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 to hedge the economic risks related to climate change. Our *quantity-based* approach exploits information on how mutual fund managers trade in response to idiosyncratic changes in their climate change beliefs. We exploit two types of idiosyncratic belief shocks: (1) instances when fund advisers experience local extreme heat events that are known to shift climate change beliefs, and (2) instances when mutual fund managers change the language in the shareholder disclosures to specifically mention climate risks. We use the funds' observed portfolio changes following such idiosyncratic belief shocks to predict how investors will reallocate their capital when aggregate climate news shocks occur, news that shift the beliefs and asset demands of many investors and thus move equilibrium prices. We show that a portfolio that holds stocks that investors tend to buy after experiencing idiosyncratic climate belief shocks appreciates in value in periods with aggregate negative climate news shocks. Our quantity-based approach yields superior out-of-sample hedging performance compared to traditional methods of identifying hedge portfolios. The key advantage of the quantity-based approach is that it learns from 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 efficacy and versatility of the quantity-based approach by constructing successful hedge portfolios for aggregate unemployment and house price risk.

---

*This version: June 14, 2022. We thank Daron Acemoglu, Michael Barnett, Richard Berner, Robert Engle, Nicola Gennaioli, Itay Goldstein, Wei Jiang, Annette Vissing-Jorgensen, Dexin Zhou, and seminar and conference participants at the NBER, SFS Cavalcade, WFA, NYU Stern, Columbia University, Cornell, MIT Sloan, the New York Fed, Dallas Fed, San Francisco Fed, the Federal Reserve Board, the Central Bank of Chile, UCLA, UC Boulder, Fordham University, DIW, UNC-Duke, PRI, SFS Cavalcade, WFA, 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.

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

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

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

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

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

Climate change presents a major challenge to humanity: in addition to a wide range of social implications, both the physical effects of climate change and the regulatory efforts to slow carbon emissions have the potential to substantially disrupt economic activity. As investor awareness of the economic and financial risks of climate change has increased, there has been rising demand for financial instruments that hedge these risks (see Krueger et al. 2020*a*, Giglio, Kelly & Stroebel 2021, Stroebel & Wurgler 2021). At present, there exists at most a limited set of instruments that are targeted directly at hedging various climate risks, most prominently the relatively small market for "catastrophe bonds" (Tomunen 2021). However, investors interested in hedging climate risks can still build a hedging portfolio using other assets, like stocks, by exploiting the assets' exposures 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 climate hedge (e.g., Engle et al. 2020).

The asset pricing literature has explored different approaches to build hedge portfolios for a variety of macro risks, the most prominent of which is the mimicking portfolio approach of Lamont (2001). These existing approaches, however, strongly 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 the realizations of the hedge target. This makes existing approaches poorly suited in cases in which the targeted risks are new and the time series data is limited, as in the case of climate risk.

In this paper, we propose a novel methodology to build portfolios to hedge climate risks. Our *quantity-based* approach uses cross-sectional information on investors' trading activity to identify which stocks to hold in the hedging 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. Such shocks do not affect asset prices, because they are idiosyncratic, but they still affect the affected investors' trading activity. Based on this insight, the quantity-based approach studies systematically how investors *trade* in response to these shocks, thus learning how investors' stock demand is shifted by these idiosyncratic belief shocks. Specifically, it identifies which stocks investors tend to buy and sell after a 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 for this 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 shock similarly to how they respond to the idiosyncratic shock, we expect the correlated shift in demand of many investors to move prices. Therefore, the hedge portfolio will rise in value when aggregate climate risks materialize.

We operationalize this quantity-based approach by building a hedge portfolio for climate risk using U.S. stocks, and focusing on an important group of investors whose trading we can observe: mutual funds. The first step is the identification of the "idiosyncratic belief shocks".

We propose two ways to identify these shocks. The first looks at geographically localized extreme heat events, based on prior literature that documents how such heat shocks are an important driver of beliefs about global climate risk (see Egan & Mullin 2012, Deryugina 2013, Joireman et al. 2010, Li et al. 2011, Fownes & Allred 2019, Sisco et al. 2017). At the same time, these heat shocks are by their nature very localized, so they only affect a small number of investors. We consider three different ways to measure extreme local heat shocks, based on the presence of fatalities and injuries, the extent of indemnity payments due to extreme heat, and the occurrence of extreme temperatures relative to the history of each location. For each of these three shocks, we validate that the occurrence of the shock in an area leads to increased Google searches for the term "climate change", providing suggestive evidence that these heat shocks indeed induce updates in climate beliefs.

The second approach to identifying idiosyncratic belief shocks is based instead on the coverage of transition risk (that is, the risk stemming from government responses to climate risks) in mutual funds' shareholder reports. This approach does not rely on heat events as shifters of beliefs; rather it directly identifies a change in beliefs about or attention paid to climate change at the individual fund level.

Next, we study how investors trade in response to these idiosyncratic belief shocks. We focus on U.S. active mutual funds, whose holdings we can observe at the quarterly frequency, and study how they change their portfolio allocation across industries in quarters when they experience the shocks.[1] The four belief shock measures have very low correlation with each other: different heat shocks tend to affect some areas more than others, and mutual funds' disclosures are updated in response to many other events that also affect managers' climate beliefs in idiosyncratic ways. Interestingly, however, we find that the kind of industries that funds buy and sell in response to these different shocks are significantly correlated: that is, the trading behavior of funds that experience one type of idiosyncratic climate belief shock is similar to that of funds experiencing a different shock.

There are several interesting patterns that emerge by looking at which industries are bought or sold in response to climate belief shocks. For example, the auto industry is the industry with the strongest positive quantity response (that is, mutual funds tend to buy auto stocks after increasing concerns around climate risks). While that may appear surprising at first glance—automobiles are, after all, an important source of carbon emissions—this may reflect managers' 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 for sure why mutual funds are buying or selling these industries. To build a hedge portfolio based on this information, we rely on the consistency of funds' behavior in response to idiosyncratic shocks and aggregate shocks; that is, we rely on trading patterns following idiosyncratic shocks to take place again when aggregate climate

[1] At this stage, we do not focus on individual stocks, because doing so would add significant estimation error given the large number of stocks.

risks materialize. In the paper, we provide different ways to validate this consistency assumption. First, we show that the industry-level trading activity in response to local shocks is similar across periods in our sample as well as across different investors. Second, we show that, while our different measures of belief shocks have low correlations, their corresponding quantity-based hedging portfolios are significantly correlated, indicating that investors trade similarly in response to different climate shocks. Finally, and perhaps most importantly, we show directly that this cross-sectional quantity information is indeed useful for learning about the pricing response to aggregate shocks by studying the out-of-sample hedging performance of the quantity-based portfolios for aggregate climate shocks.

In particular, we use the observed industry-level trading of mutual funds in response to managers' idiosyncratic belief shocks to build a long-short 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, and 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. (2020), and Hoepner et al. (2018). This approach will hedge climate risk if the underlying economic intuition is correct, and properly identifies which companies would stand to gain or lose from the occurrence of a climate shock. This approach has the advantage that, like ours, it does not require long time series to be implemented; however, it requires investors to have correct priors about how firm characteristics relate to performance following realizations of climate risk.

The second alternative approach is the "mimicking portfolio" approach as in Lamont (2001), where climate risk series are projected onto a set of portfolio returns using time-series information. The mimicking portfolio approach relies strongly on time-series data: since it does not take an a priori view on which assets gain or lose when climate shocks occur, it needs to learn this from assets' 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 in the second half of our sample (2015–2019). For the mimicking portfolio approach and the quantity-based approach, we construct the hedge portfolios using rolling 5-year windows of price and quantity data, respectively.[2] To evaluate the hedging performance 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 different approaches 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

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

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 approaches give 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 methodology based on fatalities and injuries yields positive out-of-sample correlations with *all* of our aggregate climate shock series, with maximum correlations of above 30%. The other quantity-based portfolios do almost as well, and 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 to hedge aggregate climate shocks.

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: is the broader stock market likely to increase or decrease upon the realization of climate risks? This relationship turns out to be unstable over time: during 2010-2014, the S&P 500 comoved positively with climate risk realizations, while during 2015-2019, it comoved negatively, highlighting the challenges of the mimicking portfolio approach to constructing successful climate hedges. Adding more base assets in the construction of the mimicking portfolio can help better target the hedge, but requires estimating more parameters, again an issue in short samples. Narrative-based portfolios are immune to the short-sample issue, since historical data is not used to determine positions. However, deciding on positions in an a priori way is hard: as an example, for many industries, the different co-authors of this paper would have picked widely divergent holdings. 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 good hedge, but shorting traditional energy companies provides a bad hedge, with mostly negative out-of-sample correlations with the various targets.

