0.23	0.22	-0.80	-0.44	-0.28
6010	Real Estate	-0.25	-0.08	-0.94	0.38	-0.37	-0.11	-0.10
2020	Commercial & Prof. Serv.	-0.64	-1.00	-0.10	-1.00	-0.44	-1.00	-1.00

Note: Industry-specific climate quantity betas as in Equation (3). The coefficients are sorted by the average coefficient across the four individual shocks and are based on data from 2015 to 2019 inclusive.

ral disasters. However, if insurance companies are able to adjust premia appropriately—something aided by the fact that most policies reprice annually—increased physical climate risk might not impact expected profits. Indeed, by increasing the overall demand for insurance—something predicted by [Lot & Haegeli \(2021\)](#)—an increase in physical climate risk might even raise insurance firms’ profitability, as limited insurance capital strengthens the pricing power of incumbents. This mechanism was highlighted by McKinsey analysts in [Grimaldi et al. \(2020\)](#), who wrote that:

[Property & casualty insurers] can use the annual policy cycle and their sophisticated understanding of evolving risks to reprice and rearrange portfolios to avoid long-term exposure to climate events. And the growth in the value at risk—and possibly volatility—should increase the demand for new and different insurance solutions and services, which, in turn, could expand the industry’s opportunities.

Even the traditional energy sector, which one might have expected to be among the primary losers from a negative realization of climate transition risk, given its reliance on polluting fossil fuels ([van Benthem et al. 2022](#)), actually displays positive quantity betas for four of the six measures. Several mechanisms might be consistent with such an exposure, in particular given that concerns about “stranded assets” have been reflected in prices for some time.

First, fear of tighter future regulation can discourage entry from competitors. The increased market power can raise the present discounted value of incumbents’ profits from selling hydrocarbons while renewable alternatives remain unreliable, even if a faster transition might reduce industry-wide profits in a calculation that includes potential entrants (see [Ryan 2012](#), [Elliott 2022](#), [Magolin & Santino 2022](#)). Second, large energy companies play an important role in innovation in the clean energy space (see [Pickl 2019](#), [Cohen et al. 2020](#)).

These discussions highlight the difficulties of predicting industry exposures to climate shocks using a narrative approach. Instead, our quantity-based approach relies on the “wisdom of the crowd” by using mutual fund managers’ trading response to infer their average narratives about which stocks would gain or lose in response to climate shocks.

Correlation Across Measures of Idiosyncratic Belief Shocks. An interesting pattern in Table 4 is that the ordering of the industry quantity betas is broadly correlated across different measures of climate belief shocks, despite the fact that the raw correlations of the shock measures are close to zero (see Table 1). Panels A and B of Table 5 formally test this by reporting the correlation and rank-correlation, respectively, of the industry-specific climate loadings obtained from running the regression in Equation 3 for the period of 2015-2019.

The industry-specific climate quantity betas are generally correlated across the various shocks, though they are more consistent within the heat-based shocks than between heat-based and disclosure-based shocks (perhaps because disclosure-based shocks only measure belief shifts about transition risks, while heat-based shocks also capture shifts in beliefs about physical risks). This indicates that mutual funds change their portfolios in a broadly consistent way in response to these different climate belief shocks. Of course, the correlations are far from perfect, in part due to non-trivial noise in the estimation. Ultimately, the way to assess whether our procedure is effective despite this noise is to verify that it generates portfolios that successfully hedge aggregate climate risk, as we do in the next section.

Estimate Stability Across Subsamples. One way to investigate the magnitude of estimation error is to look at the stability of the quantity coefficients with respect to the same idiosyncratic shock measure across subsamples. We explore this by randomly splitting the sample into two mutually exclusive subsamples, and looking at how the β^I estimates correlate across these subsamples. We repeat the random splitting 100 times, and report the average correlations in Table 6. The two panels of the table split the sample in different ways. Panel A 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. All approaches except for the high indemnity payments achieve substantial average coefficient correlations, ranging from 0.19 to 0.42, indicating that the sample consistently picks up a common signal.

Panel B of Table 6 reports the average correlations from randomly splitting the sample either by funds, periods, or counties. For example, with the fund split, each fund fully belongs to one of the two selected mutually exclusive subsamples. Again, the positive correlations highlight that our β^I coefficients reflect significant signals in addition to estimation error.

Table 5: Across-Shock Correlations of Industry-Specific Climate Quantity Betas

<i>Panel A: Pearson Industry Climate Beta Correlations</i>						
	Fat./Inj.	High Ind. Pay.	Extr. Temp.	Trans. Risk	Pooled: All	Pooled: Heat
Heat: Fat./Inj.	1.00					
Heat: High Ind. Pay.	0.17	1.00				
Heat: Extr. Temp.	0.33	0.03	1.00			
Report: Trans. Risk	0.32	−0.48	0.06	1.00		
Pooled: All	0.93	0.06	0.48	0.46	1.00	
Pooled: Heat	0.96	0.31	0.52	0.16	0.92	1.00

<i>Panel B: Spearman (Rank) Industry Climate Beta Correlations</i>						
	Fat./Inj.	High Ind. Pay.	Extr. Temp.	Trans. Risk	Pooled: All	Pooled: Heat
Heat: Fat./Inj.	1.00					
Heat: High Ind. Pay.	0.35	1.00				
Heat: Extr. Temp.	0.09	0.01	1.00			
Report: Trans. Risk	0.23	−0.46	0.12	1.00		
Pooled: All	0.79	0.20	0.32	0.44	1.00	
Pooled: Heat	0.91	0.47	0.23	0.12	0.84	1.00

Note: Panel A shows the Pearson correlations among the industry-specific climate quantity betas. Similarly, Panel B shows the Spearman *rank* correlation among the industry-specific climate quantity betas. The coefficients are based on estimating Equation 3 using data from 2015 to 2019 inclusive.

Finally, we investigate the stability of industry climate coefficients over time. Appendix Figure A.4 shows the industry ranking based on data for each (rolling) five-year window between 2010 to 2019; low numbers correspond to the biggest “long” positions in the implied hedge portfolio. The automobile, technological hardware, and software and services industries are generally long positions over the entire period, but there are also notable shifts in estimated industry exposures over time. Such shifts can occur for multiple reasons. Obviously, each estimate includes non-trivial estimation noise which inevitably moves around our estimates of each industry’s climate risk exposure. However, there are also numerous fundamental explanations for changing industry exposures. For example, the focus of government policies changes frequently, which could affect industries differentially due to shifting tax and subsidy policies. Similarly, industries’ exposures might change due to strategic adaption along the transition path. For example, many traditional fossil fuel firms now have substantial investments in renewable energies (van Benthem et al. 2022), reducing their transition risk exposure. The quantity-based approach has the potential to learn these changing underlying exposures faster, since each period delivers multiple data points through providing cross-sectional information. As before, the ultimate test of whether changing portfolio weights broadly correspond to actual changes in the underlying climate risk exposures or whether they are largely the result of estimation noise involves assessing the ability of the resulting portfolios to hedge aggregate climate risks, which we turn to next.

Table 6: Across-Sample Split Correlations of Industry-Specific Climate Quantity Betas*Panel A: Random Split within Groups*

Climate Shock	Stratified		Fully Random	
	Spearman	Pearson	Spearman	Pearson
Heat: Fatalities/Injuries	0.34	0.36	0.33	0.36
Heat: High Indemnity Payments	−0.07	−0.21	−0.00	−0.14
Heat: Extr. Temperature	0.33	0.39	0.33	0.42
Report: Transition Risk	0.35	0.37	0.36	0.37
Pooled: Heat	0.19	0.23	0.20	0.23
Pooled: All	0.30	0.25	0.27	0.22

Panel B: Random Split between Groups

Climate Shock	Fund Split		Period Split		Location Split	
	Spearman	Pearson	Spearman	Pearson	Spearman	Pearson
Heat: Fatalities/Injuries	0.31	0.36	0.15	0.23	0.33	0.34
Heat: High Indemnity Payments	0.10	0.18	−0.05	−0.36	−0.13	−0.06
Heat: Extr. Temperature	0.34	0.43	0.17	0.25	0.32	0.38
Report: Transition Risk	0.31	0.31	0.02	0.06	0.02	−0.03
Pooled: Heat	0.17	0.20	0.06	0.05	0.37	0.47
Pooled: All	0.25	0.22	0.11	0.05	0.36	0.39

Note: This table shows the average Spearman (rank) and Pearson correlations of the industry-specific climate quantity betas 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 period split, the quarters are split randomly without considering the chronological order; each period fully belongs to exactly one of the two mutually exclusive subsamples.

3 Quantity-Based Climate Hedging Portfolios

We next describe how we use the estimated climate quantity betas to build our climate hedge portfolios. Then, we evaluate the out-of-sample hedging performance of these portfolios, and compare this performance against that of other approaches proposed in the literature.

3.1 Portfolio Construction and Description

We build our hedging portfolio for each time t by estimating $\hat{\beta}_t^{I,S}$ as described in the previous section, using data from the five years prior to t . We compute excess returns of each of the 24 industries as Equation 4, where R_t^I is the value-weighted industry return and R_t^f denotes the risk-free rate. We use the estimated $\hat{\beta}_t^{I,S}$ as the portfolio weights. (Since each component of the portfolio is an excess return, we do not need to scale the $\hat{\beta}_t^{I,S}$, since they do not need to sum to one). The excess return of the quantity-based hedge portfolio will be:

$$QP_t^S = \sum_I \hat{\beta}_{t-1}^{I,S} (R_t^I - R_t^f), \quad (4)$$

Panel A of Table 7 shows the correlations of monthly returns among the quantity-based hedge portfolios based on the different idiosyncratic climate belief shocks. Given that the idiosyncratic shocks are largely uncorrelated (see Table 1), the high correlations in the return series provide additional evidence that our various shocks are picking up a common signal.

