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

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

This version: August 12, 2021.
PRELIMINARY AND INCOMPLETE

Abstract

We propose a new methodology to build hedging portfolios for climate change risks. Our procedure, which we refer to as *quantity-based* approach, aggregates cross-sectional information on investor portfolio reallocations in response to "local" climate shocks to build a hedging portfolio for "global" shocks. Specifically, we study how mutual funds change their portfolio allocation when they experience a local heat shock in the fund adviser's geographical proximity, which has been shown to shift their beliefs about climate risks. We use this to anticipate how investors will reallocate their investment when "global" climate shocks occur, which shift the beliefs and asset demand of many investors simultaneously and ultimately move equilibrium prices. The quantity-based hedging portfolio appreciates in value when a global climate shock occurs because it holds stocks that investors tend to buy after experiencing a (local) climate shock; it therefore provides a hedge for climate risks. We show empirically that this approach yields superior out-of-sample hedging performance compared to traditional methods, like ESG-score-based portfolios or standard mimicking portfolios. The key advantage of the procedure is that it learns from cross-sectional responses rather than from purely time-series information, which is limited in the case of climate change. Finally, we also demonstrate the efficacy and versatility of the quantity-based approach by constructing successful hedge portfolios for unemployment and house price indices from regional unemployment and house price levels.

*This is a working paper version as of August 12, 2021. Send correspondence to Johannes Stroebel, NYU Stern, 44 West 4th Street, New York, NY 10012. Email: johannes.stroebel@nyu.edu.
†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.

1 Introduction

Climate change is one of humanity's most significant challenges in the 21st century, and the increased frequency of natural disasters in recent years has made this more evident than ever. At the same time, investors awareness' of the risks related to climate change has developed fast over the last decade. This has created a strong demand for financial products that can be used to hedge climate risk in its different forms of physical or transition risk. At the moment, there is a scarce supply of financial instruments like derivatives that might offer a direct and targeted hedge against specific climate risks. As a consequence, the most viable way for investors to insure against climate change is to build hedging portfolios using available assets like equities. These assets are exposed to many more risks beyond the climate ones, which makes the construction of climate hedging portfolios – that isolate climate risk exposures to exposure from all other risks – extremely hard in practice.

In this paper, we propose a novel methodology to build hedging portfolios for specific risks (like climate risks) that exploits the availability of *cross-sectional* information on investor responses to "local" shocks to build a hedge portfolio against "global" shocks. The approach is particularly suitable for the climate-change application as it relies mostly on cross-sectional, rather than time-series information, which is quite limited for the case of climate change. Our approach, which we refer to as a *quantity-based* approach, studies how mutual funds adjust their portfolios after experiencing a regional heat shock in their adviser's geographical proximity – a "local" heat shock. It then uses this information to predict the aggregate response of investors when "global" climate shocks occur (climate shocks that are large enough to move equilibrium prices). Intuitively, by anticipating which stocks will be bought and sold when a climate shock occurs, we can construct a global climate-hedging portfolio by over- and under-weighting those stocks appropriately. While the main application of our methodology in this paper is to hedging climate risks, we also show that it can be applied to various other risks: we build quantity-based hedging portfolios for unemployment and housing price shocks.

Our approach for building the hedging portfolio starts from identifying a "local" driver of investor beliefs about climate change, a shock that only affects a geographically small area and whose occurrence can make the local investors update their views on climate risks. While there are various calamities associated with climate change, previous literature identified extreme heat most consistently as a weather event that shifts climate change beliefs. Not only do Egan & Mullin (2012) and Deryugina (2013) show that prolonged periods of heat

shift climate change beliefs, Joireman et al. (2010) and Li et al. (2011) show that single days of extreme heat can suffice for changing opinions. Fownes & Allred (2019) and Sisco et al. (2017) highlight the advantages of tracking relative temperature deviations over absolute temperatures. All of these papers show that extreme temperatures, in one way or another, affect opinions on the existence and severity of climate change.[1]

We then study how investors that are exposed to these local shocks change their portfolio allocations. Specifically, we identify which mutual funds whose advisers reside in afflicted areas, and we contrast them with funds whose advisers were not affected by the event. By looking at the changes in mutual fund portfolios around the "local" shocks, we identify which stocks are mostly bought (or sold) by investors that experience local shocks.

Finally, we use this information to build a hedging portfolio for "global" climate shocks. If the changes in portfolios following "local" shocks is informative about how investors reallocate their investments as their climate beliefs shift, we should expect similar reallocations to occur (but on a larger scale) when a global climate shock occurs, since the occurrence of a global shock will shift many investors' beliefs. In the event of a global shock, in fact, the correlated shift in beliefs (and therefore change in demand) from many investors will ultimately affect the equilibrium prices of the assets. A portfolio that has bought and sold stocks based on the response of investors to the local shocks (buying those stocks that investors tend to buy after a bad climate shock and selling those that the investors tend to sell) should yield positive returns after a bad global climate shock (when a mass of investors will implement similar trades), thus providing a hedge against the realization of aggregate climate shocks.

As an example, suppose that fund managers that observe a (local) heat shock, and update negatively about climate change, on average, increase their allocation in green energy companies as a response. Based on this information, we would build a hedging portfolio going long green energy companies. When a global shock occurs, we predict that many investors would demand green energy companies, increasing their price and therefore yielding a positive return on the hedging portfolio. This idea is the foundation of our quantity-based portfolio construction.

The literature, as well as the industry, has proposed several alternative ways to build hedging portfolios. We can classify them into two groups. The first group uses firm characteristics (e.g., ESG scores or industry classifications) together with prior information about

[1] In the Online Appendix, we also build quantity-based climate hedges using other types of weather events. Although previous research also showed that droughts, tropical storms, wildfires, floods, or tornadoes can shift climate change beliefs (Konisky et al. 2016, Sisco et al. 2017), they are not as strong or as consistent a driver of beliefs as extreme temperatures. This is confirmed by our empirical analysis, that shows worse hedging performance for quantity-based portfolios built using those events.

the climate exposures related to those characteristics to build the hedging portfolio. For example, ESG scores have been studied extensively in the literature as a way to capture apriori climate change exposures (Engle et al. 2020, Pástor et al. 2020, Hoepner et al. 2018). The simplest method to build such a portfolio is to buy high-ESG-score and sell low-ESG-score companies. Sometimes this method builds long-short portfolios of groups or stocks or industries; for example, a simple climate-hedging strategy might involve buying a clean-energy ETF like PBD, or selling a traditional energy ETF like XLE. This approach has the advantage that it does not require long time series to be implemented but has the important downside of relying on a relatively arbitrary choice of characteristics (even among ESG scores, there are many different versions often in disagreement with each other).

The second alternative method to constructing hedge portfolios is the mimicking portfolio approach, where climate change risk factors are projected onto a set of portfolio returns using time-series information. This approach necessitates the definition of a climate change hedge target. Engle et al. (2020) pioneered this method's application to climate change risks by constructing hedge targets from innovations in climate news series. Then, they use portfolios sorted on climate characteristics and the size, value, and market factors as projection portfolios. Intuitively, they constructed hedge portfolios using in-sample fitting on price data. The hedging portfolio is built by replicating what portfolio has been successful – in historical data – in hedging the global climate shocks.

In the paper, we produce the quantity-based hedging portfolio for climate risks and compare its performance with the other two approaches used in the literature. We use data from 2010 to 2014 to train the models, and data from 2015 to 2019 to evaluate them out of sample (prior to 2010, climate risks were hardly incorporated into market prices or affected investor behavior, making all of these approaches difficult to implement). We construct different "global" climate shocks to be used as the target of the hedge, drawing on a recent but rapidly expanding literature (for example, we look at climate change mentions in national newspapers, Google searches for climate change, and so on), and evaluate how the different approaches perform in hedging each series.

We find several interesting patterns. First, at a broad level, hedging climate risks is hard. Few approaches manage to achieve more than 20% out-of-sample correlation with the climate shock series, thus extending the finding in Engle et al. (2020). In addition, the vast majority of the hedging portfolio methodologies do not achieve consistent performance across the various climate shocks. That is, some approaches provide better hedging for some measures of global climate risks, and other approaches for other measures. This, of

course, could be a reflection that climate risk is an inherently complex phenomenon, and the different procedures might be better able to hedge different aspects of climate risks. On the other hand, it may also again reflect the inherent difficulty in building these hedge portfolios. Last but not least, our quantity-based portfolio manages to hedge *all* of the measures, yielding positive out-of-sample correlations between the hedging portfolio returns and each climate shock, and reaching correlations above 30%. This provides direct evidence of the good out-of-sample performance of the quantity-based approach.

We structure our paper as follows. First, in Section 2, we describe a simple model of quantity-based information that justifies our approach to constructing hedge portfolios. We then describe the regional heat shocks that we include in our analysis in Section 3, and we list the climate news indices, which we use to validate the quantity-based hedge portfolios, in Section 4. Sections 5 and 6 describe our novel quantity-based approach exemplary for the case of climate change realizations. We validate the quantity climate change hedge portfolios in Section 7 and further demonstrate the efficacy of our methodology by hedging unemployment and housing prices in Section 8. We conclude this paper by summarizing the results and listing ideas for future research.

2 A Simple Model of Quantity-Based Information

In this section, we describe a simple model of quantity-based information that can be used to interpret and justify our quantity-based approach to building climate change hedge portfolios.