The primary focus of our paper is to use our new quantity-based approach to construct portfolios that hedge realizations of climate risk. This is a natural application of our methodology; climate change is a first-order risk that has attracted investor attention only recently, and therefore there is not enough time-series data to allow precise estimation of 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 belief 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). 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.

Finally, it is useful to discuss what conditions are necessary for the quantity-based approach to produce a good hedge to climate risk, and how they compare with those required by the alternative approaches. The quantity-based approach relies on the information contained in mutual fund trades about various stocks' exposures to climate change. If mutual funds are wrong in their evaluation, the hedging portfolio might still work in the short term to hedge climate news (as long as mutual funds keep behaving in a consistent way across time periods); however, in the long run, when climate shocks ultimately materialize, mistakes made by mutual funds managers would still lead to imperfect hedging by the quantity-based portfolio. Note that the mimicking portfolio approach similarly relies on information from market participants (because it uses market prices to determine which stocks to buy and sell), whereas the narrative approach relies on the investors' own views.

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, reallocating 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 strong this effect is (see Pedersen et al. 2021, Pástor et al. 2021, Goldstein et al. 2022, Berk & van Binsbergen 2021, Bolton & Kacperczyk 2021*b*). On the other hand, hedging climate risk decreases its economic cost and can lessen the incentives to mitigate these risks. Ultimately, the effect of hedging 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.

Our work contributes to a growing literature that studies the interaction between climate change and asset markets. In equity markets, Bolton & Kacperczyk (2021*a*) and Hsu et al. (2022) find 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, while Choi et al. (2020) find that the stocks of carbon-intensive firms underperform during periods of abnormally warm weather, where investors' attention to climate risks is likely heightened. Barnett (2020) finds that increases in the likelihood of future climate policy action lead to decreased equity prices for firms with high exposure to climate policy risk. Koijen et al. (2020) find that shocks to demand for green firms tend to benefit long-term investors. The pricing of climate risk in real estate markets is explored by Baldauf et al. (2020), Bakkensen & Barrage (2021), Bernstein et al. (2019), Giglio, Maggiori, Rao, Stroebel & Weber (2021), and Murfin & Spiegel (2020).

More generally, this paper also builds upon the literature 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, Conlin et al. 2007, Busse et al. 2015, Chang et al. 2018, Cen 2021) and the extent to which 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, Stroebel & Utkus 2021, Roth & Wohlfart 2020).

Finally, our paper relates to a recent literature which uses quantity and holdings data in asset pricing. Berk & van Binsbergen (2016) uses flows data to test common asset pricing models. The demand-based asset pricing literature, started by Koijen & Yogo (2019), uses holdings data to predict price responses to fundamental shocks. We contribute to this literature by providing reduced form evidence that quantity information is useful for predicting price movements.

We structure our paper as follows. First, in Section 1, we describe a simple model of quantity-based information that justifies our approach to constructing hedge portfolios. We then explain the regional heat shocks and the fund report measure that we include in our analysis in Section 2.1. We describe the climate news indices, which we use to validate the quantity-based hedge portfolios, in Section 3.2. In sections 2 and 3, we apply the novel quantity-based approach to the case of climate change realizations. We validate the quantity climate change hedge portfolios in Section 3.4 and further demonstrate the efficacy of our methodology by hedging unemployment and housing prices in Section 5. We conclude this paper by summarizing the results and suggesting ideas for future research.

1 Quantity-Based Portfolios: A Simple Model

In this section, we describe a simple model that illustrates the mechanism behind 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 $q_A\big(p_A, \epsilon_A(i)\big)$, where p_A is the (relative) price of security A, and $\epsilon_A(i)$ gives investor i's beliefs about the (relative) future payoffs of security A. For simplicity, assume that $q_A\big(p_A, \epsilon_A(i)\big) = f(p_A) + g\big(\epsilon_A(i)\big)$, with f and g continuously differentiable, and $\frac{\partial f}{\partial p_A} < 0$. The market-clearing condition is:

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

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

An individual investor's beliefs can be decomposed into a common component ν_A and an investor-specific 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 events that we 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,

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

Thus, $\frac{\partial p_A^*}{\partial \omega_A(i)} = 0$. However, since investor i's demand changes, $\frac{\partial q^*}{\partial \omega_A(i)} \neq 0$. In words, if investor i 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 an aggregate shock about stock A, that affects all investor beliefs (i.e., a change in ν_A). For example, all investors believe that climate change regulation has become more likely, reducing their expected profitability of firm A. By the implicit function theorem and the chain rule, equilibrium price responses are given by:

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

In words, the sensitivity of prices to national news is directly proportional to average quantity sensitivities, $\int_{i=0}^{i=1} \frac{\partial q_A}{\partial \epsilon_A(i)} di$.[3] 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.

[3] The 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)

2 Idiosyncratic Belief Shocks & Portfolio Changes

Our quantity-based approach to constructing climate risk hedge portfolios requires identifying 'idiosyncratic belief shocks' that satisfy several criteria. First, the shocks should shift asset demands of affected investors through influencing their beliefs about climate risks (or their attention to these risks).[4] Second, the shocks should only affect a small subset of investors, so that they influence those investors' portfolios without inducing a large price response. Third, changes in asset demands in response to the idiosyncratic belief shocks should closely correspond to 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 directly measure the changes in the climate beliefs of different investors.

In the following section, we 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*

[4] Giglio, Maggiori, Stroebel & Utkus (2021) document that differences in beliefs across investors generally translate into differences in asset demands and portfolio holdings.

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.

Table 1: Summary Statistics on Extreme Heat Measures

Panel A: Local Heat Shocks: Summary

Climate Shock	Event Description	Frequency Monthly	Frequency Sample
Heat: Fatalities/Injuries	Injuries or fatalities	0.10%	1.32%
Heat: High Indemnities	90th percentile indemnity payments	1.08%	0.54%
Heat: Extr. Temperature	Monthly maximum temp. 4°C above avg maximum	0.22%	1.21%
Discl.: Transition Risk	Change in fund disclosures about transition risk	-	0.46%

Panel B: Local Shocks: Monthly Jaccard Correlations

	Fatalities/Injuries	Indemnities	Extr. Temperature
Heat: Fatalities/Injuries	1.00		
Heat: Indemnities	0.01	1.00	
Heat: Extr. Temperature	0.01	0.00	1.00

Panel C: Local Shocks: Sample Jaccard Correlations

	Fatalities/Injuries	Indemnities	Extr. Temperature	Discl.: Trans. Risk
Heat: Fatalities/Injuries	1.00			
Heat: Indemnities	0.06	1.00		
Heat: Extr. Temperature	0.04	0.00	1.00	
Discl.: Transition Risk	0.00	0.00	0.00	1.00

Note: Panel A provides an overview of the constructed 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. For heat events, the differences in "monthly" and "sample" frequencies arise from mutual fund advisers generally being located in high population density counties; injuries and fatalities from heat also disproportionately occur in areas with high population densities, while heat-related crop indemnity payments are more common in rural areas. The frequency of the disclosure-based measure is computed as the frequency of any changes happening, i.e., either an increase or a decrease in transition risk disclosures. Panel B shows the Jaccard correlation between the constructed heat measures across all county-months from 2010 to 2019, whereas Panel C shows the Jaccard correlation among the shock measures in our final sample. Intuitively, the Jaccard correlation measures the likelihood of observing both shocks conditional on observing one of them.

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 identify an extreme heat indemnity event when the monthly heat-related crop indemnity payments in a given county exceed the 90th percentile of non-zero payments across all U.S. county-months in the past 10 years; about 1.08% of county-months between 2010 and 2019 had such an event. Panel B of Table 1 highlights that the correlation of high crop indemnity

[5] Crop indemnity payments are insurance payments to farmers, which are paid when external disruptions lead to crop yields or revenues below the agreed amount in the insurance contract. The U.S. Department of Agriculture reports these payments for several private insurance companies, covering more than 100 crops.

heat events with heat-related fatalities and injuries is essentially zero. Crop indemnity payments are more frequent in low-density rural areas, whereas fatalities and injuries due to heat are more frequent in urban areas. Crop indemnity shocks therefore provide a source of variation for our analysis that is 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 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 prior two 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 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]

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]

[6] A 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 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.

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

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

where $\widetilde{G_{t,s}}$ is the 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 Indemnities		0.06^{**}	
		(0.03)	
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 is associated with an increase in relative interest in climate change of 5%. Reported crop indemnity payments increase the relative Google searches by 6%, whereas record temperatures are associated with an increase of 5%. These findings highlight that all of the extreme heat measures affect climate change awareness or beliefs.