Table 7: Portfolio Return Correlations

<i>Panel A: Pearson Portfolio Return Correlations</i>						
	Fat./Inj.	High Ind. Pay.	Extr. Temp.	Trans. Risk	Pooled: All	Pooled: Heat
Heat: Fatalities/Injuries	1.00					
Heat: High Ind. Pay.	0.23	1.00				
Heat: Extr. Temp.	0.51	0.15	1.00			
Report: Transition Risk	0.51	−0.14	0.03	1.00		
Pooled: All	0.92	0.18	0.64	0.55	1.00	
Pooled: Heat	0.93	0.26	0.68	0.39	0.95	1.00

<i>Panel B: Orthogonalized to Fama-French Three-Factors</i>						
	Fat./Inj.	High Ind. Pay.	Extr. Temp.	Trans. Risk	Pooled: All	Pooled: Heat
Heat: Fatalities/Injuries	1.00					
Heat: High Ind. Pay.	0.21	1.00				
Heat: Extr. Temp.	0.56	0.23	1.00			
Report: Transition Risk	0.23	−0.55	0.06	1.00		
Pool. All	0.86	0.16	0.72	0.36	1.00	
Pool. Heat	0.89	0.28	0.74	0.13	0.93	1.00

Note: Panel A shows the monthly return correlations—constructed as in Equation 4—for the period 2015 to 2019 among our quantity-based hedge portfolios. Panel B shows the corresponding return correlations after orthogonalizing each portfolio with respect to the Fama-French market, size, and value factors.

We next investigate how much of the portfolio return correlations are driven by a potential common loading of the different quantity-based hedge portfolios on standard asset pricing factors (here, we study specifically the three Fama-French factors). To identify the factor loadings of the quantity portfolios, we regress the portfolio excess returns on the returns of the market, size, and value factors:

$$QP_t^S = \alpha + \beta_c^M (R_t^M - R_t^f) + \beta_c^{SMB} SMB_t + \beta_c^{HML} HML_t + \epsilon_{t,S}. \quad (5)$$

Table 8 shows the regression results. Four of the six portfolios have a significant loading on the market, five of the six portfolios have a significant loading on the value, and all portfolios have no loading on size (note that the magnitude of these exposures is not interpretable since the scale of the quantity portfolio is arbitrary). Overall, the time-series variation in the Fama-French factors captures on average around 40% of the variation in the quantity portfolios.

Panel B of Table 7 shows the correlations among the returns of the quantity-based hedge portfolios after orthogonalizing the returns to the three Fama-French factors (i.e., taking the correlations of the residuals from the regression in Equation 5). The correlations are similar to those in Panel A, suggesting that a common loading on the Fama-French factors is not the main driver of the high return correlations across the different hedge portfolios.

Table 8: Factor Exposures of Hedge Portfolios

	Quantity-Based Hedge Portfolio					
	Fat./Inj.	High Ind. Pay.	Extr. Temp.	Trans. Risk	Pooled: All	Pooled: Heat
$R^M - R^f$	0.09** (0.04)	0.10* (0.05)	-0.03 (0.06)	0.19*** (0.04)	0.09** (0.04)	0.07 (0.04)
SMB	0.02 (0.04)	-0.01 (0.05)	-0.01 (0.06)	-0.02 (0.03)	-0.01 (0.04)	0.01 (0.04)
HML	0.21*** (0.04)	-0.03 (0.04)	0.11* (0.05)	0.10*** (0.03)	0.22*** (0.05)	0.20*** (0.04)
Constant	-0.01 (0.11)	0.05 (0.13)	0.02 (0.15)	0.02 (0.09)	-0.01 (0.11)	-0.01 (0.12)
R^2	0.44	0.12	0.09	0.53	0.43	0.34
N	60	60	60	60	60	60

Note: Regression of monthly returns of the quantity-based climate hedge 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 choosing 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 regulations to curb emissions. Different risks may be relevant for different investors, and these risks are imperfectly correlated. In addition, climate change is a long-run threat, and we would thus ideally build portfolios that hedge the long-run realizations of climate risk, something difficult to produce in practice.

To overcome these challenges, [Engle et al. \(2020\)](#) argue that the objective of hedging long-run realizations of a given climate risk can be achieved by constructing a sequence of short-lived hedges against *news* (one-period innovations in expectations) about future realizations of these 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—after all, the right hedge target will vary across investors based on their background risk exposure¹⁴—but instead 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 below. Each measure is signed such that a larger number corresponds to “bad news.” Building on [Engle et al. \(2020\)](#), we use the AR(1) innovations of each climate news series as the targets. For a given climate news series c , we denote these AR(1) innovations in month t as $CC_{c,t}$. The series we consider are the following:

¹⁴For example, one would imagine that the sovereign wealth funds of Norway and the UAE would be particularly interested in hedging transition risks, since the economies they represent are heavily dependent on an ability to sell hydrocarbon fuels. Similarly, one might expect the sovereign wealth fund of Singapore, a relatively low-lying island, to focus more on hedging physical climate risk realizations.

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 one captures the number of news articles in the WSJ dedicated to climate change (broadly assuming that “no news is good news”), the second one 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 news about physical and transition risks. These news indices are monthly. The WSJ index covers the period of February 1984 to June 2017. The CHNEG index covers the period of July 2008 to May 2018.

Ardia et al. (2021). [Ardia et al. \(2020\)](#) build on the WSJ index of [Engle et al. \(2020\)](#) by including additional media outlets and differentiating between positive and negative news. Their daily Media Climate Change Concerns index is available between January 2003 and June 2018. We aggregate to the monthly level by taking the average of the daily news series.

Kelly (2021). [Kelly \(2021\)](#) creates three climate news series that reflect general, physical, and transition risk, respectively. These series are constructed as the product of the number of relevant WSJ articles published in a month and the articles’ sentiment, such that higher levels correspond to more “bad news” in the respective categories. These indices are monthly and cover the period of January 1985 to December 2020.

Faccini et al. (2021). We include four of [Faccini et al. \(2021\)](#)’s climate news indices: international climate summits, global warming, natural disasters, and narrative indices. The international climate summits, global warming, and natural disasters indices measure news coverage of the respective topics; the narrative index is constructed by manually reading and classifying 3,500 articles. The international climate summits and narrative indices capture news about transition 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 the daily frequency between January 2000 and November 2019. We aggregate them to the monthly frequency by taking the average of the daily series.

National Google searches. This climate news series is the national Google search interest in “climate change”, capturing attention paid to climate change and its associated risks by the general population. This index does not differentiate between positive and negative news, and could be associated with various climate risks.

National Temperature Deviations. Just as local extreme temperatures increase local climate change awareness, U.S.-wide extreme heat events have the potential to drive national awareness (e.g. [Barnett \(2017\)](#) shows that monthly temperature innovations from a rolling one-sided Christiano-Fitzgerald bandpass filter induce significant stock market reactions). We therefore include such innovations as one of the climate news series.

3.3 Alternative Approaches to Building Hedge Portfolios

We want to compare the hedging performance of our quantity-based portfolios against that of two alternative approaches to constructing hedge portfolios: the narrative approach and the mimicking portfolio approach. All approaches share the same goal: to be long stocks that do well in periods with unexpectedly bad news about climate risks, and short stocks that do badly in those scenarios. The approaches differ in how they identify those stocks.

Narrative approach. The first alternative approach we consider selects portfolio weights of different assets based on an *ex-ante* view of the possible exposures of those assets to climate risks. To do so, one can identify firm characteristics that are associated with high predicted exposures. For example, one characteristic could be the firms’ environmental scores constructed by ESG data providers, under the prior that high-ESG-score companies will fare better when climate risks materialize (see [Engle et al. 2020](#)). Another example of this approach is grouping energy stocks based on whether they focus on renewable energy or fossil fuels, and then building a long-short portfolio of the two groups, motivated by the different potential regulatory exposure of the two groups to transition risk. Overall, the narrative approach requires identifying *ex ante* the economic forces that determine firms’ climate exposures, something that we argue is quite difficult in practice.

We build several portfolios using such a narrative-based approach. Our first narrative-based portfolio takes positions in all U.S.-listed stocks covered by the Sustainalytics ESG scores: the portfolio’s position in each stock is the stock’s ESG score percentile in each period, minus 50. For example, 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.

Our second narrative-based strategy uses industries to take a directional view. We build 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. XLE’s largest holdings are the two U.S. integrated oil companies, ExxonMobil and Chevron, followed by EOG Resources and ConocoPhillips. We would expect that realizations of climate change news should increase PBD’s returns and decrease XLE’s returns. Therefore, the hedge portfolio would go long PBD and/or short XLE.

Our third narrative-based portfolio is the stranded asset portfolio as in [Jung et al. \(2021\)](#) based on XLE, the VanEck Vectors Coal (KOL), and SDPR S&P 500 (SPY) ETFs, using the following weights: $0.3XLE + 0.7KOL - SPY$. To hedge climate risks, economic reasoning would suggest going short this portfolio.

While the three strategies described above all make intuitive sense, it is hard to take a stand on which should be the most successful hedge for climate risk, as many complex factors play a role in determining climate exposures. The approach we turn to next, the mimicking portfolio approach, uses purely statistical methods to choose the portfolio weights, and does not rely at all on economic priors.