2.1 Setup

Consider a unit continuum of investors $i \in [0,1]$ who can choose to invest in either of two securities, security A and security B. Investor i has demand for security A given by: $q_A\big(p_A, \epsilon_A(i)\big)$, where p_A indicates the (relative) price of security A, and $\epsilon_A(i)$ indicates investor i's beliefs on the (relative) future profitability 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 $\frac{\partial f}{\partial p_A} < 0$.

This setup implies the market clearing condition: $\int_{i=0}^{i=1} q_A\big(p_A, \epsilon_A(i)\big)di = \bar{A}$, where $\bar{A}$ is the supply of the security. The equilibrium of this economy is characterized by an equilibrium price p_A^* and equilibrium asset allocations $q_A^*(i)$. We focus on the equilibrium in market A, as market B clears by Walras' law.

Finally, we decompose individual investors' beliefs into a common component ν_A and an investor-specific component $\omega_A(i)$, so that $\epsilon_A(i) = \nu_A + \omega_A(i)$, where ν_A. We interpret the

common belief ν_A as capturing the beliefs of investors that all observe a "global" (common) shock. $\omega_A(i)$, instead, represents a "local" shock that is observed by only a small number of investors (in the model, only by investor i), and is not observed by any other investor.

Note that we deliberately do not impose assumptions on the origin of the global and local beliefs; what differentiates them is whether they are common to all or specific to only some investors. Also, note that in this simple model investors simply "agree to disagree": there is no learning from prices about the beliefs or information of other investors. While it would be easy to extend this model to incorporate a rational expectations or behavioral mechanism driving these results, this is not the main focus of our analysis. All that matters is that investors experience local shocks that shift beliefs.

2.2 Local Shocks

We start by studying what happens to equilibrium prices and quantities when one investor changes their beliefs about stock A, e.g., in response to a "local" shock. By the chain rule:

$$\frac{\partial q}{\partial \omega_A(i)} = \frac{\partial q}{\partial \epsilon_A(i)}.$$

We can also calculate the change in equilibrium prices. Note that:

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

Thus, $\frac{\partial p_A^*}{\partial \omega_A(i)} = 0$. However, since investor i's demand changes, $\frac{\partial q^*}{\partial \omega_{i,A}} \neq 0$. Given market clearing, this demand is absorbed by some other agent in the economy (who exactly is not relevant for our purposes).

In words, if a single investor receives a "local" belief shock:

(i) This agents' equilibrium asset allocation changes.
(ii) Equilibrium prices do not change.
(iii) Because prices don't change, differences in their equilibrium asset allocation are directly informative about the demand sensitivity to beliefs $\frac{\partial q}{\partial \epsilon_A(i)}$.

2.3 From Local Shocks to Common Shocks

Suppose now there is global news about stock A. By the implicit function theorem and the chain rule, equilibrium price responses are given by:

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

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

Putting this together with the earlier result shows how idiosyncratic quantity responses can be used to predict national price responses. Intuitively, by studying how investors react to local shocks (that have no effect on the equilibrium price), we can predict how their demand shifts in response to news. Global shocks move the demand function of many investors simultaneously, leading to price movements that can be predicted by the response to local news.

3 Local Climate Shocks

In this section, we describe the local extreme heat measures we use to discipline the quantity hedge portfolios. Additionally, Table 1 provides a more simplified overview of the constructed heat shocks, and Appendix Figures A.2 to A.5 visualize the event distributions in a heatmap. We validate the climate shocks' abilities to impact climate change awareness in Subsection 3.2.

Climate Shock	Event Description	Frequency	
		Monthly	Sample
Heat: Any	Property damage, injuries, or fatalities	0.15%	1.65%
Heat: Fatalities/Injuries	Injuries or fatalities	0.13%	1.32%
Heat: Indemnities	90th percentile indemnity payments	0.79%	0.54%
Heat: Record Temperature	10-year record of 3-month-average county temperature	0.91%	1.06%

Table 1: An overview of the constructed heat measures. All shocks are coded as dummy variables, i.e., they indicate the existence of a severe heat event. The event description lists the dummy criteria. The "monthly" frequency shows the share of months across all US counties from 2010 to 2020 that experience the climate event. The "sample" frequency shows the share of observations in our final sample that experience the climate event. The differences in "monthly" and "sample" frequency arise from mutual fund advisers being concentrated in relatively few, high population density counties. Also, the "sample" frequency uses three-month windows instead of monthly, generally leading to a higher prevalence.

3.1 Shock Descriptions

SHELDUS - Climate Event Data

The Spatial Hazard Events and Losses Database for the United States (SHELDUS) provides county-level information on various events. We include their reported extreme heat events, which are maintained by the National Centers for Environmental Information (NOAA). We classify our baseline heat shock, "Heat: Any", as resulting in either property

damages, injuries, or fatalities. We also consider a more severe version, "Heat: Fatalities/Injuries", that requires recorded fatalities or injuries.

SHELDUS - Crop Indemnity Payments

We construct another measure of SHELDUS local heat shocks from crop indemnity payments. The US Department of Agriculture maintains the crop indemnity payments data. We define an extreme heat indemnity event when the crop indemnity payments exceed the 90th percentile of 10-year historical monthly payments. Interestingly, the correlation of high crop indemnity events with recorded fatalities is tiny, with a Pearson coefficient of 3%; see Appendix Table A.1. This is for two reasons. First, while we construct both event types with SHELDUS, the events are retrieved and maintained by two different data providers. This leads to differences in methodology and coverage. And second, crop indemnity payments are more frequent among low-density rural areas, whereas fatalities and injuries are more frequent in urban areas.

Temperature

While SHELDUS extreme heat shocks capture the most devastating events, they do not necessarily capture prolonged periods of heat in regions that normally experience more chilly temperatures or have little agricultural activity. Therefore, we construct an additional county-level heat shock measure based on average temperatures. In particular, for each three-month window, we compute the average daily temperature. We then classify a heat shock by exceeding the 10-year historical record average temperature for the given county. And indeed, as Appendix Figure A.5 shows, the temperature shock measure adds counties that are otherwise overlooked by SHELDUS.

3.2 Impact on Google searches

In this subsection, we validate the ability of our proposed local heat shock measures to affect regional climate awareness. We use monthly state-level Google searches for "climate change" in this test.[2] This is a common technique employed by Choi et al. (2020) and others. Google search series measure relative interest in a topic, i.e., the fraction of all Google searches

[2]In principle, one could use the regional Google search interest as a direct measure of local climate change awareness in the following analysis. However, there are two problems with that. First, Google search interests are not available at the county level. While they are available at the MSA level, MSA's cover only a subset of the country and are not always directly mappable to counties. Second, at higher granularity, such as MSA level, Google often reports missing values for the Google search interests. That is due to insufficient data. Both these reasons make Google search interests an unreliable and imprecise measure of county-level climate change awareness.

in a region that search for a given term. Then, Google scales the relative interest cross-sectionally between 1 and 100. This means that, in each period, the region with the most relative searches 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.

We overcome the multiplicative scaling factor by testing the ability of our local climate shocks to affect the *logarithm of the Google searches*.[3] Since our local climate shocks are measured at county level, we aggregate them for each state. Our validation regression is then

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

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

	Log(Relative Google Search Volume)				
	(1)	(2)	(3)	(4)	(5)
Heat: Fatalities or Injuries	0.04*				
	(0.02)				
Heat: 50th pct. Fatalities or Injuries		0.06**			
		(0.03)			
Heat: Property Damages			0.11**		
			(0.05)		
Heat: 90th pct. Indemnities				0.07**	
				(0.03)	
Heat: Record Temperature					0.08**
					(0.04)
R^2	0.77	0.77	0.77	0.77	0.77
State FE	Y	Y	Y	Y	Y
Month FE	Y	Y	Y	Y	Y
N	5,506	5,506	5,506	5,506	5,506

Table 2: Regression of relative Google searches for climate change on various heat shock indicators. Standard errors are below the coefficients in parenthesis. Standard errors are clustered at the month and state level, and observations are weighted by the state's population size. Significance levels: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

[3]To visualize this, 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 $\log(\widetilde{G_{t,s}}) = \beta x_{t,s} + \epsilon_{t,s}$, we ensure that the error term captures the scaling factor: $\log(\widetilde{G_{t,s}}) = \log(G_{t,s}) - \log(\eta_t)$ so that $\log(G_{t,s}) = \beta x_{t,s} + u_{t,s}$, where $u_{t,s} = \epsilon_{t,s} + \log(\eta_t)$.

Table 2 reports the coefficients β_S from running this regression separately for various indicators of extreme local heat. We use an indicator for fatalities or injuries and a more severe version that requires having fatalities or injuries per capita over the 50th percentile to avoid focus on more populated states. Reported property damages from heat are rare in the Sheldus data and, therefore, we include an indicator for those too. Because indemnity payments do reach a lower level of attention and are frequent, we include an indicator for exceeding the 90th percentile in per capita terms. Finally, the last predictor is an indicator that any county in the state reached a new temperature record as defined in Subsection 3.1. All coefficients are at least somewhat significant. Intuitively, experiencing any fatalities or injuries from heat is associated with an increased interest in climate change by 4%. Reported property damages increase the relative Google searches by 11%, whereas record temperatures are associated with an increase by 8%. In summary, all heat measures can affect climate change awareness and are, therefore, suitable candidates to discipline the quantity hedge portfolio.