2.2 Idiosyncratic Belief Shocks: Investor Disclosures

In addition to using extreme heat events as shifters of investors' climate risk concerns, we also attempt to directly measure changes in 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). We first extract sentences from the reports that use one of the following phrases: climate change, carbon emission(s), greenhouse gas(es), and global warming.[8] We then manually classify each of these sentences to determine whether it expresses concerns about transition risk. As an example, we identify the following passage from Parnassus Fund's Q4 2018 N-CSR filing as expressing transition risk concerns:

[8]These phrases were selected based on their cosine similarity to "climate change" in Google's word2vec implementation.

> The Fund's weakest performer was National Oilwell Varco, a global supplier of oilfield equipment and technology. The stock subtracted 28 basis points from the Fund's return, as its total return from our average. We sold our position in the third quarter after the company reported steep order declines for its oil-field equipment and significantly lower earnings guidance. As we re-evaluated our thesis on the stock, we became concerned about declining relevancy of the company's products as the global economy continues to decarbonize and shift toward renewable fuel sources and electric transportation. Following our exit of the stock, Parnassus adopted a firmwide fossil fuel free policy, which means that our funds will avoid investing in companies that derive significant revenues from the extraction, exploration, production or refining of fossil fuels.

Funds whose reports do not discuss climate transition risk get assigned a score of 0. Funds whose reports discuss transition risk get assigned a score of 1 if most sentences express concern about increasing transition risk (i.e., a climate policy tightening) and -1 if most sentences express views consistent with decreasing transition risk. In our baseline analysis, we use the *change* of this measure between consecutive reports as a measure of idiosyncratic changes in beliefs about climate transition risks, though the approach is robust to using alternative ways of aggregating this text-based information.[9]

This disclosure-based measure allows us to capture shifts in climate change awareness or beliefs that could by driven by forces other than extreme heat events. In addition, this measure is less 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).

2.3 Holdings and Location Data

There are a number of reasons that mutual fund managers are a natural focus of our study of trading responses due 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.

For our approach to work, mutual fund managers must care about climate risk. Managers may believe equilibrium valuations do not fully account for climate risks and thus they can earn alpha (Krueger et al. 2020*b*). Alternatively, managers may view climate 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. 2021). Notably, all of these

[9] In 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 constrution throughout our sample.

motivations lead to similar predictions for investors' responses to idiosyncratic belief shocks; when a manager becomes 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.[10] 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,[11] which focuses our hedge assets 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, keeping funds that have Investment Objective Code equal to 2 ("Aggressive Growth"), 3 ("Growth"), 4 ("Growth & Income"), or missing, and have CRSP Objective Code "Equity Domestic Non-Sector".[12]

We obtain stock-level characteristics from CRSP and Compustat. We 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 digits of the GICS code determine the stock's classification into the 24 "industry groups" that are the main focus of our analysis.[13]

Measuring Active Portfolio Changes. In our main analysis, we explore how idiosyncratic climate belief shocks induce changes in the share of the portfolio invested into industry I by fund f through active trading. We perform our analysis at the industry level, since the sparsity of the stock-level holding matrix would lead to potentially noisy estimates. For every fund f and time t, we define the active change in industry I holdings as:

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

where $P_{j,t-1}$ denotes the previous period price for stock j, $S_{f,j,t}$ denotes the number of shares of stock j held by fund f in period t, and $Share^{I}_{t}$ measures the market capitalization share of industry I as a fraction of the U.S. stock market. The term in square brackets therefore

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

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

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

[13] The 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.

captures the active change of the share of industry I in fund f's portfolio.[14] The reason for scaling by industry size is that a given increase in the portfolio share of a particular industry (i.e., shift of a given dollar amount invested) is likely to induce larger price movements for smaller industries. Since most funds report their holdings quarterly, we measure fund composition changes in three-month intervals. Finally, 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.[15]

For 67.6% of funds, all advisers reside in the same county. For our local 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.

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 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 total market

[14]Note that prices only appear as of $t-1$: we only study the change in portfolio allocation due to active trading as opposed to changes in value due to price movements during the quarter. Alternatively, one could analyze a separate variable, $PassiveChanges^{I}_{f,t}$, where the first fraction uses $P_{j,t}$ instead of $P_{j,t-1}$, i.e., current period holdings are valued at current period prices. This alternative approach takes price changes into account, and would be a more suitable model if we assume that funds constantly rebalance their portfolio. We verify that both approaches generate similar hedge portfolios.

[15]As 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

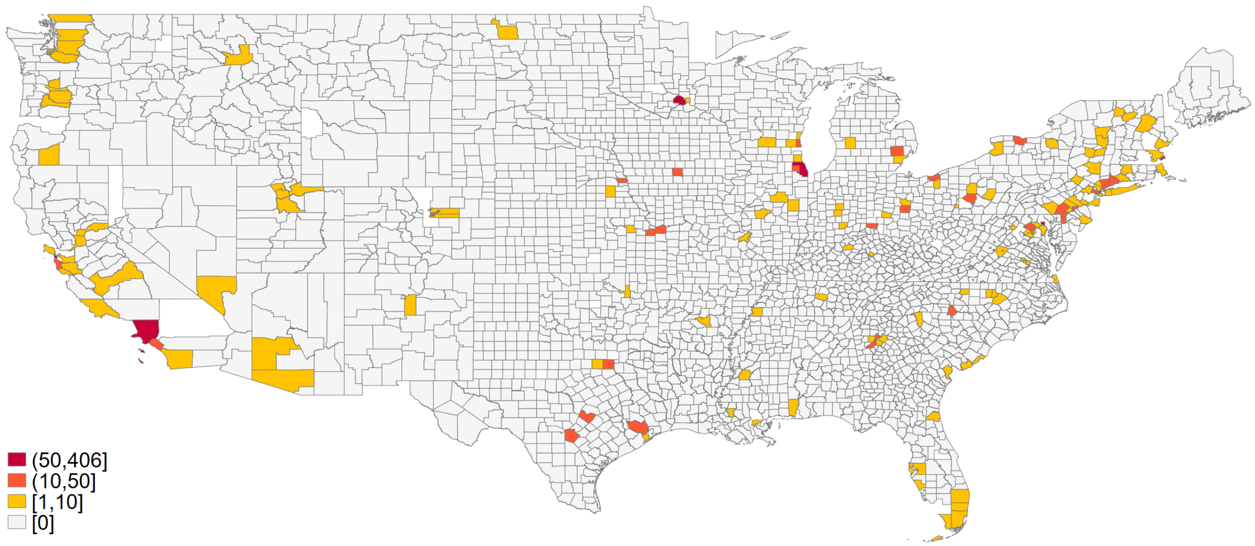

Panel A: Adviser Locations - Largest Counties

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.

capitalization. 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 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 $ActiveChanges^I_{f,t}$. 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:

$$ActiveChanges^I_{f,t} = \beta^{I,S,t} S_{f,t} + \delta^I_t + \epsilon^I_{f,t}, \tag{3}$$

where $ActiveChanges^I_{f,t}$ is defined as in equation 2. $S_{f,t}$ is one of the four idiosyncratic belief shocks described in Section 2.1: three local extreme heat shocks and the disclosure-based change in transition risk concern. δ^I_t captures time fixed-effects. The main object of interest are the climate-shock-specific estimates of $\beta^{I,S}$: for each industry I, this represents the differential change in fund holdings of that industry for 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 $\beta^{I,S}$ vary over time as we estimate regression 3 over different time horizons.

Table 4 reports the estimated $\beta^{I,S}$ coefficients, for each of the four idiosyncratic shocks when estimating regression 3 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 all four 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.

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. While many of these implied climate risk exposures might not initially appear intuitive, the key value added of the quantity-based approach to creating hedge portfolios is precisely to help identify industry exposures that go beyond what a simple narrative approach might have identified. However, before documenting that the resulting portfolios do indeed provide comparatively successful climate hedges—suggesting that the local quantity responses are successful at predicting

Table 3: Sample Summary Statistics

Panel A: Industry Summary Statistics

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

Panel B: Mutual Fund Summary Statistics

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

Panel C: Active Changes Summary Statistics

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

Note: Panel A shows, 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 report) and the sample size is 72,550 (note that active changes require two consecutive reports, which are not always available). Panel C shows summary statistics for the active industry changes as defined in Equation (2). The unit of observation is a fund-quarter-industry change and the sample size is 1,156,344.