Mimicking portfolio approach. The mimicking portfolio approach combines a pre-determined set of assets into a portfolio that is maximally correlated with a given climate change shock, using historical data. To obtain the mimicking portfolios, we estimate the following regression separately for each climate risk news series:

$$CC_{c,t} = w_c R_t + \epsilon_{c,t}$$

where $CC_{c,t}$ denotes the (mean zero) climate hedge target of type c in month t , w_c 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 five-year rolling window.

We consider different sets of excess returns to build mimicking portfolios. First, we use the market alone (the SPY ETF). A mimicking portfolio built using only one asset is effectively equivalent to studying whether historically a long or a short position in that asset was correlated with climate risk. Second, we use the three Fama-French factors (Market, SMB, and HML). Third, we use the two ETFs described above, PBD and XLE, in combination with the Fama-French factors. Fourth, we add to the Fama-French factors the excess returns of the 24 GICS industry portfolios. Given the short time series available for estimation, we regularize the estimation using 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 proposed portfolios. For the quantity-based and mimicking portfolio approaches, for every month in our testing period of 2015-2019, we construct the portfolios as described above using five-year rolling windows of data. The narrative portfolios are constant over time. We focus on the post-2010 period to train our models, as investors likely paid very little attention to climate risks before 2010. As a result, we do not expect information on prices and quantities from before 2010 to be useful in building hedge portfolios today.¹⁵ 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 in the same month, $CC_{c,t}$.¹⁶

¹⁵To explore this assumption, in Appendix Figure A.5 we attempt to hedge aggregate climate news between 2000 and 2010. Specifically, we compare the performance of the quantity-based and mimicking portfolio approaches. We do not explore the narrative approach, since the PBD ETF and Sustainalytics ESG scores are not available for that period. In the early 2000s, neither the quantity-based approach nor the mimicking portfolio approach were successful at hedging aggregate climate news innovations, consistent with the view that such innovations were not priced during this period. This is consistent with findings from Acharya et al. (2022) that physical climate risks only started being priced in municipal bond markets after 2010.

¹⁶This approach evaluates the hedging ability of the portfolio up to a scaling parameter. Our quantity-based methodology and the narrative approach do not identify the scale 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. 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 several weeks; similarly, news coverage can sometimes predate policy announcements by writing in anticipation of international summits.

Table 9 reports these out-of-sample correlations. Each row in the table represents a different hedge 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 correlations imply successful hedges. The same information is displayed in Figure 2, which reports out-of-sample correlations on the horizontal axis, and has one row for each hedge portfolio. Each point in the dot plot is the out-of-sample correlation coefficient of a hedge portfolio return with one of the climate news series. The different colors represent the different news series described above. The red rhombus shows the unweighted average among all correlations, and portfolios are sorted top-to-bottom by this value.

The first six rows of Table 9 show the hedging performance of the quantity-based climate hedge portfolios. These portfolios tend to produce relatively high out-of-sample correlations for a large variety of climate news series (the blue rows of Figure 2 show the same results).¹⁷ At the top of the figure, the “Pooled: All Shocks”, “Pooled: Heat Shocks”, and “Heat: Fatalities/Injuries” portfolios correlate positively with *all* climate news innovations. “Pooled: All Shocks” has the highest average out-of-sample correlation with our series. The “Heat: Extreme Temperature” and “Report: Transition Risk” portfolios only fail for two out of twelve climate risk series. All quantity-based portfolios provide excellent hedges for Faccini et al. (2021)’s international climate summits and global warming indices, Kelly (2021)’s general and physical risk indices, Engle et al. (2020)’s WSJ and CHNEG indices, and Ardia et al. (2020)’s MCCC index. This suggests that our various quantity-based portfolios perform well in terms of hedging a range of climate risks, spanning both physical and transition risks. Given that our quantity-based approaches are not tailored to hedge specific climate targets, their good performance against a variety of targets suggests that they are providing a hedge against some common component of climate risks that is shared by the measures we consider.

Rows 7-10 of Table 9 (and red rows of Figure 2) show the performance of the different narrative-based portfolios. The main advantage of these portfolios is that they do not require estimating the portfolio weights from historical data, since the direction of the trades is based on ex-ante information and beliefs. For example, the first portfolio features a long position in the ETF PBD, motivated by a belief that a clean energy fund should gain upon transition risk realizations. The second narrative portfolio is a short position in the stranded asset portfolio, motivated by the fact that this portfolio is dominated by polluting companies in the oil sector; the same reasoning pins down the sign of the trading strategy of the remaining portfolios in this group. The hedging performance of these narrative portfolios is mixed, with the worst results given by the short XLE position, which has a negative hedge performance. The best hedging result is given by the long PBD position. The uneven hedging performance highlights just how difficult it is to predict, based only on economic intuition, which stocks will gain or lose in response to climate shocks.


¹⁷When evaluating this out-of-sample hedging performance, it is worth keeping in mind that hedging macroeconomic shocks using stocks is generally difficult. As a reference, when building mimicking portfolios for macro risks using a regularized projection method, Giglio & Xiu (2021) report in-sample R^2 s of 2.25% for industrial production growth and 4.07% for consumption growth at the monthly frequency, corresponding to *in-sample* correlations between the target and the hedging portfolio return of 0.15 and 0.20 respectively.

Table 9: Climate Hedge Performance of Various Portfolios

	Faccini et al.				Kelly			Engle et al.		Ardia et al.	Google	Temp.
	IntSummit	GlobWarm	NatDis	Narrative	General	Transitional	Physical	WSJ	CHNEG	MCCC	National	National
Pooled: All Shocks	.37	.23	.06	.00	.25	.02	.32	.22	.13	.37	.11	.17
Pooled: Heat Shocks	.32	.18	.03	.07	.27	.07	.34	.23	.13	.32	.09	.21
Heat: Fatalities/Injuries	.38	.18	.05	.06	.30	.13	.34	.05	.10	.25	.03	.13
Report: Transition Risk	.19	.14	.31	−.07	.22	.08	.19	.10	.05	.26	.21	−.11
Heat: Extreme Temperature	.38	.11	−.19	.06	.09	−.11	.16	.35	.05	.23	.12	.15
Heat: High Indemnity Payments	.05	.28	−.09	.15	.02	−.07	.12	.15	.04	.30	−.18	.00
Long PBD	.06	.09	.19	.04	.21	.17	.23	−.02	−.03	.23	.20	−.03
Short Stranded Assets	−.04	.08	.31	.09	−.12	.02	−.15	.06	.17	−.08	.10	−.01
Short XLE	−.09	−.16	−.12	.14	−.32	−.13	−.29	−.12	.05	−.27	−.01	.03
Sustainalytics	.13	−.08	.06	.07	.25	.28	.09	−.25	−.23	−.20	.13	.14
Reg: Fama-French Three-Factors	−.02	.06	.06	−.11	.06	−.06	.03	.20	−.27	−.26	.13	−.17
Reg: ETFs + Fama-French	−.09	−.15	.06	.02	.41	.22	.14	−.28	−.22	−.10	.33	−.17
Reg: SPY	−.12	−.10	−.10	−.00	−.20	.12	−.24	−.10	−.28	.18	.15	.10
Lasso Reg: All-Industries + FF	.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 indicate successful hedges. While the narrative and quantity portfolios stay identical along the rows. For the mimicking portfolios, we show in each cell the portfolio that was specifically trained on the respective climate news series.

Figure 2: Climate Hedge Performance of Various Portfolios



Note: Dot plot of monthly out-of-sample 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 of climate news series. The red rhombus shows the unweighted average among all correlations, and portfolios are sorted top-to-bottom by this value.

The remaining 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, but is poor on average. 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 -0.27 from Engle et al. (2020), and similarly negative correlations with national temperatures and with the MCCC index from Ardia et al. (2020). All of the other correlations are close to zero. Note that the mimicking portfolios have a relatively weak hedging performance despite the fact that they are estimated separately for each hedge target, giving them additional flexibility compared to the other methodologies (which instead build a single hedge portfolio for all climate news series).

Overall, the results show that our quantity-based approach to forming hedge portfolios consistently delivers the best out-of-sample climate hedging performance. Among the alternative approaches, with only little historical data available for periods when climate risk was potentially priced, mimicking portfolio approaches do not deliver successful climate hedges. In contrast, the narrative-based approach to building hedging portfolios is poten-

tially promising—for example, PBD is able to hedge all but three news series—especially because it does not require estimating portfolio weights using historical data. However, there is often 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 using the narrative approach do not perform consistently well across measures. In fact, the short XLE position—a very intuitively appealing trade *ex ante*—has the worst performance of all proposed climate hedging portfolios. While more systematic narrative approaches, such as using ESG scores, currently suffer from mixed data signals—[Billio et al. \(2021\)](#) and [Berg et al. \(2022\)](#) highlight the low degree of correlation of ESG ratings by different providers—increased firm-level disclosure requirements, such as those included in the SEC’s proposed rule on the “The Enhancement and Standardization of Climate-Related Disclosures for Investors”, may improve the performance of narrative approaches over time.