4 National Climate News Series

In contrast to our macro hedge targets in Section 8, there is no established standard for climate change risk. While the recent years brought forth a variety of climate news series that do well at capturing the various threats posed by climate change, there cannot be a single complete climate risk index. For example, physical risk can threaten coastal properties or inland forests, just as transitional risk can affect coal producers or car manufacturers. What risks should the index capture? How should the index weight the different risks? We believe there is no straightforward answer to this question, and future research in this direction is required. However, we also believe that all of the currently existing climate news series have their merits, and each captures a distinct mix of threats.

Therefore, we do not pick one climate news series to validate our portfolios but rather show that they can hedge multiple indices. Moreover, this approach allows us to determine the flaws and merits of our climate hedge portfolios. For example, climate hedge portfolios could be grouped by their ability to hedge physical or transitional risk. We list all included climate news series in the following subsections.

Faccini, Matin, and Skiadopoulos (2021)

We include four of Faccini et al. (2021)'s climate news indices: the international summits, global warming, natural disasters, and narrative index. We exclude their climate policy index

because it does not differentiate between good and bad news. That is especially relevant for the recent years in the US because former President Donald Trump notoriously reversed many climate protection laws. The international summits, global warming, and natural disasters indices measure news coverage of the respective topics, whereas Faccini et al. (2021) constructed the narrative index by manual inspection. Further, the international summits and narrative indices capture transitional risk, and the global warming and natural disasters indices capture physical risk.

Kelly (2021)

Kelly (2021) made three climate news series that reflect general, physical, and transitional risk, respectively. Each of the climate news series is the product of the number of relevant Wall Street Journal articles in a given month multiplied by their sentiment.

[This subsection will change.]

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 recognized work on constructing climate news series. While the first series builds upon WSJ news articles, the latter builds upon proprietary news aggregations from Crimson Hexagon. Note that only the latter index is capable of separating good from bad news. However, since both indices capture a mix of physical and transitional risks, we believe this is acceptable.

National Google searches

This climate news series is simply the national Google search interest in "climate change". The national Google climate news index does not differentiate between positive and negative news, and it could be associated with any type of climate risk. However, its main advantage is the ability to directly measure what we are most interested in: national climate change awareness. That is in contrast to all the other indices, which capture possible catalysts of climate change awareness.

National Temperature Deviations

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

5 Construction of Industry Climate Betas

Having identified county-level climate shocks, we can use them to measure mutual fund's portfolio reallocations after experiencing these events. The following subsections list our data sources, preprocessing steps, and the model we use to determine the industry-level climate betas.

5.1 Portfolio, stock, and industry data

We use the Thomson Reuters Mutual Fund Holdings s12 database to get a time series of portfolio holdings. We combine the holdings with fund characteristics obtained from CRSP. The mutual funds are linked from Thomson Reuters to CRSP on their Wharton Financial Institution Center Number (WFICN), which we got from WRDS MFlinks. Most funds report their holdings in three-month intervals. Therefore, in the following analysis, we will only analyze holding changes in three-month intervals.

We restrict asset holdings on share codes 10, 11, 12, and 18, and exchange codes 1, 2, and 3.[4] This ensures that our hedge portfolio spans North American common stocks. Since we want to identify deliberate fund manager asset reallocations in response to local climate shocks, we restrict our analysis to actively managed funds. Therefore, we only keep funds that have Investment Objective Code (IOC) 2, 3, 4, or missing, and have CRSP Objective Code 'Equity Domestic Non-Sector.'[5]

Our stock data is from CRSP, and we use end-of-month prices to determine the value of each holding. We then obtain stock GICS codes from Compustat by merging the assets on their CUSIP identifiers. The first four digits of the GICS industry code then determine the stock's industry classification.[6]

[4]Share codes 10, 11, 12, and 18 represent Ordinary Common Shares that are 'not further defined,' 'need not be further defined,' 'incorporated outside the US,' or 'REITs (Real Estate Investment Trusts).' Exchange codes 1, 2, and 3 represent the New York Stock Exchange, American Stock Exchange, and Nasdaq Stock Market, respectively.

[5]Investment Objective Codes 2, 3, and 4 represent 'Aggressive Growth,' 'Growth,' and 'Growth & Income,' respectively.

[6]The Global Industry Classification Standard (GICS) is developed by MSCI and S&P for use by the financial community. We decided to use this classification system because it ensures that within-industry return correlations tend to be high.

5.2 Mutual fund adviser locations

US mutual funds must publicly file with the SEC in regular intervals, and we use these filings to determine their fund advisers and their fund adviser locations. We share the belief that fund advisers are primarily responsible for making asset-allocation decisions (Chang 2019). Therefore, we parse the adviser locations from N-SAR filings until 2017 and parse them from N-CEN filings from 2018 onward, which replaced the N-SAR format.[7]

However, SEC filings cannot be matched directly with Thomson Reuters or CRSP mutual fund data. Therefore, we apply a fuzzy match algorithm to match SEC filings with mutual funds. We only continue with almost perfect name matches, successfully matching 84.1% of fund-holding observations.

For most funds, all advisers reside in the same county (67.6%). Whenever funds have multiple advisers that are not all located in the same county, we assign fund shock exposure as an average of fund adviser shock exposures. For example, if a fund has two advisers in county A and one adviser in county B, and county A is affected by a climate shock, we assume the fund is affected by a 2/3 climate shock. Appendix Figure A.1 shows the geographical distribution of the fund advisers for the subset of funds where all advisers reside in the same county.

5.3 Industry Climate Beta Regressions

This subsection describes how we measure mutual fund reactions to regional climate shocks and how we can use that information to determine industry climate betas from the view of professional investors. For that purpose, we define the active change in industry I holdings, as of period t, for mutual fund f, as

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

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

In other words, we are looking at the change in the share of value invested into industry I by fund f, using previous period prices, and weight the change by the market share of industry I. The idea behind the industry size scaling is that an increase in holdings of an

[7] [Refer to Github repository here.]

industry that is a tiny part of the economy is more meaningful than an increase in holdings of a relatively large, and common, industry.[8]

Since most funds report their holdings quarterly, we *always* measure fund composition changes in three-month intervals, i.e., there are three-months between period t and $t-1$, $\forall t$. We winsorize the $ActiveChanges^I_{f,t}$ at the 1% level to mitigate outliers posed by major strategy overhauls or mandate changes. Industry I's climate beta is then determined by regressing

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

where $S_{loc(f),t}$ is a local, time-varying climate shock and δ^I_t captures time fixed-effects. In other words, we regress active changes on the local, time-varying climate shock, controlling for time fixed-effects. Therefore, β^I captures the differential change in holdings of a given industry, when affected by a local climate shock, compared to funds that were not affected by a local climate shock.

From running the regression for each industry, we obtain industry-specific climate loadings β^I, which we can use for hedging climate change news.

6 Construction of Climate Hedge Portfolios

In this section, we first present the construction of the quantity hedge portfolio. Then, we continue with constructing alternative hedge portfolios that represent the ESG, narrative, or mimicking portfolio approach. These alternative approaches will be used as a benchmark for comparing our quantity methodology.

6.1 Quantity portfolio

For each month in our sample, we re-estimate β^I as defined in equation (3) using the previous five years of data on mutual fund portfolio compositions and climate shocks. We do this separately for each industry I and each of the local climate shocks $S_{loc(f),t}$ detailed in Section 3. Thereby, we create a series of five-year rolling industry climate betas $\{\widehat{\beta^I_{S,t}}\}$.

Then, we create industry portfolios whose weights will be determined by the industry

[8] Alternatively, one could continue the analysis with $PassiveChanges^I_{f,t}$, where the first fraction uses $P_{j,t}$ instead of $P_{j,t-1}$, i.e., current period holdings are valued at current period prices. This alternative approach takes price changes into account, and would be a more suitable model if we assume that funds can freely and constantly rebalance their portfolio. However, we noticed that there is not much difference between $ActiveChanges^I_{f,t}$ and $PassiveChanges^I_{f,t}$.

climate betas. In particular, for each industry I, we calculate the monthly industry return R_t^I as

$$R_t^I = \sum_{j \in I_t} Share_t^{j,I} R_t^j, \tag{4}$$

where $Share_t^{j,I}$ denotes firm j's share of industry I's market capitalization in month t and R_t^j denotes firm j's monthly return. So for each industry, we calculate a marketcap-weighted monthly return, based on the industry's constituents.

Finally, for each month t and for each type of local climate shock S, we then construct the quantity portfolio using these estimated industry betas and returns:

$$QP_{S,t} = \sum_I \widehat{\beta_{S,t}^I}(R_t^I - R_t^f), \tag{5}$$

where $\widehat{\beta_{S,t}^I}$ denotes the estimated beta for industry I for month t using the previous five years of data and R_t^f denotes the risk-free rate.

6.2 ESG approach: Sustainalytics portfolio

Our first benchmark portfolio takes positions in all US-listed stocks covered by Sustainalytics ESG scores. In particular, the portfolio's position in each stock is the stock's ESG score percentile minus 50 in each period. In other words, the portfolio takes a long position of 50 into the company with the highest ESG score and a short position of -50 into the company with the lowest score in each month. Stocks with exactly the median ESG score are neither invested nor shorted. This portfolio follows the ESG approach.

6.3 Narrative approach: ETFs

For the narrative approach, we use the returns on ETFs to construct climate hedge portfolios. The first index, Invesco Global Clean Energy ETF (Ticker: PBD), invests in firms focused on the development of cleaner energy and conservation. On the other hand, the second index, Energy Select Sector SPDR Fund (Ticker: XLE), tracks a market-cap-weighted index of US energy companies that are constituents of the S&P 500 index. According to the narrative, realizations of climate change news should increase PBD's returns and decrease XLE's returns.