Table 4: Industry-Specific Climate Quantity β Coefficients

GICS	Description	Avg.	Fat./Inj.	Indemnities	Extr. Temp.	Tran. Risk
2510	Auto & Components	0.74	1.00	0.80	1.00	0.18
4520	Tech. Hardw. & Equip.	0.48	0.74	1.00	0.59	−0.39
3010	Food & Staples Retailing	0.27	0.58	0.48	0.09	−0.05
3020	Food, Bev. & Tobacco	0.27	0.34	0.57	−0.09	0.25
2010	Capital Goods	0.26	0.27	0.46	−0.13	0.45
5010	Communication Services	0.23	0.67	0.29	−0.34	0.31
4530	Semiconductors & Equip.	0.17	0.71	0.17	−0.07	−0.14
4030	Insurance	0.13	−0.07	0.42	0.07	0.11
4020	Diversified Financials.	0.12	0.47	0.34	−0.17	−0.16
5510	Utilities	0.10	0.35	0.31	−0.08	−0.17
4010	Banks	0.09	0.65	0.21	−0.20	−0.30
3520	Pharma., Biotech., & Life Sc.	0.08	0.34	0.09	−0.03	−0.08
1510	Materials	0.07	−0.00	0.09	0.19	−0.00
1010	Energy	0.07	0.49	0.45	−0.45	−0.21
3030	Household & Pers. Prod.	0.06	0.38	−0.12	−0.21	0.21
4510	Software & Services	0.03	0.38	0.03	−0.14	−0.14
2520	Consum. Durables & Apparel	0.00	0.50	−0.56	−0.92	1.00
6010	Real Estate	−0.02	−0.08	0.17	0.20	−0.37
5020	Media & Entertainment	−0.04	−0.13	0.39	−0.11	−0.33
2030	Transportation	−0.05	0.49	0.73	−0.76	−0.64
2530	Consumer Services	−0.18	−0.65	0.05	−0.10	−0.01
2550	Retailing	−0.27	−0.44	0.15	−0.01	−0.80
3510	Health Care Equip. & Serv.	−0.41	0.03	−0.14	−0.52	−1.00
2020	Commercial & Prof. Serv.	−0.86	−1.00	−1.00	−1.00	−0.44

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

price responses to aggregate climate news shocks—we want to discuss possible economic mechanisms behind some of the observed quantity betas.

The largest positive climate quantity responses are 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 traditional 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 natural 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 profits in expectation. Indeed, by increasing the overall demand for insurance—something predicted by Lot & Haegeli (2021)—an increase in physical climate risks might even increase insurance firms' profitability, as limited insurance capital raises the pricing power of incumbents. This was highlighted by McKinsey analysts in Grimaldi et al. (2020), who wrote that:

> [Property and casualty (P&C) 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 energy sector, which one would expect to be among the 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 two of the four measures. First, regulation can discourage new innovation and entry from competitors; this can actually increase the present discounted value of incumbent's profits by increasing market power (see Ryan 2012). In addition, large energy companies play an important role in the transition to clean energy, either due to short-term reliability concerns about clean energy (see Elliott 2022, Magolin & Santino 2022) or their role in innovation in the clean energy space (see Pickl 2019, Cohen et al. 2020).

These discussions highlight the difficulties with 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 result visible in Table 4 is that the ordering and the sign of the industry betas are broadly correlated across different measures of climate belief shocks, despite the fact that the raw correlations of the shock measures are close to zero (Table 1). Panels A and B of Table 5 report the correlation and rank-correlation, respectively, of the industry-specific climate loadings obtained from running regression 3 for the period of 2015-2019. The industry coefficients are relatively correlated across the various shocks, though they are more consistent within the heat-based shocks than between heat-based and disclosure-based shocks. 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 only way to verify that our procedure is effective is to verify that it produces a portfolio that indeed hedges climate risk out of sample, something that we tackle in the next section.

Table 5: Across-Shock Correlation of Industry-Betas

Panel A: Pearson Industry Climate Beta Correlation

	Fat./Inj.	Indemnities	Extr. Temperature	Discl.: Tran. Risk
Fat./Inj.	1.00			
Indemnities	0.57	1.00		
Extr. Temperature	0.33	0.62	1.00	
Discl.: Transition Risk	0.32	−0.06	0.06	1.00

Panel B: Spearman (Rank) Industry Climate Beta Correlation

	Fat./Inj.	Indemnities	Extr. Temperature	Discl.: Tran. Risk
Fat./Inj.	1.00			
Indemnities	0.44	1.00		
Extr. Temperature	0.09	0.30	1.00	
Discl.: Transition Risk	0.23	0.06	0.12	1.00

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

Estimate Stability Across Subsamples. One way to investigate the magnitude of estimation error is to look at the stability of the quantity coefficients across subsamples. We explore this by splitting the sample randomly into two mutually exclusive subsamples, and looking at how the β^I estimates correlate across the two subsamples. We repeat the random splitting 100 times, and report the average correlation we obtain 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. Both approaches achieve relatively high average coefficient correlations, ranging from 0.26 to 0.43, indicating that the sample consistently picks up a common signal.

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

Finally, we investigate the shifts in industry climate coefficients over longer time horizons. While Table 4 shows the resulting coefficients when using data from 2015 to 2019,

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

Panel A: Random Split within Groups

Climate Shock	Stratified Spearman	Stratified Pearson	Fully Random Spearman	Fully Random Pearson
Heat: Fatalities/Injuries	0.35	0.38	0.33	0.35
Heat: High Indemnities	0.26	0.38	0.26	0.34
Heat: Extr. Temperature	0.35	0.43	0.30	0.38
Report: Transition Risk	0.34	0.38	0.34	0.36

Panel B: Random Split between Groups

Climate Shock	Fund Split Spearman	Fund Split Pearson	Period Split Spearman	Period Split Pearson	Location Split Spearman	Location Split Pearson
Heat: Fatalities/Injuries	0.35	0.38	0.14	0.20	0.23	0.21
Heat: High Indemnities	0.28	0.35	0.22	0.24	0.11	0.10
Heat: Extr. Temperature	0.31	0.39	0.17	0.26	0.30	0.35
Report: Transition Risk	0.31	0.30	0.05	0.06	0.02	0.01

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

i.e., the very end of our rolling 5-year sample, Appendix Table A.1 shows the coefficients when using data from 2010 to 2014. Notably, the automobile industry achieves the largest climate beta coefficients in both periods, suggesting a strong market belief in the opportunities arising from the transition to green energy; technological hardware and semiconductors are also strongly favored after idiosyncratic climate belief shocks during both periods. Similarly, retailing and real estate remain at the lower end of the ordering in both periods. Other industries (like consumer services and insurance) have instead shifted more during this period.

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

3 Quantity-Based Climate Hedging Portfolios

We next describe how we use the industry-specific climate quantity betas estimated in the previous section to build our quantity-based climate hedge portfolios. We also evaluate the out-of-sample hedging performance of these portfolios over the period of 2015-2019, and compare this performance against that of other approaches in the literature.

3.1 Portfolio Construction and Description

For each month t, industry I, and idiosyncratic belief shock S, we estimate $\beta_t^{I,S}$ from equation 3 using quarterly data on mutual fund portfolio compositions and idiosyncratic climate shocks from the five years prior to t. Using these estimates, we then construct the excess returns of the corresponding quantity-based hedge portfolio as:

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

where R_t^I is the value-weighted industry return and R_t^f denotes the risk-free rate. Note that each component of the portfolio is an excess return, so we do not need to scale the $\beta_t^{I,S}$.

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

Table 7: Portfolio Return Correlations

Panel A: Pearson Portfolio Return Correlation

	Fatalities/Injuries	Indemnities	Extr. Temperature	Discl.: Trans. Risk
Heat: Fatalities/Injuries	1.00			
Heat: Indemnities	0.66	1.00		
Heat: Extr. Temp.	0.51	0.38	1.00	
Discl.: Transition Risk	0.51	0.58	0.03	1.00

Panel B: Orthogonalized to Fama-French 3-factors

	Fatalities/Injuries	Indemnities	Extr. Temperature	Discl.: Trans. Risk
Heat: Fatalities/Injuries	1.00			
Heat: Indemnities	0.66	1.00		
Heat: Extr. Temp.	0.61	0.54	1.00	
Discl.: Transition Risk	0.27	0.31	0.11	1.00

Note: Panel A shows the monthly return correlation—constructed as in equation 4—for the period 2015 to 2019 among our four quantity-based hedge portfolios. Panel B shows the corresponding return correlation 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 four quantity-based hedge portfolios on the Fama-French factors. To identify the factor loadings of the quantity portfolios, we regress the portfolio returns on the returns of the market, size, and value factors:

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

Table 8 shows the regression results. Three of the four portfolios have a significant loading on the market factor and no loading on size; the "Fatalities/Injuries" and the disclosures-based portfolios have a significant positive loading on HML (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 30% of the variation in the quantity portfolios.