3.5 Robustness

We now consider the robustness of our results with respect to a variety of choices made in the construction of the quantity-based climate hedge portfolios. Figure 3 shows out-of-sample correlations similar to Figure 2 for variations of the quantity-based portfolio constructed using all four idiosyncratic belief shocks (Panel A) and the quantity-based portfolio constructed using the three heat shocks (Panel B). Appendix Figure A.6 shows similar robustness checks for hedge portfolios built separately using the four idiosyncratic belief shocks. We consider the following variations to our baseline portfolio construction choices:

- (i) Add the interaction of time fixed effects with fund type or fund family fixed effects in the regression in Equation 3 (“Fund Type \times Time FE”; “Fund Family \times Time FE”);
- (ii) Measure changes in investors’ industry-level portfolio holdings using current prices, defined as *PassiveChanges* in Footnote 11. This allows for portfolio changes to be driven by price changes in addition to active trading (“Passive Changes”);
- (iii) Change how we handle extreme changes in investors’ industry-level portfolio holdings from the the baseline procedure of winsorizing at the 1% level (“1% Trimming”; “0.5% Winsorizing”; “2% Winsorizing”; “No Winsorizing”);
- (iv) Do not weight changes in investors’ industry-level portfolio holdings by the industries’ relative market size (“No Industry Weighting”);
- (v) Change the relevant universe of funds to be defined using only CRSP or Thomson Reuters IOC data (“CRSP Equity Domestic”; “IOC in 2, 3, or 4”), and to exclude ESG funds that are likely to have a large climate risk focus even before receiving idiosyncratic belief shocks (“No ESG”);
- (vi) Only keep funds where all advisers reside in the same county, and for which the geographic allocation of heat shocks is thus easier (“Unique County”); only keep funds


that have been hit by at least one of the four four idiosyncratic belief shocks (“Shocks hit fund ≥ 1 ”);

- (vii) Measure extreme local heat shocks and investor locations at the commuting zone level instead of the county level (“Commuting Zone Shocks”);
- (viii) Use three-year and one-year rolling windows of trading activity to identify industry-level climate quantity betas, instead of five-year rolling windows as in the baseline (“Three-Year Rolling Window”; “One-Year Rolling Window”);
- (ix) Orthogonalize each hedge portfolio with respect to the market factor (“Orthogonal to MktRF”) or the Fama-French market, size and value factors (“Orthogonal to MktRf, HML, and SMB”); and
- (x) For news indices at a higher frequency than monthly, first calculate daily innovations and then aggregate to the monthly level instead of the baseline approach of aggregating to the monthly level before taking innovations (“Aggregate daily innovations”)


Across the different hedge portfolios, most of these changes have little effect on the overall hedge performance, and multiple changes appear to even improve the hedge performance relative to our baseline data construction choices. The exceptions are the move to a one-year rolling window of trading activity to construct climate quantity betas, and the inclusion of fund family \times time fixed effects in the regression in Equation 3. This is perhaps unsurprising. Adding fund family \times time fixed effects removes the across-fund family variation that powers much of our main results. Similarly, moving from five-year rolling windows to one-year rolling windows reduces by 80% the amount of quantity data that our approach can use to identify the industries’ quantity betas. This loss in data appears to outweigh any possible gains from being able to better detect time-varying exposures with one-year rolling windows.

Similarly, going from five-year to three-year rolling windows also reduces the hedge performance somewhat, though by much less than moving to a one-year rolling window does. To explore this further, in Figure 4 we construct all quantity-based and mimicking portfolios using only three-year rolling windows of data while leaving the narrative-based portfolios unchanged. The quantity-based portfolios achieve slightly lower average out-of-sample hedge performance than when constructed using five-year rolling windows, but they generally maintain their positions as the most effective hedge portfolios.

Figure 3: Climate Hedge Performance - Robustness of Portfolio Construction Choices



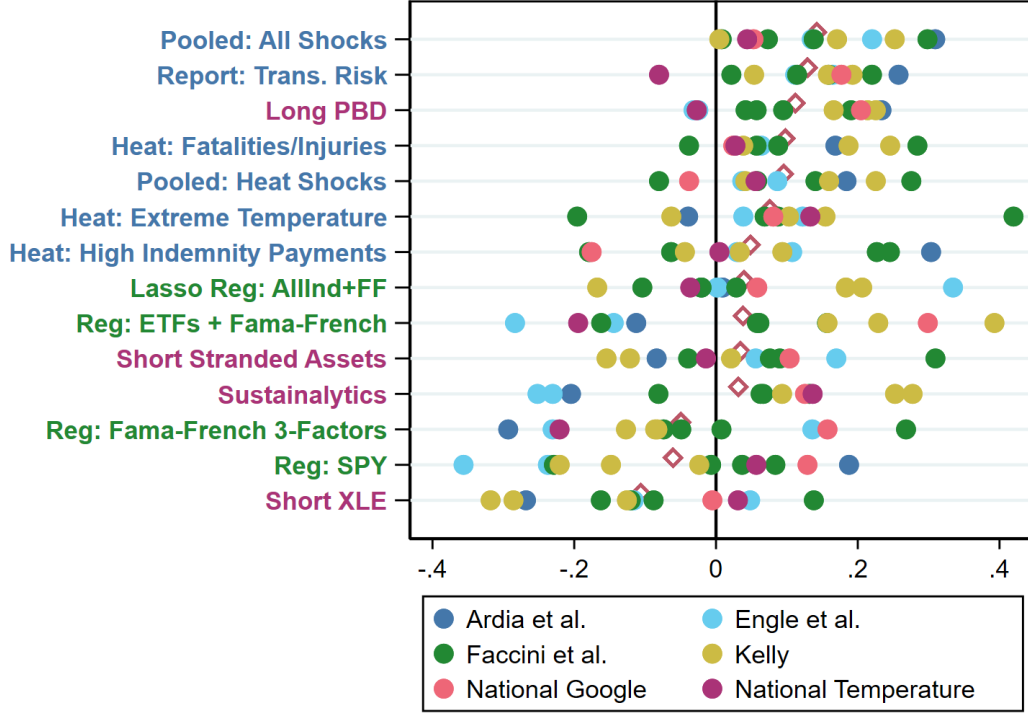
(a) Pooled: All Shocks



(b) Pooled: Heat Shocks

Note: Dot plots of monthly out-of-sample return correlations for variations of the quantity-based portfolio constructed using all four idiosyncratic belief shocks (Panel A) and the quantity-based portfolio constructed using the three heat shocks (Panel B). Each row corresponds to a different way to build the hedge portfolio, described in the text. Each dot represents one correlation coefficient. Different colors represent different groups of climate news series. The red rhombus shows the unweighted average among all correlations, and portfolios are sorted top-to-bottom by this value.

Figure 4: Climate Hedge Performance - Three-Year Rolling Windows



Note: Dot plot of monthly return correlations similar to those in Figure 2, but constructing the quantity-based and mimicking portfolios using Three-year rolling windows of data.

Finally, as mentioned in Section 3.4, in Appendix Figure A.5 we perform a placebo test where we repeat the analysis on the earlier decade, 2000-2010. As expected, in a decade in which investors did not price in climate risks, none of the approaches that rely on market information (the mimicking portfolio approach and our quantity-based approach) produce successful hedges (we do not study the performance of the narrative approaches as the ETFs we use are not available in that period).

4 Advantages and Disadvantages of the Different Hedging Approaches

In this section we discuss the main advantages and disadvantages of the various approaches to constructing hedge portfolios. We focus on three important elements: (1) the extent to which the approaches can deal with short time series and time-varying climate risk exposures, (2) the data requirements of the different approaches, and (3) whether the procedures require identifying the true climate risk exposures of different assets.

Short time series and time-varying exposures. The mimicking portfolio methodology takes a purely statistical approach to constructing hedge portfolios. It requires little input from the researcher beyond the choice of the base assets used in the projection, and instead

relies on the time series information (the historical covariance between the hedge target and asset returns) to choose the portfolio weights. This approach works well as long as the time series is sufficiently long ($T \rightarrow \infty$) and asset risk exposures are stable over time: in that case, a sufficiently large set of base assets (i.e., as $N \rightarrow \infty$) asymptotically generates the optimal hedge portfolio (see Giglio & Xiu 2021). The mimicking portfolio approach suffers particularly when T is small, as is the case with newly emerging risks such as climate risk, since a small T lead to noisy estimates of the covariance of prices with the hedging target. For similar reasons, the mimicking portfolio approach also struggles when asset exposures are time-varying, for example because firms’ strategic decisions affect their risk exposures over time.¹⁸ In fact, one can think of a change in exposure over time as having the same effect as reducing the time periods T that can be used to learn about the new exposure.

At the other end of the spectrum, the narrative approach does not rely on historical time series at all. Rather, it requires the investor to specify the different assets’ exposures to the hedge target based on their understanding, for example, of how each industry’s business model would be affected by the different types of climate risk. As the previous section showed, this process of identifying ex ante which stocks would stand to gain or lose when climate risk materializes can be difficult.

The quantity-based approach relies neither on having prior knowledge of which stocks will gain or lose when climate risks materialize, nor on having a long time series T or highly stable risk exposures. Instead, it weakens these requirements by using *cross-sectional* information on trading behavior instead of time-series information on prices to choose the portfolio weights. This allows investors and researchers to obtain many signals of asset exposures every period (in principle, one from each investor receiving an idiosyncratic climate belief shock), enabling them to construct climate hedge portfolios based on fewer time periods. It also allows investors to learn more quickly when asset exposures have changed.

Data Requirements. While the quantity-based approach has important benefits relative to the mimicking portfolio approach, it has stronger data requirements. First, the quantity-based approach requires the identification of idiosyncratic belief shocks, i.e., shocks that move investor beliefs about the aggregate risk yet only affect a few investors at the same time. While the need to identify such shocks might appear daunting, we believe that both the location-based and disclosure-based approaches to identifying idiosyncratic belief shocks can be applied to other types of risks beyond climate change. To highlight this, we show in Section 5 that both approaches can be used to identify reasonable hedge portfolios for other macro risk factors such as house price changes and unemployment rates.