Moreover, we also construct a stranded asset portfolio as in Jung et al. (2021) based on the XLE, VanEck Vectors Coal (KOL) and the SDPR S&P 500 (SPY) ETF:

$$Stranded\ Asset\ Portfolio = 0.3XLE + 0.7KOL - SPY. \tag{6}$$

6.4 Mimicking portfolio approach

The strategy of the factor mimicking portfolios is to approximate the climate change innovations with a combination of asset or portfolio returns. In the following paragraphs, we describe how we construct the factor mimicking portfolios with a projection on the Fama-French factors, ETFs, or industry portfolios.

Fama-French factor projection

For the Fama-French factor projection, we use the market, size, and value factors. Accordingly, we estimate the weights by regressing the climate change news innovations on the excess return of the market, and the HML and SMB factors, as follows:

$$CC_{c,t} = \beta^{M}_{c,t}(R^{M}_{t} - R^{f}_{t}) + \beta^{SMB}_{c,t} SMB_t + \beta^{HML}_{c,t} HML_t, \tag{7}$$

where $CC_{c,t}$ denotes the residual from an AR(1) model fitted on the climate change news series of type c in month t. The excess market return, SMB and HML factors are denoted by $(R^{M}_{t} - R^{f}_{t})$, SMB_t, and HML_t, respectively.

We estimate new betas each month, using a rolling five-year window of data. Using these estimated betas, we then construct the Fama-French factor portfolios for each climate news series as:

$$FF_{c,t} = \widehat{\beta^{M}_{c,t}}(R^{M}_{t} - R^{f}_{t}) + \widehat{\beta^{SMB}_{c,t}} SMB_t + \widehat{\beta^{HML}_{c,t}} HML_t, \tag{8}$$

where $\widehat{\beta^{M}_{c,t}}$, $\widehat{\beta^{SMB}_{c,t}}$, and $\widehat{\beta^{HML}_{c,t}}$ denote the estimated factor betas for month t and climate news series c.

ETF projection

To estimate the weights to construct the ETF projection, we regress the climate change news series innovations on the returns of PBD, XLE, and the three Fama-French factors, as follows:

$$CC_{c,t} = \beta^{PBD}_{c,t} R^{PBD}_{t} + \beta^{XLE}_{c,t} R^{XLE}_{t} + Z_{c,t}, \tag{9}$$

where $CC_{c,t}$ denotes the AR(1) climate news residuals, R^{PBD}_{t} and R^{XLE}_{t} are the returns of the PBD and XLE ETF, and $Z_{c,t} = \beta^{M}_{c,t}(R^{M}_{t} - R^{f}_{t}) + \beta^{SMB}_{c,t} SMB_t + \beta^{HML}_{c,t} HML_t$, i.e., $Z_{c,t}$ is shorthand notation for the three Fama-French factors.

Again, using a five-year rolling window regression, we construct the ETF projection portfolio with the predicted factor loadings as:

$$ETF_{c,t} = \widehat{\beta^{PBD}_{c,t}} R^{PBD}_{t} + \widehat{\beta^{XLE}_{c,t}} R^{XLE}_{t} + \widehat{Z_{c,t}}. \tag{10}$$

Industry (lasso) projection

Our final benchmark portfolios are based on a five-year rolling all-industry regression. To construct the weights for the portfolios, we use the monthly industry returns R_t^I as in equation (4) together with the three Fama-French factors. We run lasso regressions of the climate change news innovations on the industry returns and the Fama-French factors:

$$\begin{aligned} CC_{c,t} &= \sum_{i \in I} \beta_{c,t}^i (R_t^I - R_t^f) + Z_{c,t}, \\ \text{s.t.} \quad &\left(\sum_{i \in I} |\beta_{c,t}^i| \right) + \left|\beta_{c,t}^M\right| + |\beta_{c,t}^{SMB}| + |\beta_{c,t}^{HML}| \leq t, \end{aligned} \tag{11}$$

where t is a specified free parameter that determines the degree of regularization. We tune the hyperparameter t by cross-validation. In particular, t is chosen such that the cross-validated MSE generalization error is at most one standard deviation above the t parameter. This is a common rule that balances in-sample fit with the risk of overfitting. Using the estimated weights, we then construct the lasso portfolios for each climate news series as:

$$Lasso_{c,t} = \sum_{i \in I} \widehat{\beta_{c,t}^i} (R_t^i - R_t^f) + \widehat{Z_{c,t}}. \tag{12}$$

Alternatively, we also construct a portfolio that uses all industries in a standard regression framework. However, this introduces the risk of overfitting and, therefore, the lasso variation is our preferred portfolio. For completeness, however, we include both.

7 Hedging Climate Change News

In this section, we validate the effectiveness of our climate change hedge portfolios. In particular, we contrast the hedging ability of the quantity portfolios with various alternative hedging portfolios built using the ESG, narrative, or mimicking portfolio approach (as described in Sections 5 and 6). We build the hedging portfolios in our training sample (2010-2014), and evaluate them out of sample over the period 2015-2019. We focus on the post-2010 data because, before then, very little attention was paid in the investment world to climate risks, so that we wouldn't expect information in prices and quantities to be useful in building hedging portfolios.

7.1 Validation of the climate change hedge portfolios

In this section we evaluate the hedging performance of the different portfolios out of sample. As a criterion to compare the various approaches, we use the out-of-sample correlation

between the hedging portfolio return and the innovation (from an AR(1) model) in the climate news indices. As discussed in the previous sections, in this paper we consider several different "target" climate news measures. Given that climate change is a complex and multi-faceted phenomenon, it is possible that different portfolios are better suited to hedge different types of climate news.

Table 3 reports the out-of-sample correlation of the hedging portfolios and the climate shocks, at the monthly frequnecy.[9] Each row in the table represents a hedge portfolio, whereas each column corresponds to a different climate news series.

Note that some of the hedge portfolios (corresponding to different rows) are tailored for each climate news series, that is, the portfolio weights corresponding to a particular news series (column) are different across columns. For example, the first row builds a hedging portfolio using the Fama-French 3 factors. The portfolio weights on the three long-short portfolios are selected using the training data separately for each news series. This gives additional flexibility to the traditional hedging portfolios.

On the other hand, for other methodologies, the portfolio is built without using any information about the climate news series to hedge. In this case, the portfolios are not tailored to specific news, and the portfolio weights are the same across columns. This is the case for all portfolios that are built using "prior" information, like those based on the ESG scores and our quantity portfolios.

Finally, all climate news series are coded such that high numbers are tantamount to negative climate news. Therefore, positive correlation coefficients show successful hedges.

The same information, in a different format, is displayed in Figure 1. The figure shows a dot plot, where each point is the correlation coefficient with one of the climate news series. We color-coded different groups to facilitate interpretability. In particular, the groups show the Faccini et al. (2021), Kelly (2021), Engle et al. (2020) and Ardia et al. (2020), and Google searches and national temperature climate news series in different colors. The red rhombus shows the average among all correlations. In other words, the dot plot allows a quick and neat comparison of the effectiveness of the various hedge portfolios, while still maintaining all information contained in Table Table 3.

We begin by studying the hedging ability of traditional hedging portfolios for climate risks. The first 4 rows of the Table report the results for the portfolios built via mimicking

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

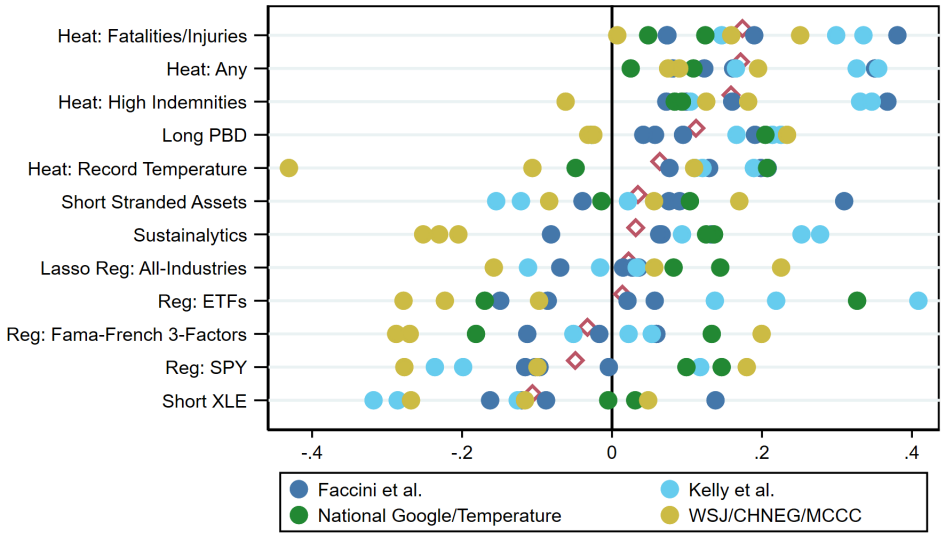


Figure 1: Dot plot of monthly return correlations for various climate hedge portfolios with various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Different colors represent different groups to which the climate news series belong. In particular, we group Faccini et al. (2021), Kelly (2021), Engle et al. (2020) and Ardia et al. (2020), and Google searches and national temperature climate news series. For more details, refer to Table 3.