Table 8: Factor Exposures of Hedging Portfolios

	Heat: Quantity Hedge Portfolio			
	Fatalities/Injuries	High Indemnities	Extr. Temperature	Discl.: Trans. Risk
$R^M - R^f$	0.09**	0.16***	−0.03	0.19***
	(0.04)	(0.03)	(0.05)	(0.04)
SMB	0.01	−0.02	−0.03	−0.02
	(0.04)	(0.04)	(0.06)	(0.03)
HML	0.18***	0.03	0.06	0.10***
	(0.05)	(0.06)	(0.06)	(0.03)
Constant	0.01	−0.01	0.04	0.02
	(0.12)	(0.12)	(0.15)	(0.09)
R^2	0.36	0.31	0.05	0.52
N	60	60	60	60

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

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

3.2 Climate Hedge Targets

One challenge with designing portfolios that hedge climate risks is that there is no unique way of defining the hedge target. Climate change is a complex phenomenon and presents a variety of risks, including physical risks such as rising sea levels and transition risks such as the dangers to certain business models from regulation to curb climate change. Different risks may be relevant for different investors, and their realizations do not always co-move. In addition, climate change is a long-run threat, and we would 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) proposed 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 own risk exposures—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. Building on Engle et al. (2020), we consider innovations in these climate news indices from an AR(1) model as our hedge targets. Specifically, for a given climate news series c, we define these innovations—and hence our hedge targets—in month t as $CC_{c,t}$.

Engle et al. (2020). The Wall Street Journal (WSJ) and Crimson Hexagon Negative News (CHNEG) climate news indices created by Engle et al. (2020) are, to our knowledge, the first climate news series used as hedge targets. The first series captures the number of news articles in the WSJ dedicated to climate change (broadly assuming that "no news is good news"), the latter series 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 and cover the period of July 2008 to June 2017.

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 Media Climate Change Concerns (MCCC) index is available daily between January 2003 and June 2018. We aggregate to the monthly frequency by taking the average of the daily news series.

Faccini, Matin, and Skiadopoulos (2021). We include four of Faccini et al. (2021)'s climate news indices: news about international climate summits, global warming, natural disasters, and narrative indices. The international summits, global warming, and natural disasters indices measure news coverage of the respective topics; the narrative index is constructed by manually reading and classifying 3,500 articles. The international 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.

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" about risk realizations in the respective categories.

National Google searches. This climate news series is the national Google search interest in "climate change", capturing attention to climate change and its associated risks in 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 the hedging performance of two alternative approaches to constructing hedge portfolios: the narrative approach and the mimicking portfolio approach.

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. One example of such an approach would be to use environmental scores constructed by ESG data providers to build the portfolios, based on the prior view that high-ESG-score companies will fare better when climate risks materialize (see Engle et al. 2020). An alternative narrative approach would be to use specific groups of stocks (e.g., green energy stocks) under the prior that those groups' exposures to certain types of climate risks are predictable ex-ante. We build several portfolios using such a narrative-based approach.

Our first narrative portfolio takes positions in all U.S.-listed stocks covered by Sustainalytics ESG scores: the portfolio's position in each stock is the stock's ESG score percentile in each period, minus 50. 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.

A second strategy within the narrative approach 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 ConocoPhilips. This approach builds on the prior that realizations of climate change news should increase PBD's returns and decrease XLE's returns. Therefore, the hedging portfolio would go long PBD and short XLE.

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

Mimicking portfolio approach. The mimicking portfolio approach combines a predetermined 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 news series c:

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

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 5-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 one portfolio only is effectively equivalent to studying whether historically a long or a short position in that portfolio 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 industries. For this industry-based portfolio, we estimate both the standard mimicking portfolio, and a regularized version via LASSO, choosing the tuning parameter by cross-validation in an attempt to minimize the dangers from in-sample overfitting (we report the latter as it performs better).

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 a five-year rolling window of data; the narrative-based portfolios are unchanged over time. We focus on the post-2010 period to train our models, as investors likely paid very little attention to climate risks before 2010. As a result, we would not expect information on prices and quantities from before 2010 to be useful in building hedge portfolios today.

As a criterion to evaluate the various hedging approaches, we compare the out-of-sample correlations between the hedging portfolio returns and the AR(1) innovations to the various climate news series, $CC_{c,t}$.[16] Table 9 reports these out-of-sample correlations at the monthly frequency.[17] Each row in the table represents a 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 correlation coefficients 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 correlation coefficient of a hedge portfolio return with one of

[16] This 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.

[17] 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. Public announcements on policy changes can happen outside of market hours, such as when the EU introduces laws that affect U.S. companies. Moreover, sometimes news coverage can predate policy changes by writing in anticipation of international summits.

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 hedging portfolios are sorted top-to-bottom by this value.

Figure 2: Climate Hedge Performance of Various Portfolios

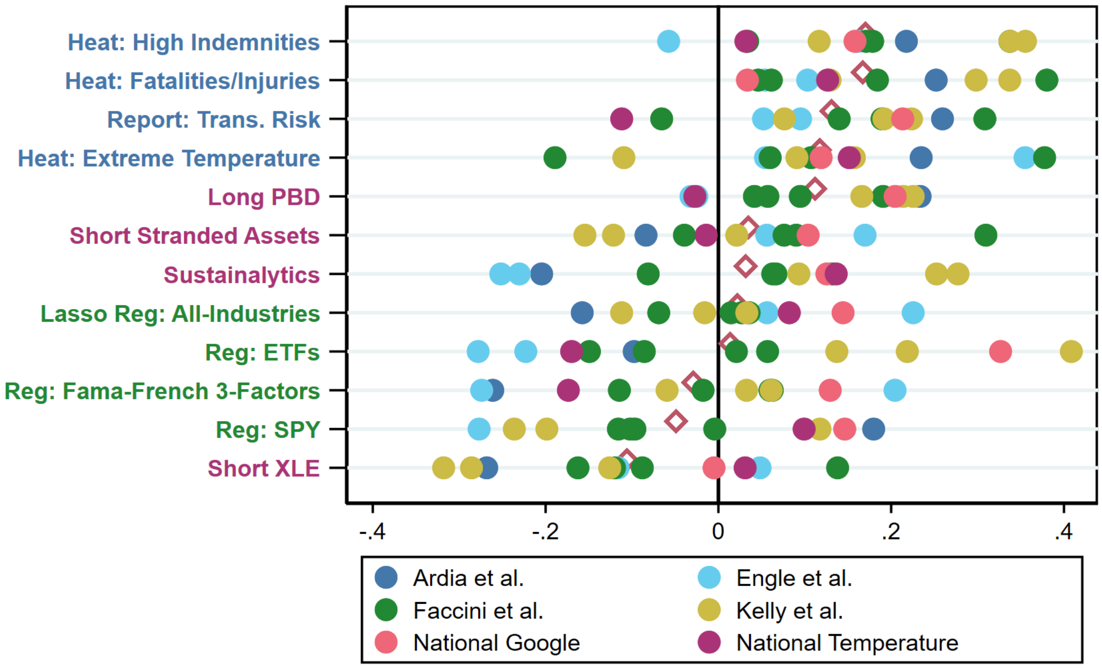

Note: Dot plot of monthly out-of-sample return correlations for various climate hedge portfolios with various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Different colors represent different groups of climate news series.

The first four 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. Figure 2 (blue rows) shows the same results. At the top of the figure, the "Heat: Fatalities/Injuries" portfolio correlates positively with *all* climate news innovations, while "Heat: High Indemnities" can hedge all but the CHNEG series and has the highest average out-of-sample correlation with our series. The "Heat: Extreme Temperature" and "Disclosure: 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 summits and global warming indices, Kelly (2021)'s general and physical risk indices, Ardia et al. (2020)'s MCCC index, Engle et al. (2020)'s WSJ index and the Google national searches series. This suggests that our various quantity-based portfolios perform well in terms of hedging a variety of climate risks, spanning 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.

Next, rows 5-8 of Table 9 show the performance of the different narrative-based portfolios.