The quantity-based approach also requires the researcher to observe portfolio holdings or trading data. Such quantity data is generally less widely available than price data, in particular for assets beyond equities. Expanding the base assets to include, for example,

¹⁸For example, van Benthem et al. (2022) discuss that European IOCs such as Shell and BP have announced ambitious net-zero targets, combined with substantial investments in renewable energies. Perhaps the most striking example is Orsted, Denmark’s largest power company, which has transformed itself from a largely hydrocarbon based firm to the largest offshore wind farm company in the world. Over time, Orsted’s exposure to transition risk would thus have shifted from negative to positive.

commodities or derivatives on the European Emissions Trading System (ETS) would require researchers to also observe investor holdings in these assets. While various regulators have access to portfolio holdings data that would allow them to implement our approach with a wider range of base assets, those data sets are not publicly accessible.

Accuracy of Exposure Measures. The quantity-based and mimicking portfolio approaches both aim to identify the average investors' perceptions of assets' climate risk exposures, one using quantity information and one using price information. But what if investors are wrong on average in their assessment of different assets' true climate risk exposures?

Even if investors misperceive assets' true risk exposures, both approaches can still build portfolios that hedge aggregate *news* about climate risks in the short run, as long as the average (incorrect) response to idiosyncratic climate belief shocks corresponds to the average (incorrect) response to global climate news shocks. For example, it may be that investors believe that car companies will benefit along the transition path (as suggested by their quantity responses to idiosyncratic climate belief shocks), but in reality, car companies will actually suffer disproportionate losses in response to transition risk realizations (perhaps because electric vehicles will largely be sold by new entrants rather than the incumbents currently trading on the stock market). In the short run, while the average investor holds this mistaken belief, it is likely that news of aggregate climate risk will push investors to buy car stocks and thus drive up their prices. Therefore the quantity-based portfolio would still hedge news about aggregate climate risk in the short term. Yet, in the long run, any portfolio that is long car stocks will ultimately lose in value once climate shocks actually materialize, and the true climate risk exposures are revealed. Said differently, the short-term ability of both quantity-based and mimicking portfolio approaches to hedge aggregate climate news only relies on *consistent* behavior of investors in response to idiosyncratic and aggregate shocks; the long-term hedging performance against actual climate risk realizations relies on markets (i.e. the average investor) being right about the risk exposures.

The narrative approach has the opposite challenge. Under the (certainly strong) assumption that the researcher constructing the hedge portfolio has a better understanding than the average investor of the true climate risk exposures of different assets, the resulting portfolio will likely have solid long-run hedging properties against aggregate climate risk *realizations*. However, in the short-run, while the researcher disagrees with the average market participant on different assets' risk exposures, the narrative approach will not hedge the arrival of *news* about climate risks. If the researcher constructing the hedge portfolio has incorrect perceptions of different assets' actual climate risk exposures, but disagrees with the (possibly also incorrect) assessments of the average market participant, the resulting narrative portfolio will neither be able to hedge the arrival of news about climate risks in the short run, nor will it be able to hedge the arrival of the actual climate risk realizations in the long run.

5 Hedging Macro Factors

While the main focus of this paper is on hedging climate risks, quantity-based portfolios can also be built to hedge other macroeconomic risks. In this section we briefly explore two such applications: hedging national unemployment rate changes and hedging national house price changes. In each case, we identify “idiosyncratic belief shocks” using local versions of the aggregate shocks as well as shocks based on disclosures in investor reports. We then construct hedge portfolios based on investors’ trading responses to these idiosyncratic shocks. Our motivation for analyzing local housing market and unemployment shocks is the connection between these local shocks and beliefs about the corresponding aggregate series documented in prior work. Most directly relevant here is the work of [Kuchler & Zafar \(2019\)](#), who show that locally experienced house price movements affect expectations about future U.S.-wide house price changes, and that personally experienced unemployment affects beliefs about the future national unemployment rate.

We view our efforts in this section as providing a “proof of concept” for the versatility of our quantity-based approach. They should not be considered as the definitively optimized approach to constructing quantity-based hedge portfolios for house price changes and unemployment rates changes. Researchers interested in implementing our approach to hedge specific aggregate risks should carefully consider how to construct the relevant local or disclosure-based belief measures in their specific contexts.

Local and National Unemployment Shocks. We obtain monthly data on county-level and national unemployment rates from the [Bureau of Labor Statistics](#). Local unemployment shocks are defined as quarterly AR(1) innovations in changes in the county-level unemployment rate (we estimate the local belief shifters over three months intervals to align them with the quarterly portfolio holdings data):

$$\Delta Unemp_{t,t-3,c} = \theta_c \Delta Unemp_{t-3,t-6,c} + \delta_{m(t)} + \epsilon_{t,c}, \quad (6)$$

$\Delta Unemp_{t,t-3,c}$ is the change in county c ’s unemployment rate between months t and $t-3$, and $\delta_{m(t)}$ are calendar month fixed effects to remove possible seasonality. We run the regression county by county and use these resulting county-level AR(1) innovations as idiosyncratic shifters of local investors’ beliefs about future changes in the national unemployment rates. The hedge targets are AR(1) innovations of national changes in the unemployment rate at the monthly frequency:

$$\Delta Unemp_{t,t-1} = \theta_{national} \Delta Unemp_{t-1,t-2} + \delta_{m(t)} + \epsilon_t, \quad (7)$$

Local and National House Price Shocks. Our seasonally adjusted house price measure is the [Zillow Home Value Index \(ZHVI\)](#). We obtain local house price shocks as AR(1) innovations in the three-month growth rate of county-level house prices, where $\Delta \text{Log}(ZHVI_{t,t-3,c})$ captures the house price growth in county c between months t and $t-3$. $\delta_{m(t)}$ are calendar month fixed effects to remove possible seasonality:

$$\Delta \text{Log}(ZHVI_{t,t-3,c}) = \theta_c \Delta \text{Log}(ZHVI_{t-3,t-6,c}) + \delta_{m(t)} + \epsilon_{t,c}. \quad (8)$$

We run the regression county by county and use the resulting county-level AR(1) innovations as idiosyncratic shifters of local investors’ beliefs about changes in national house price growth. We construct the corresponding monthly AR(1) innovations of changes in the national ZHVI as our hedge target:

$$\Delta \text{Log}(ZHVI_{t,t-1}) = \theta_{\text{national}} \Delta \text{Log}(ZHVI_{t-1,t-2}) + \delta_{m(t)} + \epsilon_t. \quad (9)$$

Note that, unlike for climate news and the unemployment rate, positive innovations in the house price growth series constitute “good” news, both at the local and the national level.

Disclosure-Based Idiosyncratic Shocks. In addition to using local unemployment and house price developments as shifters of investors’ beliefs about the corresponding aggregate series, we also attempt to directly measure changes in investor beliefs about these macro risks from mutual fund managers’ disclosures in N-CSR reports.

To measure beliefs about national movements in unemployment rates and house prices, we first extract relevant sentences from these N-CSR reports. For the unemployment rate, we focus on sentences that contain one of the following words: ‘employment’, ‘unemployment’, ‘job’, ‘hiring’, and ‘labor market’. We exclude sentences with unrelated phrases such as ‘Jobs Act’. For house price changes, we focus on sentences that contain the words ‘housing’ or ‘house’. We exclude sentences that include unrelated terms such as ‘White House’, ‘House of Representatives’, ‘in-house modeling’, or ‘clearing house’. To further restrict our sample to sentences that express beliefs about the respective national series, we only focus on sentences that also contain one of the words ‘expect’, ‘believe’, or ‘anticipate’.

We then use the Bidirectional Encoder Representations from Transformers (BERT) model, a state-of-art language model proposed by [Devlin et al. \(2018\)](#), to classify each of these sentences to determine whether it expresses a positive or negative sentiment about the labor or the housing market. The pre-trained model we used is developed by [Araci \(2019\)](#). Positive sentences get a score of “1”, and negative sentences get a score of “-1”. Table [A.1](#) presents examples of sentences relating to these risks, alongside their BERT sentiment classification.

In the final step, we add up the sentiment scores of all relevant sentences in a report to classify each report as overall positive or overall negative regarding the particular risk. As with our climate risk application, we use *changes* in this measure between consecutive reports to determine idiosyncratic changes in beliefs about macro risks.

Hedge Portfolio Construction. The construction of the hedge portfolios follows the approach described in Section [2](#). We first estimate the regression in Equation [3](#) with different fund-specific measures of idiosyncratic belief changes to obtain industry-specific quantity betas for each of the two macro risks. For the location-based belief shocks, we use the county-level innovations from the AR(1) processes estimated in Equation [6](#) and Equation [9](#) as proxies for each fund’s idiosyncratic belief shock $S_{f,t}$.

For the disclosure-based measures of changes in unemployment beliefs, we assign $S_{f,t}$ a value of “1” if the fund’s overall unemployment sentiment deteriorates between reports (e.g., when it changes from positive to negative or neutral and when it changes from neutral to negative), a value of “0” if the sentiment is unchanged, and a value of “-1” if the

fund’s sentiment about unemployment rates improves (e.g., when it changes from negative to neutral or positive and when it changes from neutral to positive). Positive quantity betas—corresponding to long positions in the hedge portfolio—therefore describe industries that investors disproportionately buy when they become more pessimistic about national unemployment rates. We would thus expect the hedge portfolio to outperform when the national unemployment rate increases unexpectedly. Positive innovations in house price growth constitute *good news*.¹⁹ Therefore, for the disclosure-based measure of changes in house price beliefs, we assign $S_{f,t}$ a value of “1” if the fund’s expressed sentiment regarding the housing market improves between reports, and a value of “-1” when the sentiment deteriorates between reports. For both changes in unemployment rate and house prices, we set $S_{f,t}$ equal to “+1” / “0” / “-1” whenever at least one of the months during which we measure the active portfolio changes lies between subsequent N-CSR reports.