	Google	Temp.	Faccini, Matin, Skiadopoulos				Kelly [et al.]			Engle et al.		Ardia et al.
	National	National	IntSummit	GlobWarm	NatDis	Narrative	General	Transitional	Physical	WSJ	CHNEG	MCCC
Reg: FamaFrench 3-Factors	**.13**	−.18	−.02	**.06**	**.06**	−.11	**.05**	−.05	**.02**	**.20**	−.29	−.27
Reg: ETF + FamaFrench	**.33**	−.17	−.09	−.15	**.06**	**.02**	**.41**	**.22**	**.14**	−.28	−.22	−.10
Reg: All-Industries	**.15**	**.13**	−.12	**.08**	−.12	**.03**	−.04	−.08	**.09**	**.09**	**.11**	**.08**
Lasso Reg: All-Industries	**.14**	**.08**	**.03**	**.01**	−.07	**.04**	**.03**	−.11	−.02	**.06**	**.23**	−.16
Long PBD	**.20**	−.03	**.06**	**.09**	**.19**	**.04**	**.21**	**.17**	**.23**	−.02	−.03	**.23**
Short XLE	−.01	**.03**	−.09	−.16	−.12	**.14**	−.32	−.13	−.29	−.12	**.05**	−.27
Short Stranded Assets	**.10**	−.01	−.04	**.08**	**.31**	**.09**	−.12	**.02**	−.15	**.06**	**.17**	−.08
Sustainalytics	**.13**	**.14**	**.13**	−.08	**.06**	**.07**	**.25**	**.28**	**.09**	−.25	−.23	−.20
Heat: Any	**.02**	**.11**	**.35**	**.16**	**.12**	**.08**	**.33**	**.17**	**.36**	**.07**	**.09**	**.19**
Heat: Fatalities/Injuries	**.05**	**.12**	**.38**	**.19**	**.07**	**.07**	**.30**	**.15**	**.34**	**.01**	**.16**	**.25**
Heat: High Indemnities	**.09**	**.08**	**.37**	**.16**	**.10**	**.07**	**.33**	**.10**	**.35**	**.13**	−.06	**.18**
Heat: Record Temperature	**.21**	−.05	**.21**	**.13**	**.20**	**.08**	**.12**	**.19**	**.11**	−.43	**.11**	−.11

Table 3: Monthly return correlations for various climate hedge portfolios 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. To increase legibility, we highlighted positive correlation coefficients in bold. Also, all climate news series were coded such that high numbers are tantamount to negative climate news. Therefore, positive correlation coefficients, i.e., the numbers in bold, show successful hedges. While the long, short, Sustainalytics, and quantity portfolios stay constant along the rows, the regression and lasso regression portfolios show in each cell the portfolio that was specifically trained on the respective climate news series.

portfolio regressions. The performance of these portfolios varies substantially across climate news series. They all produce reasonable out-of-sample correlations with the Google search innovations, but also display substantial variation with the remaining news series. For example, the portfolio built using the Fama-French 3 factors has relatively large correlation with the WSJ index from Engle et al. (2020) (.2) in addition to a .13 correlation with the Google index. But it also displays relatively large *negative* correlation with the CHNEG index from Engle et al. (2020) (-0.29), similarly negative correlations with national temperatures (-.18) and with MCCC from Ardia et al. (2020) (-0.27). All the other correlations are close to zero. Overall, the performance of the mimicking-portfolio hedging approach is mixed. Figure 1 shows graphically the poor performance of this approach.

Next (rows 5-8 of the table), we can study the performance of portfolios built using climate-related characteristics. The main advantage of these portfolios is that they do not require estimating the portfolio weights from historical data, as the direction of the trade, and the relative weights, are determined based on prior information. For example, the first portfolio of the group features a long position in the ETF PBD – the long position is motivated by economic reasoning, in that PBD is a clean energy fund where the natural hedging position is long (with respect to physical risks but most importantly transition risks). The second row is instead 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. Again, the performance of these portfolios is mixed, with the worst results given by the short XLE position, and the best results given by the long PBD position. The Sustainalytics portfolio is strong in hedging the Kelly (2021) climate news series but fails in hedging the indices of Engle et al. (2020) and Ardia et al. (2020).

These results highlight a few important points about the construction of hedging portfolios for climate risks. First, hedging climate risk is hard. We have little historical data available, which makes the mimicking portfolio approach particularly noisy in practice. Second, using climate characteristics to build a hedging portfolio can give encouraging results (PBD is able to hedge out of sample all but 3 news series), especially because it does not require estimating portfolio weights using historical data. However, as the table shows, there is an inherent difficulty in choosing the "right" climate characteristics: beyond PBD, the other 3 portfolios in this group do not perform as well or as consistently. Finally, the results vary significantly across climate news series (the targets of the hedge). Many of the hedging portfolios considered so far perform well in hedging the Google searches and the narrative

index, whereas for most of the other climate news targets, the performance is mixed and inconsistent, with some approaches better hedging some targets and not others.

Having these two approaches as baseline, we can now look at the performance of the quantity hedge portfolios. These portfolios tend to produce relatively high out-of-sample correlations for a large variety of measures. Figure 1 summarizes these results: the four heat indexes appear at the very top of the figure, with the almost entirety of correlations in the positive domain. In fact, both "Heat: Fatalities/Injuries" and "Heat: Any" positively correlate with all climate news innovations. "Heat: High Indemnities" could hedge all but the CHNEG series, whereas "Heat: Record Temperature" only failed for the national temperature deviations, WSJ, and MCCC. Given the low correlation among most of the temperature shocks (see Appendix Figure A.1), the high consistency in the results is notable. All heat quantity portfolios provide excellent hedges for Faccini et al. (2021)'s international summits and global warming indices, as well as Kelly (2021)'s general and physical risk indices. This suggests that mutual fund adviser's climate change awareness following an extreme heat event increases both in terms of physical and transitional risk.

Why does the quantity approach outperform the traditional methodologies? In our view, the answer is compatible with a variety of explanations. First, as previously stated, mimicking portfolios' main disadvantage is the backward-looking nature, combined with the fact that only recently investors have started incorporating climate risks into prices. ESG scores also take time to adjust to changes, but further, they are oblivious to regulatory developments. For example, when the government switches its climate strategy from carbon taxes to expanding recycling regulations, this can have major implications on stock valuations but is unlikely to be reflected in ESG scores. Narrative approaches are similarly static and they could suffer from oversimplification. For example, hedging climate change could require a more complex portfolio structure than simply investing into green stocks.

In contrast, we believe that the quantity approach uses a forward-looking information source. In general, investors observe climate and regulatory events in real-time and swiftly anticipate their future consequences. That is the missing feature of the traditional approaches. Moreover, the quantity approach does not rely on a single information source but instead uses crowd knowledge, in that it aggregates the responses of many different fund managers.

7.2 Robustness

During the construction of the quantity hedge portfolio, we had to take a stance on a variety of options. What funds should be included? How do we winsorize industry changes? What to do with multiple advisers in different counties? In this subsection, we explore the impact of those decisions on the performance of the "Heat: Any" quantity hedge portfolio. We chose this portfolio as a baseline because it was the 2nd best performer while achieving a relatively low variance in the correlation with the various climate news series.

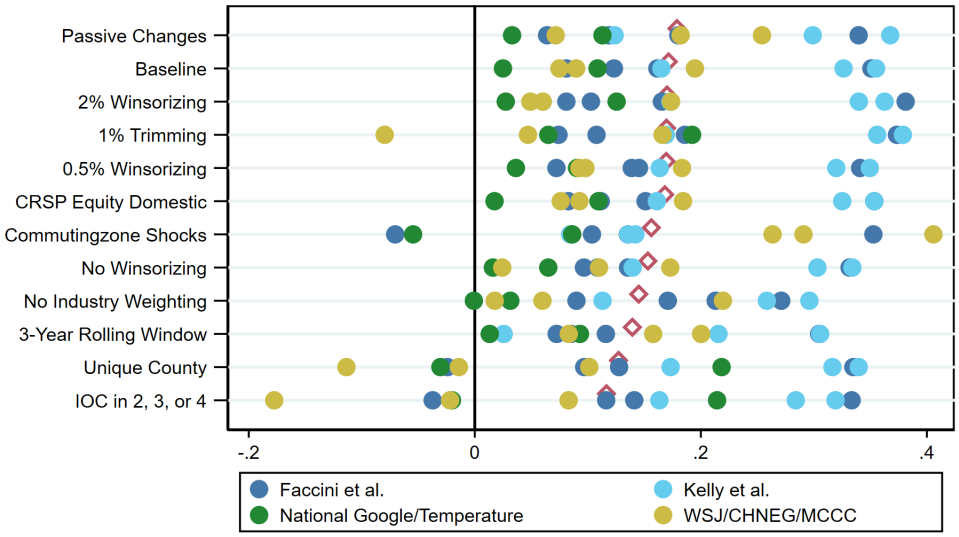


Figure 2: Dot plot of monthly return correlations for various climate hedge portfolios with various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Different colors represent different groups to which the climate news series belong. In particular, we group Faccini et al. (2021), Kelly (2021), Engle et al. (2020) and Ardia et al. (2020), and Google searches and national temperature climate news series. For more details, refer to Table 3. The climate hedge portfolios are variations of "Heat: Any" and, therefore, show numerous robustness tests.