Table 9: Climate Hedge Performance of Various Portfolios

	Faccini, Matin, Skiadopoulos				Kelly [et al.]			Engle et al.		Ardia et al.	Google	Temp.
	IntSummit	GlobWarm	NatDis	Narrative	General	Transitional	Physical	WSJ	CHNEG	MCCC	National	National
Heat: Fatalities/Injuries	**.38**	**.18**	**.05**	**.06**	**.29**	**.12**	**.32**	**.05**	**.10**	**.25**	**.04**	**.13**
Heat: High Indemnities	**.34**	**.18**	**.17**	**.03**	**.30**	**.09**	**.32**	**.16**	−.06	**.22**	**.18**	**.05**
Heat: Extr. Temperature	**.38**	**.11**	−.19	**.06**	**.05**	−.14	**.11**	**.35**	**.05**	**.23**	**.15**	**.17**
Report: Transition Risk	**.19**	**.14**	**.31**	−.07	**.22**	**.07**	**.19**	**.10**	**.05**	**.26**	**.22**	−.11
Long PBD	**.06**	**.09**	**.19**	**.04**	**.21**	**.17**	**.23**	−.02	−.03	**.23**	**.20**	−.03
Short XLE	−.09	−.16	−.12	**.14**	−.32	−.13	−.29	−.12	**.05**	−.27	−.01	**.03**
Short Stranded Assets	−.04	**.08**	**.31**	**.09**	−.15	−.02	−.18	**.06**	**.17**	−.08	**.14**	**.01**
Sustainalytics	**.13**	−.08	**.06**	**.07**	**.24**	**.26**	**.08**	−.25	−.23	−.20	**.14**	**.14**
Reg: FamaFrench 3-Factors	−.02	**.06**	**.06**	−.11	**.06**	−.06	**.03**	**.20**	−.27	−.26	**.13**	−.17
Reg: ETF + FamaFrench	−.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**
Reg: All-Industries	−.12	**.06**	−.13	**.03**	−.10	−.11	**.06**	**.00**	**.14**	**.10**	**.17**	**.17**
Lasso Reg: All-Industries	**.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 constant along the rows, the mimicking portfolios show in each cell the portfolio that was specifically trained on the respective climate news series.

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. 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 row is a short position in XLE, motivated by the fact that XLE is dominated by polluting companies; the same reasoning pins down the sign of the trading strategy of the other portfolios in this group. The hedging performance of these narrative portfolios is mixed, with the worst results given by the short XLE position, and the best results 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.

The last five 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 relatively 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 MCCC 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 delivers the best out-of-sample climate hedging performance. Among the alternative approaches, with only little historical data available on periods when climate risk was potentially priced, mimicking portfolio approaches do not deliver successful climate hedges. In contrast, using observed asset-level characteristics to build hedging portfolios is potentially promising—PBD is able to hedge all but 3 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. (2019) 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 our quantity-based climate hedge portfolios. Figure 3 shows the out-of-sample correlations similar to Figure 2 for variations of the three heat-based quantity portfolios: "Heat: Fatalities/Injuries" in Panel A; "Heat: High Indemnities" in Panel B; and "Heat: Record Temperature" in Panel C. Figure 4 reports the same robustness checks for the disclosure-based hedge portfolio.

We consider the following variations: add the interaction of fund and time fixed effects to regression 3; define industry exposure changes in terms of current prices (previously defined as $PassiveChanges$); change the winsorization of the active changes (the baseline winsorizes at 1%, we report here the results using trimming instead of winsorizing at 1%, and alternatively winsorizing at 0.5% or 2%, and without winsorization); define the relevant universe of funds with only the CRSP objective code or only the Thomson Reuters IOC; do not weight industry changes by their relative market size; aggregate climate shocks from counties to commuting zones; use three-year rolling windows (instead of five); and lastly, only keep funds where all advisers reside in the same county. As Figures 3 and 4 show, most of these changes have minimal effects on the performance of the measures. Overall, the results appear robust to these variations in portfolio construction.

Particularly interesting is the robustness test based on the three-year rolling window (instead of five) because it shows the hedging ability of the quantity approach when very little time-series data is available or when frequent structural shifts occur. Appendix Figure A.4 shows how well the quantity portfolios perform when compared to other strategies based on three-year windows. In particular, we replicate the mimicking portfolios using only three-year windows of data and leave the narrative-based portfolios unchanged. The quantity portfolios achieve slightly lower correlations, on average, while still maintaining their position at the top of the chart. The mimicking portfolios continue to achieve very mixed performance. Therefore, this provides additional evidence that the quantity-based approach can be superior even when very little time-series data is available.