Using the estimated quantity betas, we then construct quantity-based hedge portfolios for unemployment rate changes and house price growth as described in Equation 4. Moreover, for comparison, we construct mimicking portfolios as described in Section 3.3.

Hedge Performance. Table 10 shows the quarterly out-of-sample correlations of various hedge portfolios with AR(1) innovations in the national unemployment rate and the national growth rate of house prices. To align with the approach to hedging climate news, we estimate the out-of-sample performance using data from 2015 to 2019, and estimate the hedge portfolios using five-year rolling windows of holdings and price data.

Table 10: Macro Hedge Performance

	Hedge Target	
	Growth in House Prices	Δ Unemployment Rate
<i>Mimicking Portfolio Approaches</i>		
Reg: Fama-French Three-Factors	.11	−.03
Reg: SPY	.09	.23
Lasso Reg: All-Industries + Fama-French	.11	.03
<i>Quantity-based Approaches</i>		
Quantity: Local Shocks	.17	.23
Quantity: Disclosure	.14	.10

Note: Monthly correlations for various hedge portfolios’ returns with AR(1) innovations of national changes in the unemployment rate and the national house price growth rate. The first three rows are mimicking portfolios, and the last two rows are quantity-based portfolios. Each row represents a hedge portfolio, whereas each column corresponds to the innovations of either the growth in house prices or changes in the national unemployment rate. Positive correlation coefficients are highlighted in bold.

The mimicking portfolio approach generally produces hedge portfolios with positive out-of-sample correlations with the corresponding hedge targets, with varying performance as we change the base assets. For example, using all industries and Fama-French factors with lasso works better for hedging house price shocks (correlation of 0.11) than unemployment

¹⁹An investor hoping to hedge surprisingly weak house price growth should therefore short the hedge portfolio we construct here.

shocks (0.03), whereas using the market alone works better for unemployment shocks (0.23) than for house price shocks (0.09). The relative success of the mimicking portfolio approach for hedging these two macro risks is indicative of comparatively stable industry exposures to these shocks during the time period considered.

The quantity-based portfolios also perform well with respect to the macro series that they target. The house-price quantity-based portfolio formed using local shocks performs well in hedging aggregate shocks to house price growth (correlation of 0.17), whereas the unemployment-based quantity portfolio formed using local shocks hedges unemployment shocks well (0.23).

The last row of Table 10 show the hedging performance of the disclosure-based macro risk hedge portfolios. These portfolios have returns that also correlate positively with their hedge targets (correlations of 0.14 for house price growth shocks, and 0.10 for shocks to the unemployment rate). Though this performance is similar to that of the mimicking-portfolio approach, there are many margins of adjustment that could be explored in order to improve the disclosure-based approach, such as further fine-tuning the model for financial sentiment classification or varying the words used to identify beliefs about macro risks. We leave such adjustments to using the quantity-based approach for hedging shocks other than climate risk to future work.

6 Conclusions and Directions for Future Research

In this paper we introduce a quantity-based approach to hedging aggregate news about climate change and other macro risks. Our quantity-based hedge portfolios outperform traditional approaches to hedging climate risks.

Despite the initial success of the quantity-based approach, we believe that investors interested in operationalizing this approach can further improve upon the resulting hedge performance by introducing portfolio holdings data from a wider range of investors, including retail investors, and by expanding the set of base assets beyond industry equity portfolios. For example, including positions in commodity or carbon futures may further improve the hedge portfolios’ ability to hedge the arrival of aggregate physical or transition risk news. Similarly, a fruitful direction for further work would be to explore whether other severe weather events beyond extreme heat (e.g., hurricanes and wildfires) can also be used as shifters of investors’ climate risk beliefs, thus potentially expanding the set of trading activities that can inform the construction of quantity-based hedge portfolios.²⁰

Our work so far has concentrated on exploring the hedging ability of various portfolios in terms of the correlations of their returns with realizations of climate news. Future work should focus both on the expected returns of these hedge portfolios as well as on the average

²⁰One possible concern with using such events, that does not apply to using heat shocks, is that wildfires and hurricanes often destroy local physical capital. To the extent that “home bias” makes local investors more exposed to the resulting decline in local stock prices (see [Huberman 2001](#), [Kuchler et al. 2022](#)), there may be other forces beyond the updating of climate beliefs that affect investors’ trading behavior.

volatilities of these portfolios. Both of these objects are informative about the overall costs of hedging climate risks.

The focus of our application is the cross-sectional allocation of climate risk across investors, taking as given the total amount of climate risk in the economy. Of course, re-allocating risks can have general equilibrium effects which in turn affect the aggregate amount of climate risk. The canonical channel for this effect is that equity market reallocation can affect the cost of capital for firms, differentially affecting investment for ‘green’ and ‘brown’ firms. In the context of climate change, there is significant debate as to how large changes in the cost of capital are, and the extent to which they have the potential to reduce overall emissions (see [Pedersen et al. 2021](#), [Pástor et al. 2021](#), [Goldstein et al. 2022](#), [Berk & van Binsbergen 2021](#), [Bolton & Kacperczyk 2021b](#)). More generally, even the direction of the overall amount of climate risk is not clear. For example, hedging climate risk decreases its economic cost and could lessen the incentives to mitigate these risks. Ultimately, the effect of individual investors’ hedging of climate risk on the aggregate amount of such risk is ambiguous, and understanding the quantitative importance of the various channels is an ongoing important area for research.

Lastly, while our focus on this paper has been on hedging climate risks, investors are also increasingly focusing on other emerging risks, such as cybersecurity risks or pandemic risks. An interesting avenue for future work would be to explore the extent to which our new quantity-based approach can allow investors to also improve investors’ ability to hedge these and other risks.

References

- Acharya, V. V., Johnson, T., Sundaresan, S. & Tomunen, T. (2022), Is physical climate risk priced? evidence from regional variation in exposure to heat stress, Technical report, National Bureau of Economic Research.
- Alok, S., Kumar, N. & Wermers, R. (2020), ‘Do Fund Managers Misestimate Climatic Disaster Risk’, *The Review of Financial Studies* **33**(3), 1146–1183.
- Araci, D. (2019), ‘Finbert: Financial sentiment analysis with pre-trained language models’, *arXiv preprint arXiv:1908.10063*.
- 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).
- Armantier, O., Bruine de Bruin, W., Topa, G., van der Klaauw, W. & Zafar, B. (2015), ‘Inflation expectations and behavior: Do survey respondents act on their beliefs?’, *International Economic Review* **56**(2), 505–536.

- Armona, L., Fuster, A. & Zafar, B. (2019), ‘Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment’, *Review of Economic Studies* **86**(4), 1371–1410.
- Bachmann, R., Berg, T. O. & Sims, E. R. (2015), ‘Inflation expectations and readiness to spend: Cross-sectional evidence’, *American Economic Journal: Economic Policy* **7**(1), 1–35.
- Bailey, M., Cao, R., Kuchler, T. & Stroebel, J. (2018), ‘The economic effects of social networks: Evidence from the housing market’, *Journal of Political Economy* **126**(6), 2224–2276.
- Bailey, M., Dávila, E., Kuchler, T. & Stroebel, J. (2019), ‘House price beliefs and mortgage leverage choice’, *The Review of Economic Studies* **86**(6), 2403–2452.
- Bailey, M., Johnston, D. M., Koenen, M., Kuchler, T., Russel, D. & Stroebel, J. (2020), Social networks shape beliefs and behavior: Evidence from social distancing during the covid-19 pandemic, Technical report, National Bureau of Economic Research.
- Bakkensen, L. A. & Barrage, L. (2022), ‘Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics’, *The Review of Financial Studies* **35**(8), 3666–3709.
- Baldauf, M., Garlappi, L. & Yannelis, C. (2020), ‘Does climate change affect real estate prices? only if you believe in it’, *The Review of Financial Studies* **33**(3), 1256–1295.
- Barnett, J. (2020), ‘Global environmental change ii: Political economies of vulnerability to climate change’, *Progress in Human Geography* **44**(6), 1172–1184.
- Barnett, M. (2017), Climate change and uncertainty: An asset pricing perspective, Technical report, Working Paper.
- Berg, F., Koelbel, J. F. & Rigobon, R. (2022), ‘Aggregate confusion: The divergence of esg ratings’, *Forthcoming Review of Finance* .
- Berk, J. B. & van Binsbergen, J. H. (2016), ‘Assessing asset pricing models using revealed preference’, *Journal of Financial Economics* **119**(1), 1–23.
- Berk, J. & van Binsbergen, J. H. (2021), ‘The impact of impact investing’, *Available at SSRN 3909166* .
- Bernstein, A., Gustafson, M. T. & Lewis, R. (2019), ‘Disaster on the horizon: The price effect of sea level rise’, *Journal of Financial Economics* **134**(2), 253–272.
- Billio, M., Costola, M., Hristova, I., Latino, C. & Pelizzon, L. (2021), ‘Inside the esg ratings:(dis) agreement and performance’, *Corporate Social Responsibility and Environmental Management* **28**(5), 1426–1445.