Figure 2 shows the correlation of the robustness portfolios with the innovations of all climate indices. We considered the following variations: define industry exposure changes in terms of current prices (*PassiveChanges*); trim industry changes at 1% or winsorize them at 0.5%, 2%, or not at all; define the relevant universe of funds with only the CRSP objective code or only the Thomson Reuters IOC; do not weight industry changes by their relative market size; aggregate climate shocks from counties to commuting zones; and lastly, only keep funds where all advisers reside in the same county.

At first glance, Figure 2 reveals that our model is not particularly sensitive to any of the choices we have taken. However, it is notable that the CRSP objective code is more suitable in picking relevant mutual funds than the Thomson Reuters classification. Moreover, it

is surprising that restricting the subset of funds to those where all advisers reside in the same county often decreases the hedging ability. This could be due to the lower sample size having less information but also due to introduced selection bias. Omitting the weighting by industry size does not largely affect the results, as does using different variations of winsorizing or trimming. The measurement of industry changes with *PassiveChanges* is just as good as using *ActiveChanges*. And finally, aggregating the shocks to commuting zones did improve the ability to hedge the news series of Engle et al. (2020) and Ardia et al. (2020) but worsened the performance when evaluated against Kelly (2021) news series. Therefore, it is not clear whether climate shocks should be measured at the more precise county-level or the, potentially more relevant, commuting zone level. In summary, the quantity hedge portfolio is robust to several variations in its construction.

7.3 Factor neutral hedge portfolios

Ultimately, investors need to analyze their portfolio in its entirety. Some investors prefer to achieve higher loadings on the market or size factor, while other strategies might involve market neutrality. In this subsection, we investigate how much the quantity hedge portfolio loads on the three Fama-French market, size, and value factors, as well as how the hedge portfolio performs when we orthogonalize it to these factors.

To identify the factor loadings of the quantity portfolios, we regress the portfolio returns on the returns of the market, size, and value factor:

$$QP_{S,t} = \beta_{c,t}^{M}(R_t^M - R_t^f) + \beta_{c,t}^{SMB} SMB_t + \beta_{c,t}^{HML} HML_t + \epsilon_{t,S}. \tag{13}$$

Table 4 shows the regression results. All portfolios have a relatively weak, albeit significant, loading on the market factor. While the size factor appears to be irrelevant for all portfolios, the "Heat: Any" and "Heat: Fatalities/Injuries" portfolios have a significant loading of 0.15 on the value factor. No portfolio achieved a significant alpha.

Overall, the relatively low factor loadings show that hedging climate change risk requires a more complex approach than selecting factor exposure. Moreover, this result is in line with the poor performance of the Fama-French mimicking portfolio. However, in a sense, this is an encouraging result too. The low factor loadings show that investors can incorporate climate risk hedges in existing portfolios without huge changes in other market risk exposures.

To test this claim, we analyzed the hedging ability of the various portfolios when orthogonalized to the three factors. In particular, we used the residuals of equation (13) and correlated them with the AR(1) innovations of each of the climate news series. Figure 3

	Heat: Quantity Hedge Portfolio			
	Any	Fatalities/Injuries	High Indemnities	Record Temperature
$R^M - R^f$	0.06*	0.06**	0.13***	0.18***
	(0.03)	(0.03)	(0.04)	(0.06)
SMB	0.01	0.02	−0.03	0.00
	(0.03)	(0.03)	(0.05)	(0.06)
HML	0.15***	0.15***	0.04	−0.01
	(0.03)	(0.03)	(0.06)	(0.05)
Constant	−0.02	0.00	−0.05	0.08
	(0.09)	(0.09)	(0.15)	(0.14)
R^2	0.35	0.38	0.18	0.30
N	60	60	60	60

Table 4: Regression of the heat quantity hedge portfolios on the market, size, and value factors as in equation (13). Standard errors are below the coefficients in parenthesis. Standard errors are heteroskedasticity robust. Significance levels: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

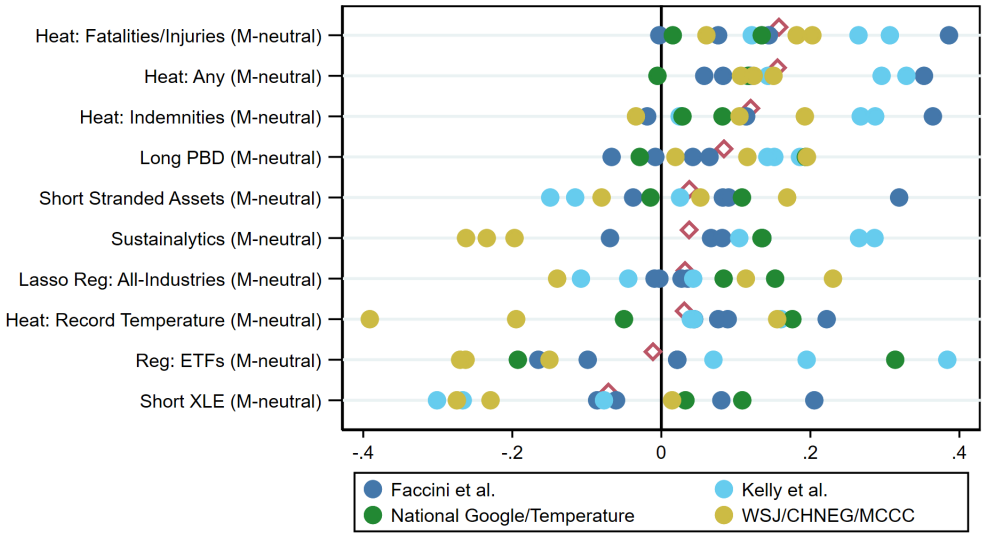


Figure 3: Dot plot of monthly return correlations for various climate hedge portfolios with various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Different colors represent different groups to which the climate news series belong. In particular, we group Faccini et al. (2021), Kelly (2021), Engle et al. (2020) and Ardia et al. (2020), and Google searches and national temperature climate news series. For more details, refer to Table 3. We orthogonalized all hedge portfolios with respect to the Fama-French market factor portfolio.

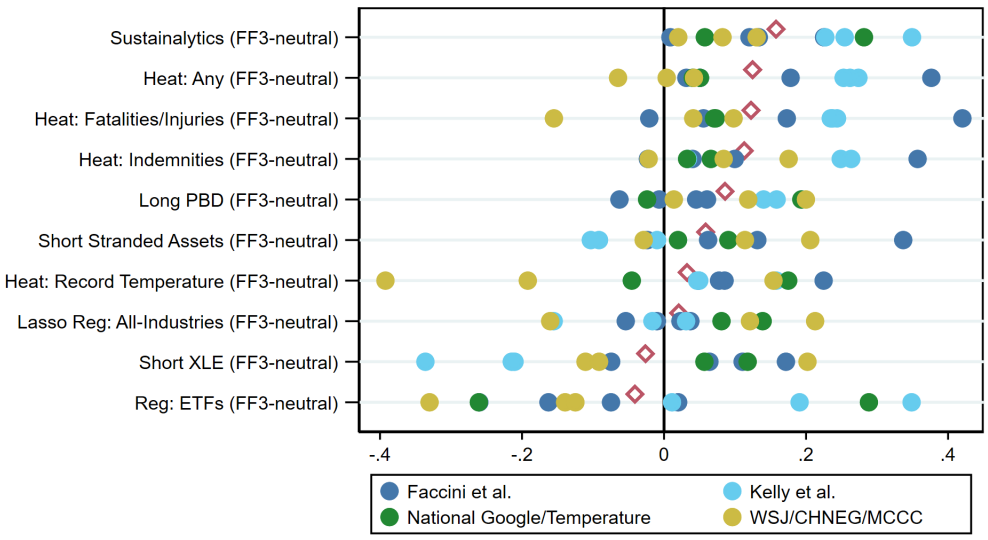


Figure 4: Dot plot of monthly return correlations for various climate hedge portfolios with various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Different colors represent different groups to which the climate news series belong. In particular, we group Faccini et al. (2021), Kelly (2021), Engle et al. (2020) and Ardia et al. (2020), and Google searches and national temperature climate news series. For more details, refer to Table 3. We orthogonalized all hedge portfolios with respect to the Fama-French market, HML, and SMB factor portfolio.

shows the results when only orthogonalized to the market factor, and 4 show the results when orthogonalized to all three. As expected, the quantity hedge portfolios continue to perform well in both cases. However, in general, the correlations are somewhat lower, especially when removing all three factor loadings. That is not surprising; imposing additional restrictions on the portfolios should make the hedging problem harder. Notably, though, the Sustainalytics portfolio substantially improved the average correlation after removing market, size, and value exposure. In fact, it is the top performer in Figure 4.

8 Hedging Macro Factors

We argued that hedging climate change news is inherently difficult in Section 4. Although our constructed quantity hedge portfolios show promising results in Section 7, we use this section to further show the effectiveness of our approach by applying the methodology to hedge national macro news. Moreover, by doing so, we demonstrate the versatility of our quantity-based approach. We use the same approach as outlined in Sections 5 and 6.

8.1 Macro shocks

We consider two macro series that we are trying to hedge: unemployment and housing prices. Both are suitable choices for our methodology because shocks in either series are highly noticeable at the national and the local level. We obtain national and county-level monthly unemployment figures from the US Bureau of Labor Statistics. Our housing index is the Zillow Home Value Index (ZHVI).