Figure 3: Climate Hedge Performance - Robustness Tests

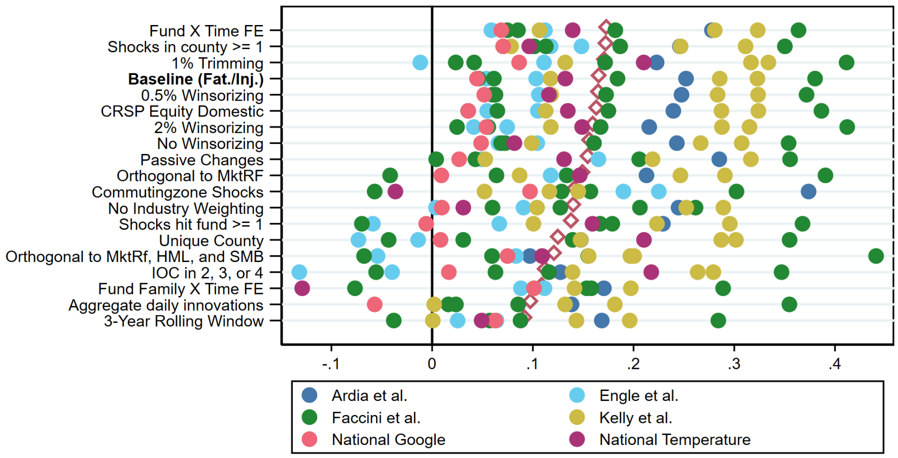

(a) Heat: Fatalities/Injuries

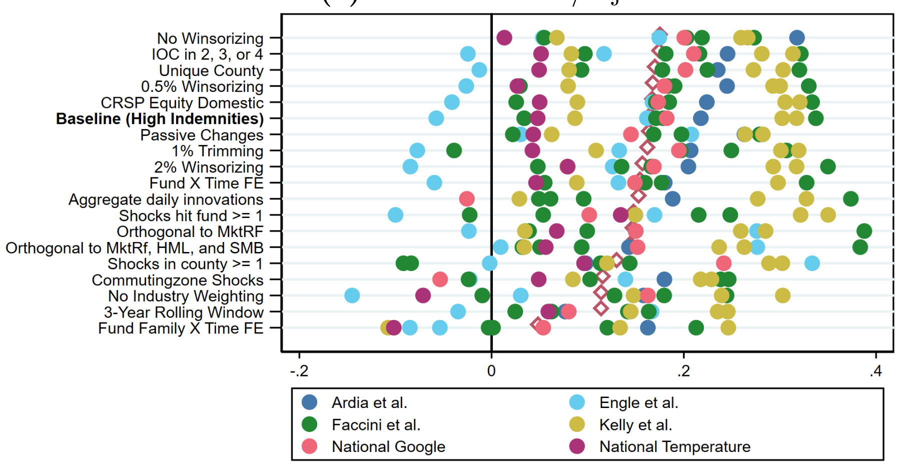

(b) Heat: High Indemnities

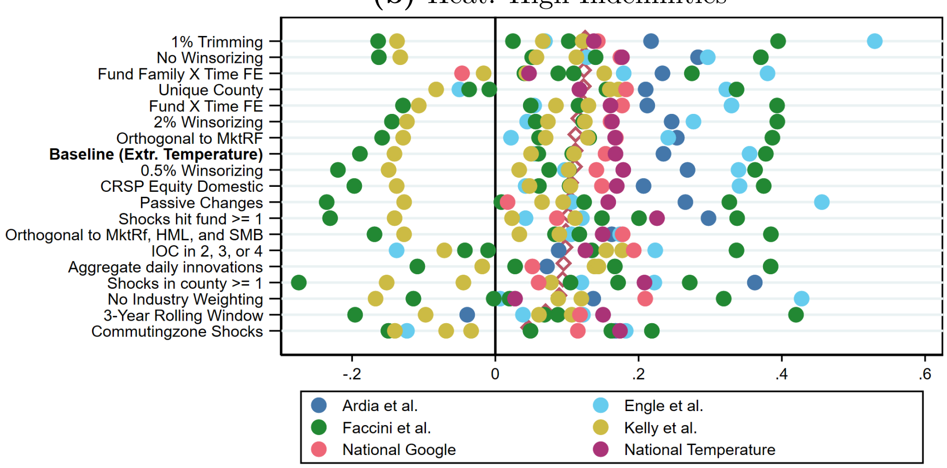

(c) Heat: Extreme Temperature

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

Figure 4: Climate Hedge Performance - Robustness Tests

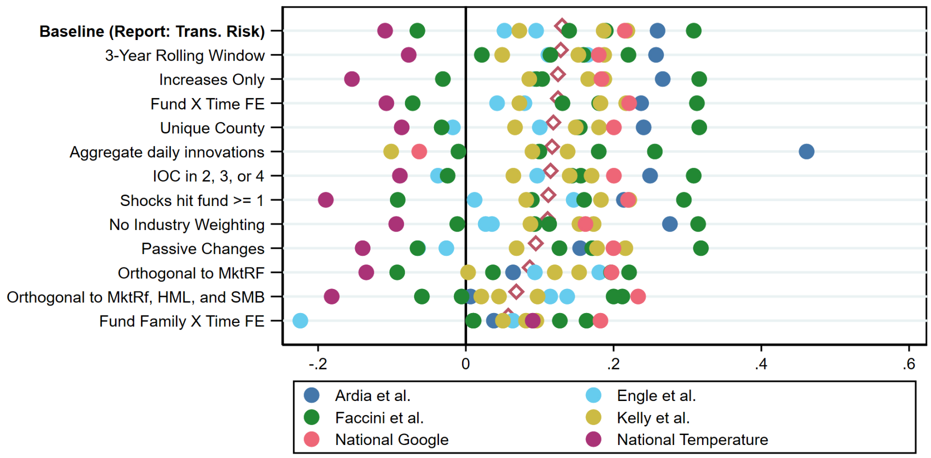

(a) Report: Trans. Risk

Note: Same as Figure 3, but for the disclosure-based hedging portfolio.

4 Different Dynamic Hedging Approaches: Pros & Cons

After showing that quantity-based approaches appear to provide the best hedge portfolios for climate risk, we next review the circumstances under which the different approaches (quantity-based, narrative, and mimicking portfolio) are expected to work better in general. Two particularly important elements on which we will focus are (1) to what extent the procedures require knowledge about the true exposures of different companies to climate risks, and (2) the extent to which these exposures are time-invariant or changing over time.

The mimicking portfolio methodology takes a purely statistical approach to portfolio construction. It requires little input from the researcher beyond the choice of the basis portfolios used in the projection, and instead relies on the time series information (the historical covariance between the target and asset returns) to choose the portfolio weights. This approach works well as long as the time series is sufficiently long ($T \to \infty$) and asset risk exposures are stable over time: in that case, one can choose a sufficiently large set of basis assets (that is, one can allow $N \to \infty$) and obtain asymptotically the optimal hedging portfolio (see derivation in Giglio and Xiu 2021). This approach suffers when T is small or asset exposures are time-varying: first, because a small T induces a large variance in the estimates of the covariance of prices with the hedging target, and second because it limits the number of basis assets that can be feasibly used (N). This is especially an issue when considering emerging risks, like climate change, for which T can be quite limited and where firms' strategic changes due to the climate transition affect their exposures over time.[18]

[18] For 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.

On the other end of the spectrum, the narrative approach does not rely on historical time series at all. Rather, it requires the investor to have correct beliefs about the different assets' exposures to the target. As the previous section showed, identifying ex-ante which stocks would stand to gain or lose when climate risk materializes can be difficult.

The quantity-based approach does not rely either on the investor 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 somewhat weakens these requirements by using *cross-sectional* information instead of time-series information 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 and quickly learn when asset exposures have changed.

While the quantity-based approach has some important benefits relative to the mimicking portfolio approach, it has stronger data requirements in a variety of dimensions. In particular, the quantity-based approach requires the identification of idiosyncratic belief shocks, i.e. shocks that move investor beliefs about the aggregate risk (e.g., aggregate climate change) yet only affect few investors concurrently. On this front, we believe that both the location-based and disclosure-based approaches to identifying idiosyncratic belief shocks will extend to other aggregate risks, something that we verify in Section 5. Second, the quantity-based approach requires the researcher to observe portfolio holdings or trading data, which are generally less widely available than price data, in particular for assets outside of equities.

Second, since the portfolio weights are determined by other investors' trading responses to the belief shocks, the quantity-based approach relies on the *average* investor's information (in our application, the average mutual fund manager's information). Indeed, one way to look at the quantity-based approach is as a way to identify the "narratives" around asset exposures of the average investor. But, what if mutual fund managers are wrong on average in their assessment of different assets' true climate risk exposures? In that case, the hedge portfolio we build can still hedge aggregate climate news in the short run, as long as the average response to idiosyncratic climate belief shocks corresponds to the average response to global climate 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. 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 climate risk in the short term. Yet, in the very long run, any portfolio that is long car stocks will ultimately lose once climate shocks materialize, and the true climate risk exposures are revealed. To sum up, the short-term hedging performance only relies on *consistent* behavior of investors; the long-term hedging performance relies on markets (i.e. the average investor) being right about the risk exposures.

It is useful to note that incorrect risk perceptions by the average investor also present challenges for the mimicking portfolio approach, which relies on price movements determined by the average investor. The benefit of the quantity-based approach is that it provides more information that allows researchers to learn more quickly when the average investor's perception of an asset's climate risk exposure has changed. The narrative approach has the opposite challenge. Under the (rather ambitious) 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, and constructs their portfolio accordingly, this portfolio will likely have somewhat better long-run hedging properties. However, in the short-run, while the researcher disagrees with the average market participant on the risk exposures, the narrative approach will not be able to provide much hedging of the arrival of news about climate risks.

5 Hedging Macro Factors

While the main focus of this paper is on hedging climate risks, quantity-based portfolios can be built to hedge any other macroeconomic risk. In this section we explore two other applications: building portfolios aimed at hedging national unemployment and house price shocks. In each case, we identify "idiosyncratic belief shocks" using local versions of the macro shocks, and study the trading response of funds located where these shocks occur. The motivation is a connection between experienced idiosyncratic shocks and beliefs about aggregate events, similar to that we described for the case of climate risks. For example, Kuchler & Zafar (2019) 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.

5.1 Unemployment & House Price Shocks

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

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

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

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

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

5.2 Validation of the macro hedge portfolios

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

Table 10 shows the out-of-sample correlation of our constructed macro risk hedge portfolios with AR(1) innovations of the national unemployment series and the growth rate of the house price index. The different versions of the mimicking portfolios mostly obtain positive out-of-sample correlations, with varying performance as we change the base assets. For example, using all industries without lasso works better for hedging house price shocks (correlation of 0.15) than unemployment shocks (-0.03), whereas using industries with lasso works better for unemployment shocks (0.19) than for house price shocks (0.07).

The quantity-based portfolios have a good performance with respect to the macro series that they target. The house-price-based quantity portfolio performs well in hedging aggregate shocks to house price growth (correlation of 0.13) but not unemployment shocks, whereas the unemployment-based quantity portfolio hedges unemployment shocks well (0.21) but not housing shocks. These results suggest that the quantity-based portfolios are able to target quite specifically the different macro risks, and offer a more consistent hedging performance compared to the mimicking portfolio approach.

Table 10: Macro Hedge Performance

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

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

6 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. 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.

References

Alok, S., Kumar, N. & Wermers, R. (2020), 'Do Fund Managers Misestimate Climatic Disaster Risk', *The Review of Financial Studies* **33**(3), 1146–1183.
URL: *https://doi.org/10.1093/rfs/hhz143*

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* .

Armona, L., Fuster, A. & Zafar, B. (2019), 'Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment', *Review of Economic Studies* .

Bachmann, R., Berg, T. O. & Sims, E. R. (2015), 'Inflation expectations and readiness to spend: Cross-sectional evidence', *American Economic Journal: Economic Policy* .

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.

Bakkensen, L. A. & Barrage, L. (2021), 'Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics', *The Review of Financial Studies* .

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* .

Barnett, J. (2020), 'Global environmental change ii: Political economies of vulnerability to climate change', *Progress in Human Geography* .

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

Berg, F., Koelbel, J. F. & Rigobon, R. (2019), *Aggregate confusion: The divergence of ESG ratings*, MIT Sloan School of Management Cambridge, MA, USA.

Berk, J. B. & van Binsbergen, J. H. (2016), 'Assessing asset pricing models using revealed preference', *Journal of Financial Economics* **119**(1), 1–23.
URL: *https://www.sciencedirect.com/science/article/pii/S0304405X1500149X*

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* .

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.
URL: *https://doi.org/10.1093/qje/qju033*

Ceccarelli, M., Ramelli, S. & Wagner, A. F. (2021), 'Low-carbon mutual funds', *Swiss Finance Institute Research Paper* (19-13).

Cen, X. (2021), 'Household wealth and entrepreneurial career choices: Evidence from climate disasters', *Available at SSRN 3803436* .

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.
URL: *https://doi.org/10.1093/restud/rdy016*

Chen, Q., Goldstein, I. & Jiang, W. (2010), 'Payoff complementarities and financial fragility: Evidence from mutual fund outflows', *Journal of Financial Economics* .