- Bolton, P. & Kacperczyk, M. (2021*a*), ‘Do investors care about carbon risk?’, *Journal of Financial Economics* **142**(2), 517–549.
- Bolton, P. & Kacperczyk, M. (2021*b*), Global pricing of carbon-transition risk, Technical report, National Bureau of Economic Research.
- Busse, M. R., Pope, D. G., Pope, J. C. & Silva-Risso, J. (2015), ‘The Psychological Effect of Weather on Car Purchases *’, *The Quarterly Journal of Economics* **130**(1), 371–414.
- Ceccarelli, M., Ramelli, S. & Wagner, A. F. (2021), ‘Low-carbon mutual funds’, *Swiss Finance Institute Research Paper* (19-13).
- Chang, S. (2019), ‘Local industry bias in investor behavior: Evidence from mutual funds’, *Working Paper*.
- Chang, T. Y., Huang, W. & Wang, Y. (2018), ‘Something in the Air: Pollution and the Demand for Health Insurance’, *The Review of Economic Studies* **85**(3), 1609–1634.
- Chen, Q., Goldstein, I. & Jiang, W. (2010), ‘Payoff complementarities and financial fragility: Evidence from mutual fund outflows’, *Journal of Financial Economics* **97**(2), 239–262.
- Choi, D., Gao, Z. & Jiang, W. (2020), ‘Attention to Global Warming’, *The Review of Financial Studies* **33**(3), 1112–1145.
- Cohen, L., Gurun, U. G. & Nguyen, Q. H. (2020), The esg-innovation disconnect: Evidence from green patenting, Technical report, National Bureau of Economic Research.
- D’Acunto, F., Malmendier, U. & Weber, M. (2022), ‘What Do the Data Tell Us About Inflation Expectations?’, *Working Paper*.
- Deryugina, T. (2013), ‘How do people update? the effects of local weather fluctuations on beliefs about global warming’, *Climatic Change* **118**(2), 397–416.
- Devlin, J., Chang, M.-W., Lee, K. & Toutanova, K. (2018), ‘Bert: Pre-training of deep bidirectional transformers for language understanding’, *arXiv preprint arXiv:1810.04805*.
- Dou, W. W., Kogan, L. & Wu, W. (2022), Common fund flows: Flow hedging and factor pricing, Technical report, National Bureau of Economic Research.
- 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.
- Elliott, J. T. (2022), ‘Investment, emissions, and reliability in electricity markets’, *Working Paper*.
- Engle, R. F., Giglio, S., Kelly, B., Lee, H. & Stroebe, 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.
- Frazzini, A. & Lamont, O. A. (2008), ‘Dumb money: Mutual fund flows and the cross-section of stock returns’, *Journal of Financial Economics* **88**(2), 299–322.
- Gennaioli, N., Ma, Y. & Shleifer, A. (2016), ‘Expectations and investment’, *NBER Macroeconomics Annual* **30**(1), 379–431.
- Giglio, S., Kelly, B. & Stroebe, J. (2021), ‘Climate finance’, *Annual Review of Financial Economics* **13**, 15–36.
- Giglio, S., Maggiori, M., Rao, K., Stroebe, J. & Weber, A. (2021), ‘Climate change and long-run discount rates: Evidence from real estate’, *The Review of Financial Studies* **34**(8), 3527–3571.
- Giglio, S., Maggiori, M., Stroebe, J. & Utkus, S. (2021a), ‘Five facts about beliefs and portfolios’, *American Economic Review* **111**(5), 1481–1522.
- Giglio, S., Maggiori, M., Stroebe, J. & Utkus, S. (2021b), ‘The joint dynamics of investor beliefs and trading during the covid-19 crash’, *Proceedings of the National Academy of Sciences* **118**(4), e2010316118.
- Giglio, S. & Xiu, D. (2021), ‘Asset pricing with omitted factors’, *Journal of Political Economy* **129**(7), 1947–1990.
- Goldsmith-Pinkham, P. S., Gustafson, M., Lewis, R. & Schwert, M. (2021), ‘Sea level rise exposure and municipal bond yields’, *Working Paper* .
- Goldstein, I., Kopytov, A., Shen, L. & Xiang, H. (2022), On esg investing: Heterogeneous preferences, information, and asset prices, Technical report, National Bureau of Economic Research.
- Grimaldi, A., Javanmardian, K., Piner, D., Samandari, H. & Strovink, K. (2020), Climate change and p&c insurance: The threat and opportunity, Technical report, McKinsey & Company.
- Grinblatt, M. & Titman, S. (1989), ‘Mutual fund performance: An analysis of quarterly portfolio holdings’, *The Journal of Business* (3), 393–416.
- Heyes, A. & Saberian, S. (2019), ‘Temperature and decisions: Evidence from 207,000 court cases’, *American Economic Journal: Applied Economics* **11**(2), 238–265.

- Hoepner, A. G., Oikonomou, I., Sautner, Z., Starks, L. T. & Zhou, X. (2018), ‘Esg shareholder engagement and downside risk’, *Working Paper* .
- Hsu, P.-H., Li, K. & Tsou, C.-Y. (2022), ‘The pollution premium’, *Available at SSRN 3578215* .
- Huberman, G. (2001), ‘Familiarity breeds investment’, *The Review of Financial Studies* **14**(3), 659–680.
- 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’.
- Koijen, R. S., Richmond, R. J. & Yogo, M. (2020), Which investors matter for equity valuations and expected returns?, Technical report, National Bureau of Economic Research.
- Koijen, R. S. & Yogo, M. (2019), ‘A demand system approach to asset pricing’, *Journal of Political Economy* **127**(4), 1475–1515.
- 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., Li, Y., Peng, L., Stroebe, J. & Zhou, D. (2022), ‘Social proximity to capital: Implications for investors and firms’, *The Review of Financial Studies* **35**(6), 2743–2789.
- 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.
- Lot, G. & Haegeli, J. (2021), More risk: The changing nature of p&c insurance opportunities to 2040, Technical report, Sigma Re Institute.
- Magolin, S. & Santino, K. (2022), Policy will engage with reality: Re-focus on secular gas theme as oil madness abates, Technical report, Wolfe Research.
- Malmendier, U. & Nagel, S. (2011), ‘Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?’, *The quarterly journal of economics* **126**(1), 373–416.

- Murfin, J. & Spiegel, M. (2020), ‘Is the risk of sea level rise capitalized in residential real estate?’, *The Review of Financial Studies* **33**(3), 1217–1255.
- Painter, M. (2020), ‘An inconvenient cost: The effects of climate change on municipal bonds’, *Journal of Financial Economics* **135**(2), 468–482.
- Pástor, L., Stambaugh, R. F. & Taylor, L. A. (2021), ‘Sustainable investing in equilibrium’, *Journal of Financial Economics* **142**(2), 550–571.
- Pedersen, L. H., Fitzgibbons, S. & Pomorski, L. (2021), ‘Responsible investing: The esg-efficient frontier’, *Journal of Financial Economics* **142**(2), 572–597.
- Pickl, M. J. (2019), ‘The renewable energy strategies of oil majors: From oil to energy?’, *Energy Strategy Reviews* **26**, 100370.
- Ritchie, H., Roser, M. & Rosado, P. (2020), ‘Co2 and greenhouse gas emissions’, *Our world in data* .
- Roth, C. & Wohlfart, J. (2020), ‘How do expectations about the macroeconomy affect personal expectations and behavior?’, *The Review of Economics and Statistics* **102**(4), 731–748.
- Ryan, S. P. (2012), ‘The costs of environmental regulation in a concentrated industry’, *Econometrica* **80**(3), 1019–1061.
- 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?’, *Journal of Financial Economics* **142**(2), 487–498.
- Tomunen, T. (2021), ‘Failure to share natural disaster risk’, *Available at SSRN 3525731* .
- van Benthem, A., Crooks, E., Giglio, S., Schowb, E. & Stroebel, J. (2022), ‘The effect of climate risks on the interactions between financial markets and energy companies’, *Nature Energy* **7**(8), 690–697.
- Wermers, R., Yao, T. & Zhao, J. (2012), ‘Forecasting Stock Returns Through an Efficient Aggregation of Mutual Fund Holdings’, *The Review of Financial Studies* **25**(12), 3490–3529.

A Appendix

A.1 Tables

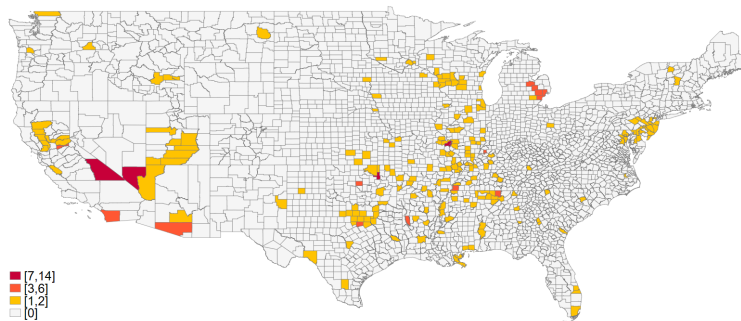
Table A.1: BERT classification examples

Shocks	Labels	Scores	Sentences
Unemployment	Positive	1	Unemployment is expected to recede significantly as the economy and business reopen, but it will take time to restore employment back to 2019 levels.
Unemployment	Positive	1	The employment situation remains sluggish, but economists anticipate an improvement in hiring in the coming months.
Unemployment	Negative	-1	However, due to pressure on growth for the U.S. economy and high unemployment expected for most of 2010, the Fed will most likely keep interest rates at exceptionally low levels which will affect interest rates for all money market mutual funds.
Unemployment	Negative	-1	And while the Fed may begin to roll back some of its bond purchases, we do not expect a change in the Federal rate policy stance, given that there is no inflationary pressure and unemployment continues above Fed targets.
House Price	Positive	1	However, with continued growth in consumer discretionary spending expected, we believe the U.S. housing market could rebound in 2019.
House Price	Positive	1	The housing market remained an area of weakness as home prices continued to fall, but we anticipate a pickup in demand as the weather improves in the months ahead.
House Price	Negative	-1	Having said that, the housing and job markets remain in poor shape, and we don't anticipate significant improvement in either until 2012.
House Price	Negative	-1	The biggest risk we see to our constructive economic view would be another sharp decline in the housing market, caused by more foreclosed houses hitting the market than is currently anticipated, thereby stifling demand for new homes.