While there are many ways of defining unemployment and housing shocks, we keep our classification deliberately simple to focus on the quantity-information model. For the unemployment series, we deseasonalize it and define shocks as the innovations of a 3-month change AR(1) process:

$$Unemp_{t,c} = \theta Unemp_{t-1,c} + \delta_m + \epsilon_{t,c}, \tag{14}$$

where $Unemp_{t,c}$ is the current employment level and $Unemp_{t-1,c}$ is the three-month lagged unemployment level for county c. Month fixed effects are captured by δ_m.

Local housing shocks are defined as the innovations of a three-month AR(1) process of the log series:[10]

$$Log(ZHPI_{t,c}) = \theta Log(ZHPI_{t-1,c}) + \epsilon_{t,c}. \tag{15}$$

Intuitively, by using these shocks, our methodology captures the response of mutual funds to unexpected rises in local unemployment or housing prices. We obtain the innovations by applying the regression from 2000 to 2019, and, to align with the climate sample, we validate the performance from 2015 to 2019 inclusive.

8.2 Validation of the macro hedge portfolios

Applying Regression (3) with the local macro shocks instead of local climate change shocks gives us industry-specific macro betas. We then construct the unemployment and housing price quantity hedging portfolios as in Equation (5). Moreover, for comparison, we construct mimicking portfolios as in Section 6.

Table 5 shows the correlation of our constructed macro hedge portfolios with AR(1) innovations of the national unemployment series and the logarithm of the housing price index. Both the Fama-French factor projection and the all-industry lasso regression portfolio achieve favorable correlations with shocks in the national indices. In other words, their returns are

[10] We use the logarithm because the index is growing exponentially. Moreover, we do not apply deseasonalizing because the index is already deseasonalized.

high when either unemployment or house prices grow unexpectedly. In comparison, the quantity portfolio achieves the second-best result for the ZHPI index and the best result for the unemployment index, highlighting the great efficacy of the quantity-based approach. Moreover, the quantity portfolios appear to be more reliable than the mimicking approach: while the all-industry lasso regression resulted in the best hedge against housing prices, it failed fully when applied to unemployment shocks.

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

Table 5: Monthly return correlations for various unemployment and housing price hedge portfolios 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. To increase legibility, we highlighted positive correlation coefficients in bold. While the quantity portfolios stay constant along the rows, the regression and lasso regression portfolios show in each cell the portfolio that was specifically trained on the respective macro series.

As previously argued, we believe one of the main advantages of the quantity-information approach is its forward-looking nature. To further test this claim, we repeat the analysis but train all models, i.e., mimicking and quantity portfolios, in a three-year rolling window (instead of five). A three-year window allows the strategies to adjust faster to structural changes, which could be due to regulatory or technological developments. Table 6 reports the results. Notably, all but the unemployment quantity hedge portfolio perform worse when disciplined only with the previous three years, making the unemployment quantity hedge portfolio the clear leader. However, the quantity ZHPI portfolio's outcome is poor. Still, the correlation is comparable with the other hedge strategies and, in fact, the second-best among the considered options. Therefore, both tables show that the quantity-information approach is highly competitive and often even better than the mimicking portfolio strategy.

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

Table 6: A variation of Table 5 where all hedge strategies are based on a three-year rolling window (instead of five).

9 Conclusion and Directions for Future Research

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

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

References

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

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

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

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

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

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

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

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

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

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

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

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

Kelly, B. (2021), 'Tbd'.

Konisky, D. M., Hughes, L. & Kaylor, C. H. (2016), 'Extreme weather events and climate change concern', *Climatic Change* **134**(4), 533–547.

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

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

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

A Appendix

A.1 Tables

	Any	Fatalities/Injuries	Indemnities	Record Temperature
Heat: Any	1.00			
Heat: Fatalities/Injuries	0.87	1.00		
Heat: Indemnities	0.03	0.03	1.00	
Heat: Record Temperature	0.04	0.04	0.14	1.00

Table A.1: Monthly Pearson correlation among the local climate shocks across the US.

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

Table A.2: Industry climate beta coefficients as in equation (3). The coefficients are sorted by "Heat: Any" and are based on data from 2015 to 2019 inclusive. Therefore, these are the most current industry climate betas in our sample.

	Any	Fatalities/Injuries	Indemnities	Record Temperature
Heat: Any	1.00			
Heat: Fatalities/Injuries	0.99	1.00		
Heat: Indemnities	0.56	0.56	1.00	
Heat: Record Temperature	0.18	0.20	0.25	1.00

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

	Any	Fatalities/Injuries	Indemnities	Record Temperature
Heat: Any	1.00			
Heat: Fatalities/Injuries	0.99	1.00		
Heat: Indemnities	0.44	0.42	1.00	
Heat: Record Temperature	0.30	0.29	0.28	1.00

Table A.4: Spearman rank correlation among the industry climate beta coefficients determined from various local heat shocks as in equation (3). The coefficients are based on data from 2015 to 2019 inclusive. Therefore, these are the most current industry climate betas in our sample.

	Any	Fatalities/Injuries	Indemnities	Record Temperature
Heat: Any	1.00			
Heat: Fatalities/Injuries	0.98	1.00		
Heat: Indemnities	0.61	0.61	1.00	
Heat: Record Temperature	0.48	0.52	0.36	1.00

Table A.5: Monthly Pearson portfolio return correlation.

	Any	Fatalities/Injuries	Indemnities	Record Temperature
Heat: Any	1.00			
Heat: Fatalities/Injuries	0.97	1.00		
Heat: Indemnities	0.62	0.61	1.00	
Heat: Record Temperature	0.51	0.55	0.18	1.00

Table A.6: Monthly Pearson portfolio return correlation. All portfolios are orthogonalized to the Fama-French market, HML, and SMB factor.

A.2 Figures

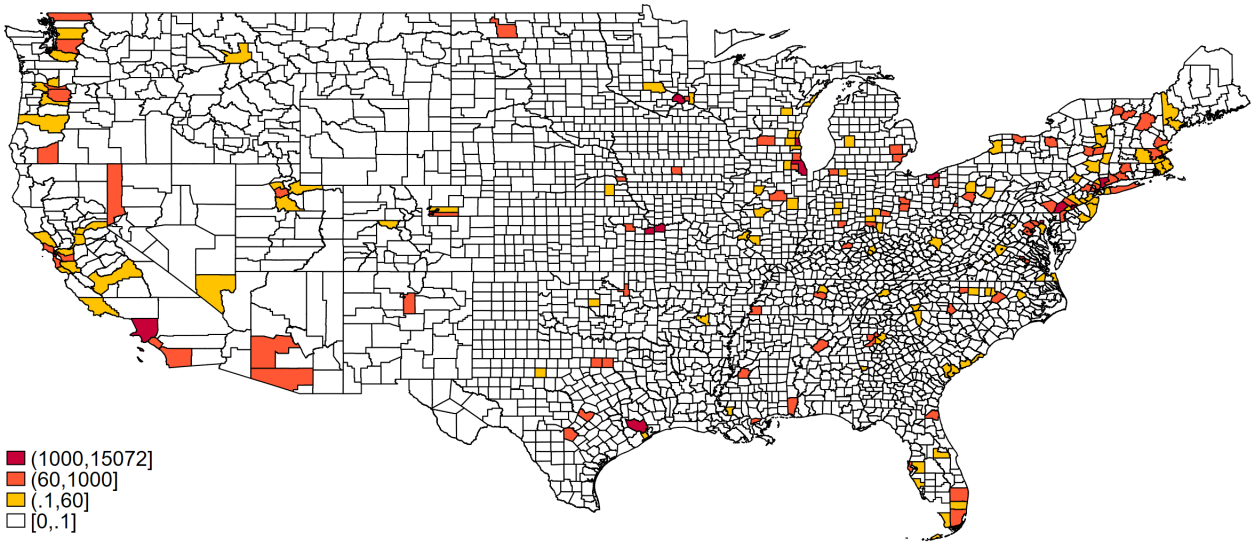


Figure A.1: Distribution of mutual fund adviser locations in our final sample. This graph only shows fund-quarter observations where all advisers reside in the same location.

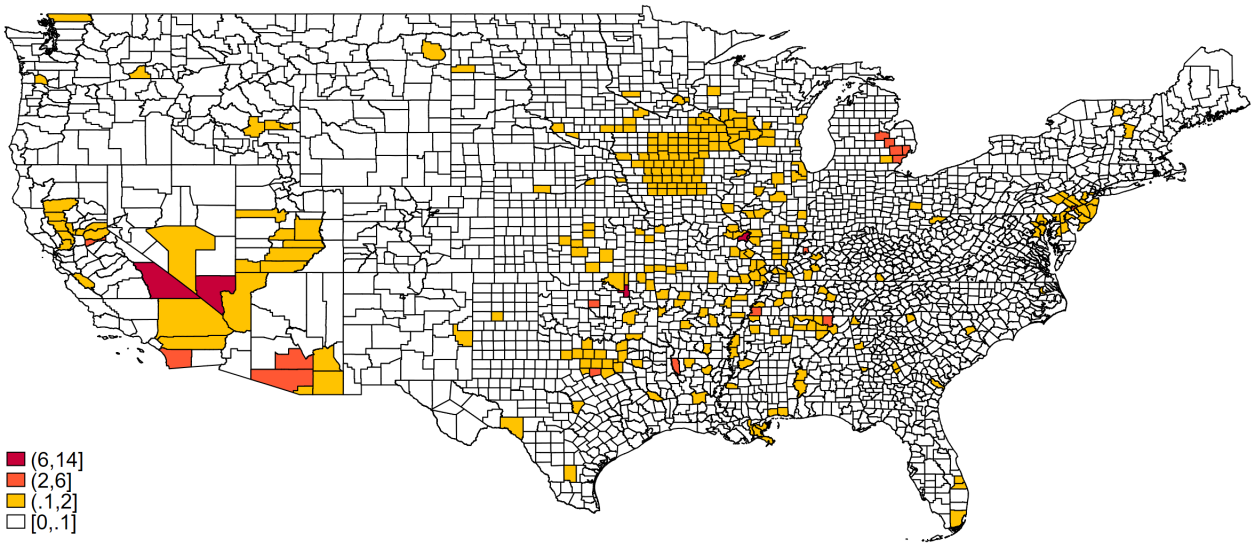


Figure A.2: Distribution of the "Heat: Any" extreme climate shock from 2010 to 2019.