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.

Conlin, M., O'Donoghue, T. & Vogelsang, T. J. (2007), 'Projection bias in catalog orders', *American Economic Review* **97**(4), 1217–1249.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.97.4.1217*

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.

Dou, W., Kogan, L. & Wu, W. (2021), 'Common fund flows: Flow hedging and factor pricing', *Jacobs Levy Equity Management Center for Quantitative Financial Research Paper* .

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'.

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

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

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

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* .

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

Giglio, S., Maggiori, M., Rao, K., Stroebel, J. & Weber, A. (2021), 'Climate Change and Long-Run Discount Rates: Evidence from Real Estate', *The Review of Financial Studies* .

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

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.
URL: *https://www.jstor.org/stable/26727317*

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

Hsu, P.-H., Li, K. & Tsou, C.-Y. (2022), 'The pollution premium', *Available at SSRN 3578215* .

Igami, M. & Uetake, K. (2019), 'Mergers, Innovation, and Entry-Exit Dynamics: Consolidation of the Hard Disk Drive Industry, 1996-2016', *The Review of Economic Studies* **87**(6), 2672–2702.
URL: *https://doi.org/10.1093/restud/rdz044*

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*a*), 'The importance of climate risks for institutional investors', *The Review of Financial Studies* **33**(3), 1067–1111.

Krueger, P., Sautner, Z. & Starks, L. T. (2020*b*), 'The Importance of Climate Risks for Institutional Investors', *The Review of Financial Studies* **33**(3), 1067–1111.
URL: *https://doi.org/10.1093/rfs/hhz137*

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* .

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.

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

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.
URL: *https://www.sciencedirect.com/science/article/pii/S2211467X19300574*

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* .

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?’.

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’.

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: Historical Industry Climate-β Coefficients (2010-2015)

GICS	Description	Avg.	Fat./Inj.	Indemnities	Extr. Temp.	Tran. Risk
2510	Auto & Components	0.83	0.31	1.00	1.00	1.00
4530	Semiconductors & Equip.	0.29	1.00	−0.51	0.27	0.40
2530	Consumer Services	0.23	0.98	−0.68	0.03	0.60
4510	Software & Services	0.22	0.27	0.27	−0.19	0.54
1510	Materials	0.10	0.72	−0.31	0.18	−0.17
1010	Energy	0.08	0.41	−0.28	0.03	0.15
2010	Capital Goods	0.08	0.27	−0.64	0.64	0.04
4520	Tech. Hardw. & Equip.	0.07	0.75	0.32	−0.49	−0.28
4010	Banks	−0.03	0.46	−0.44	0.02	−0.16
5020	Media & Entertainment	−0.05	0.05	−0.18	0.13	−0.21
3030	Household & Pers. Prod.	−0.07	−0.01	−0.48	−0.05	0.24
2520	Consum. Durables & Apparel	−0.15	0.93	0.16	−0.96	−0.74
4020	Diversified Financials.	−0.16	0.07	−0.14	0.06	−0.63
3520	Pharma., Biotech., & Life Sc.	−0.17	−0.44	−0.23	0.19	−0.21
3510	Health Care Equip. & Serv.	−0.19	0.02	−0.68	−0.21	0.12
2020	Commercial & Prof. Serv.	−0.23	0.25	−0.30	0.14	−1.00
5010	Communication Services	−0.25	−0.43	−0.46	0.06	−0.18
2030	Transportation	−0.26	0.19	−0.49	−1.00	0.25
4030	Insurance	−0.27	−0.16	−0.50	−0.07	−0.33
6010	Real Estate	−0.36	−0.56	−0.92	−0.26	0.31
3020	Food, Bev. & Tobacco	−0.39	−0.47	−0.79	−0.08	−0.22
3010	Food & Staples Retailing	−0.44	0.11	−1.00	−0.40	−0.45
5510	Utilities	−0.48	−0.62	−0.90	−0.16	−0.26
2550	Retailing	−0.66	−1.00	−0.34	−0.49	−0.83

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

A.2 Figures

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

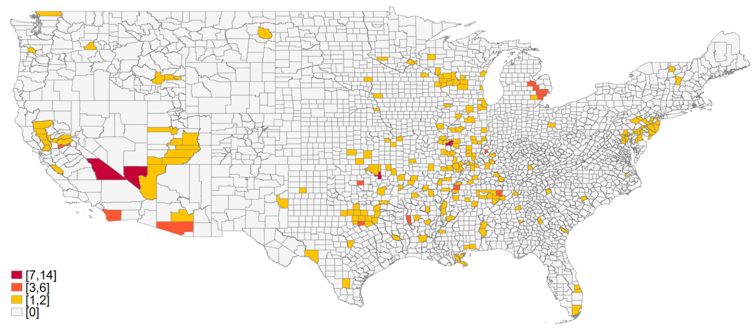

Note: Distribution of the "Heat: Fatalities or Injuries" climate shock from 2010 to 2019. The color-coding shows the number of county-months that experienced the "local" climate shock during the interval.

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

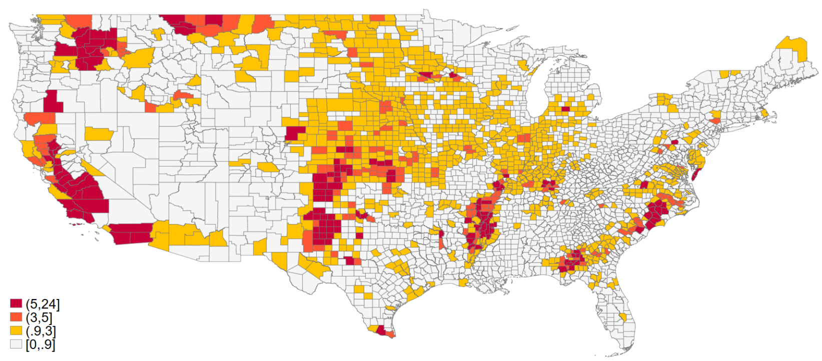

Note: Distribution of the "Heat: High Indemnities" climate shock from 2010 to 2019. The color-coding shows the number of county-months that experienced the "local" climate shock during the interval.

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

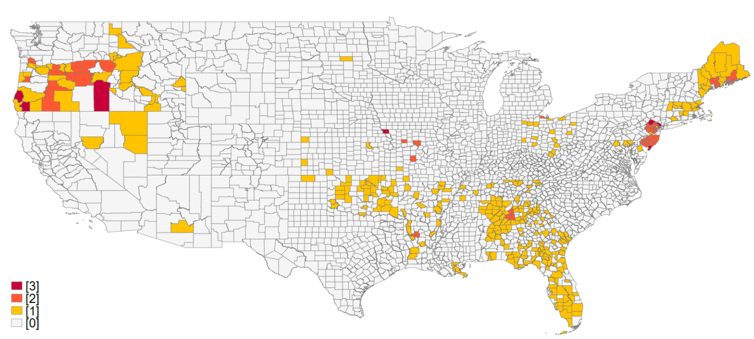

Note: Distribution of the "Heat: Extreme Temperature" climate shock from 2010 to 2019. The color-coding shows the number of county-months that experienced the "local" climate shock during the interval.

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

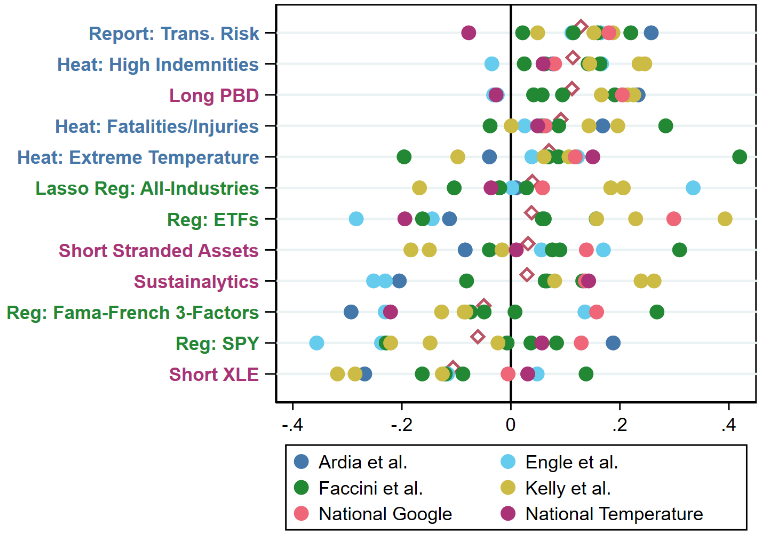

Note: Dot plot of monthly return correlations using 3-year rolling windows; see Table 9.