Note: Sentences classified as positive are assigned a score of 1 and sentences classified as negative are assigned a score of -1.

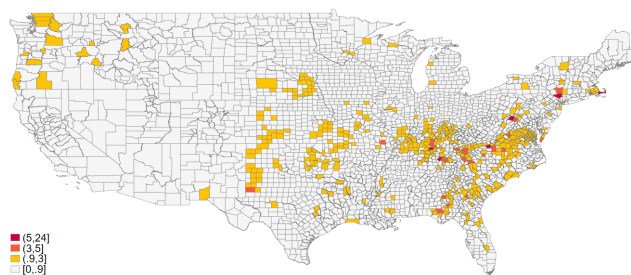
A.2 Figures

Figure A.1: Distribution of “Heat: Fatalities or Injuries”



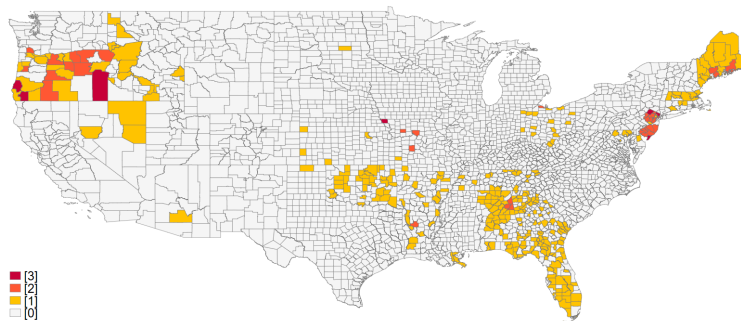
Note: Distribution of the “Heat: Fatalities or Injuries” climate shock from 2010 to 2019. The color-coding shows the number of county-months that experienced the “local” climate shock during the interval.

Figure A.2: Distribution of “Heat: High Indemnity Payments”



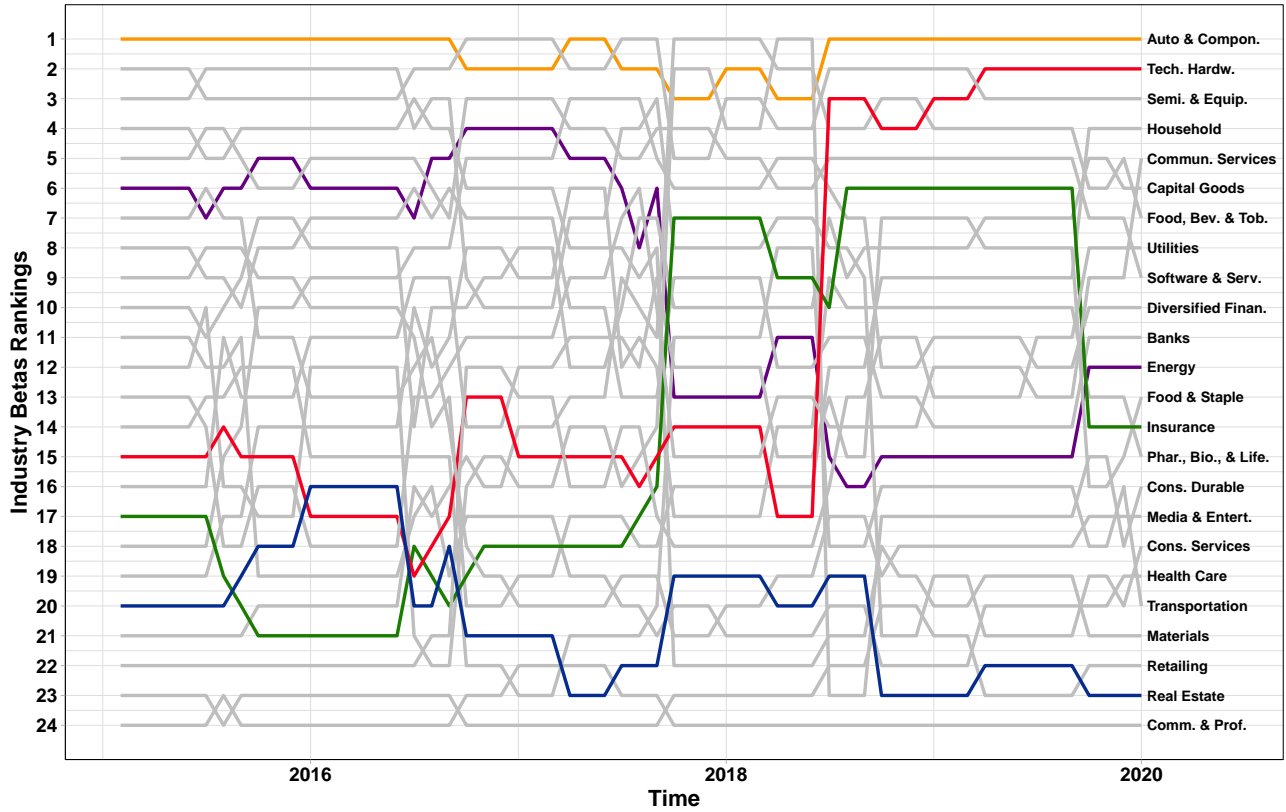
Note: Distribution of the “Heat: High Indemnity Payments” climate shock from 2010 to 2019. The color-coding shows the number of county-months that experienced the “local” climate shock during the interval.

Figure A.3: Distribution of “Heat: Extreme Temperature”



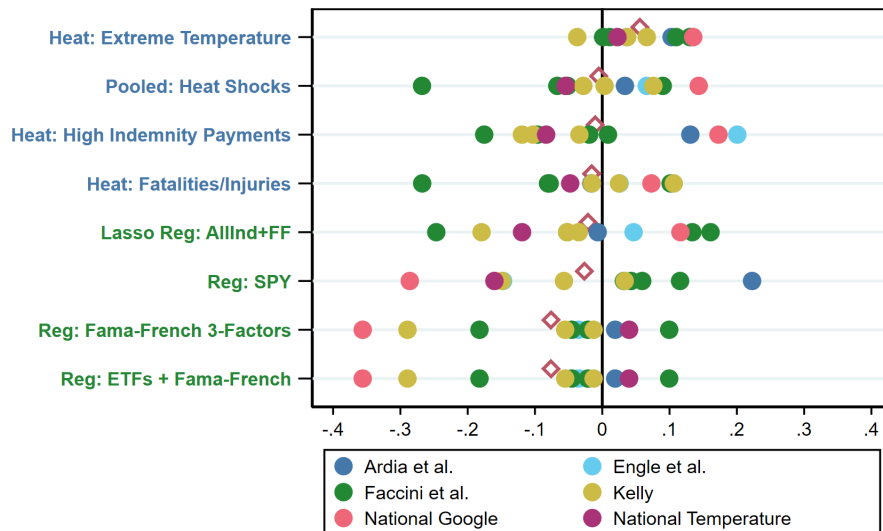
Note: Distribution of the “Heat: Extreme Temperature” climate shock from 2010 to 2019. The color-coding shows the number of county-months that experienced the “local” climate shock during the interval.

Figure A.4: Industry Betas Rankings Over Time




Note: Industry climate beta coefficients (estimated in Equation 3) rankings over time. The rankings are sorted by the average coefficient across the four individual shocks and are based on data for every five-year window from 2010 to 2019 inclusive. Each line represents an industry.

Figure A.5: Placebo test in earlier period

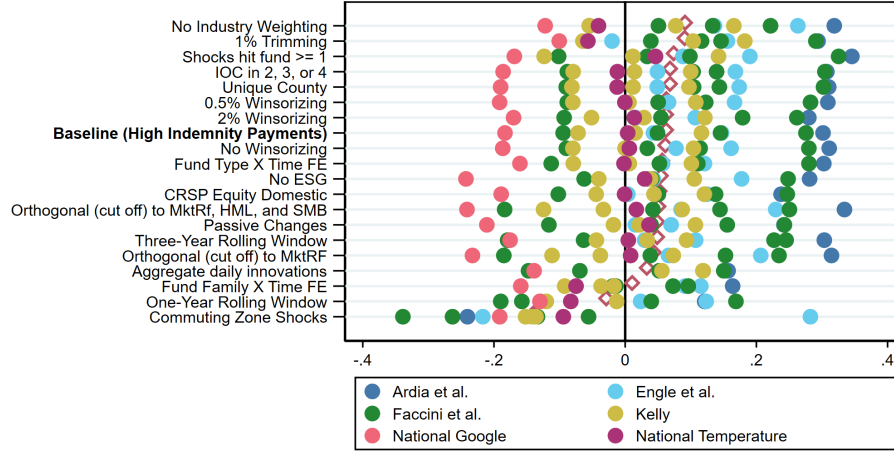


Note: Dot plot of monthly return correlations for the three heat-based hedging portfolios and mimicking portfolios with various climate news series AR(1) innovations, using data from 2000 to 2010.


Figure A.6: Climate Hedge Performance - Robustness Tests



(a) Heat: Fatalities/Injuries




(b) Heat: High Indemnity Payments



(c) Heat: Extreme Temperature

Figure A.6: Climate Hedge Performance - Robustness Tests (cont.)



(d) Report: Trans. Risk

Note: Dot plot of monthly return correlations for the three heat-based hedging portfolios and the disclosure-based hedging portfolio, with respect to the various climate targets. Each row corresponds to a different way to build the hedge portfolio, described in the text.