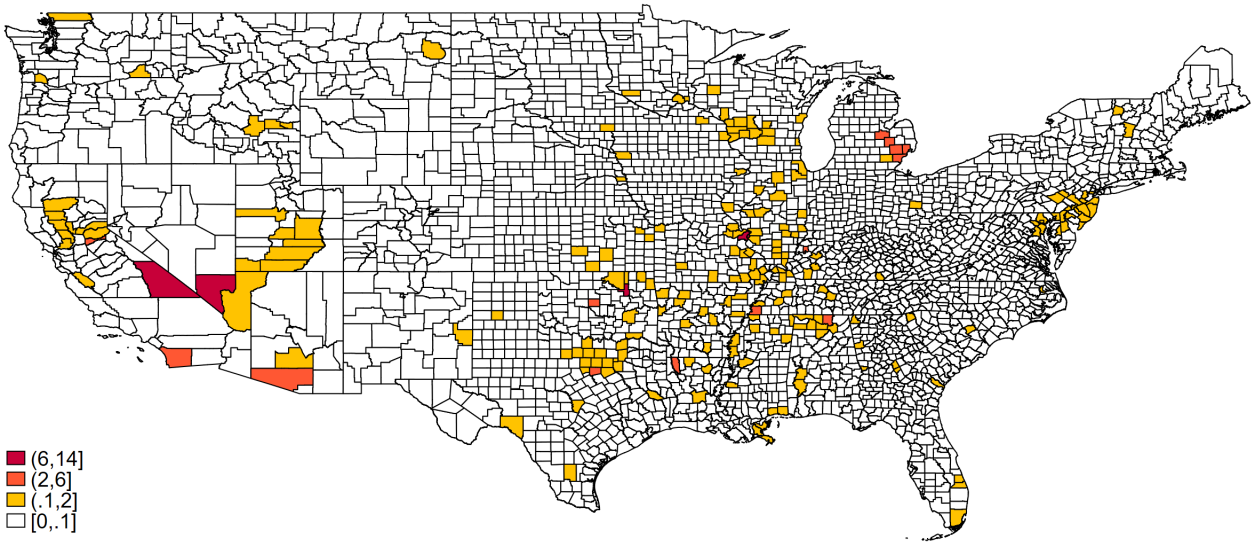


Figure A.3: Distribution of the "Heat: Fatalities or Injuries" extreme climate shock from 2010 to 2019.

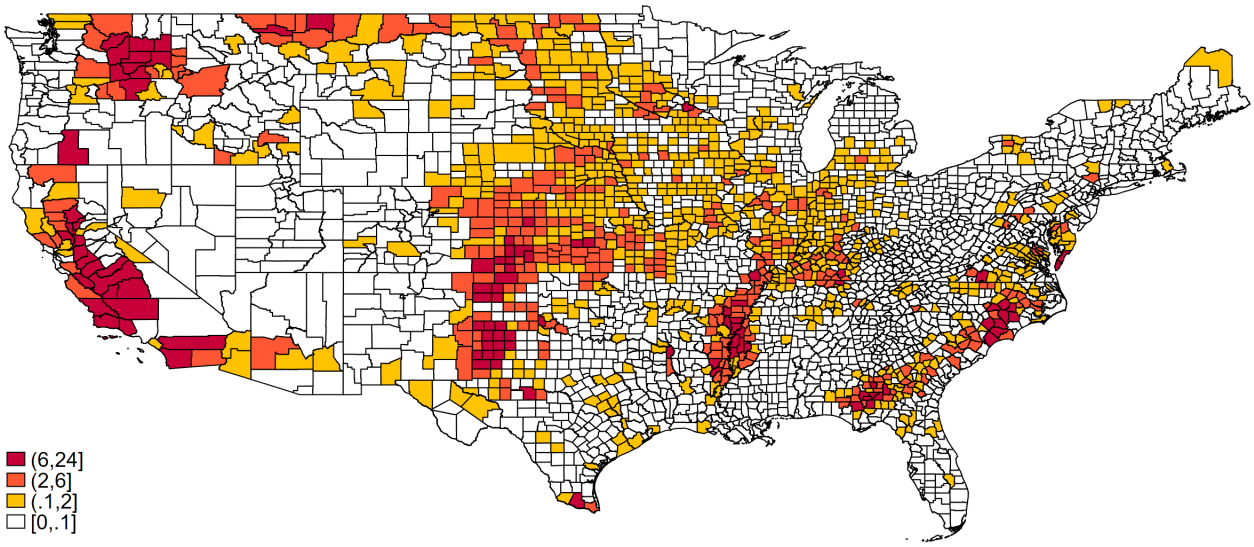


Figure A.4: Distribution of the "Heat: High Indemnities" extreme climate shock from 2010 to 2019.

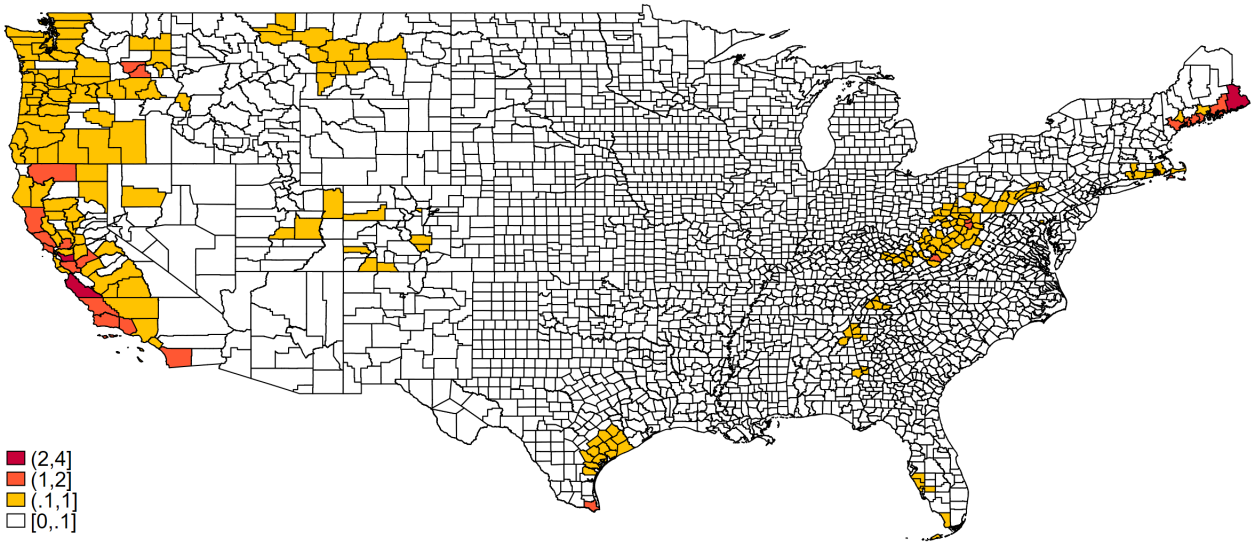


Figure A.5: Distribution of the "Heat: Record Temperature" extreme climate shock from 2010 to 2019.

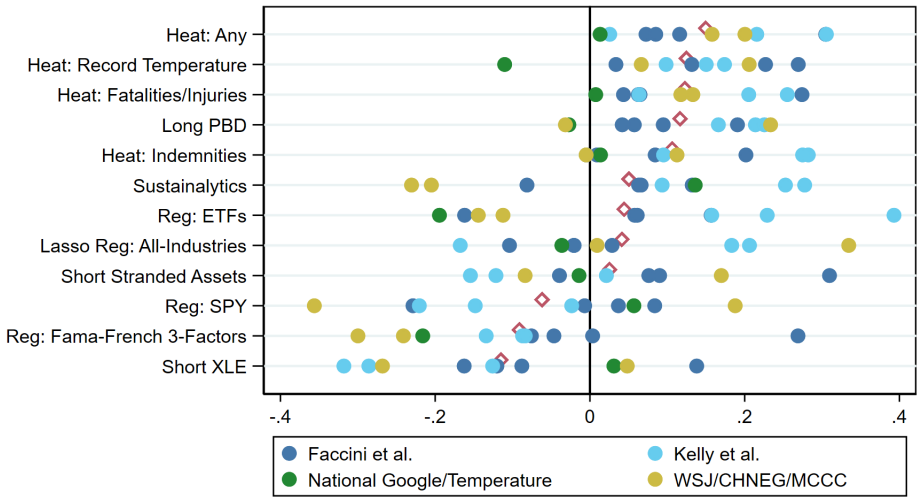


Figure A.6: Dot plot of monthly return correlations using 3-year rolling windows; see Table 3.