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 hedge portfolios for climate risks. Our quantity-based approach exploits information on mutual funds' trading responses to idiosyncratic changes in fund managers' climate beliefs. We identify these belief shocks based on (i) managers experiencing local extreme heat events that shift climate beliefs and (ii) changes in the way funds' shareholder disclosures discuss climate risks. We show that a portfolio that is long industries that investors tend to buy after experiencing negative idiosyncratic climate belief shocks, and short industries that investors tend to sell, appreciates in value in periods with negative aggregate climate news shocks.

*This version: October 2, 2024. We thank Daron Acemoglu, Michael Barnett, Richard Berner, John Campbell, Lauren Cohen, Robert Engle, Nicola Gennaioli, Itay Goldstein, Wei Jiang, Jonathan Lewellen (editor), Annette Vissing-Jorgensen, Dexin Zhou, several anonymous referees, and seminar and conference participants at the NBER Summer Institute, SFS Cavalcade, WFA, NYU Stern, Stanford GSB, Columbia University, Cornell, MIT Sloan, New York Fed, Dallas Fed, Notre Dame, Texas A&M, HEC Montreal, ARC Boston, University of Bath, MFA, Institut Louis Bachelier, SoFiE, University of Amsterdam, CKGSB, University of Luxembourg, VSCE, SITE, CREDIT, Essec, San Francisco Fed, Federal Reserve Board, Central Bank of Chile, UCLA, UC Boulder, Fordham University, DIW, UNC-Duke, PRI, Hong Kong University, Bundesbank, UCLA, NYU Shanghai Volatility and Risk Institute conference, Johns Hopkins University, and IWH Halle for comments. We thankfully acknowledge financial support from the Norges Bank through a grant to the Volatility and Risk Institute at NYU Stern and the Center for Sustainable Business at NYU Stern, and from the Tobin Center for Economic Policy at Yale University. We are grateful to Xuran Zeng and Joachim Rillo for excellent research assistance.

[†]Palantir. Email: georgij.v.alekseev@gmail.com.

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

[§]USC Marshall. Email: qmaingi@marshall.usc.edu.

[¶]Chicago Booth. Email: julia.selgrad@chicagobooth.edu.

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

Disclosure Statement

Georgij Alekseev declares that he has no conflicts of interest to disclose.

Stefano Giglio declares that he has no conflicts of interest to disclose.

Quinn Maingi declares that he has no conflicts of interest to disclose.

Julia Selgrad declares that she has no conflicts of interest to disclose.

Johannes Stroebel declares that he has no conflicts of interest to disclose.

Climate change presents a major global challenge. In addition to a wide range of social implications, both the physical effects of climate change and the regulatory efforts to slow carbon emissions have the potential to substantially disrupt economic activity. As investor awareness of the economic and financial risks from climate change has increased, there has been a rising demand by retail and institutional investors for financial instruments that hedge these risks (see Krueger et al. 2020, Giglio, Kelly & Stroebel 2021, Stroebel & Wurgler 2021, Acharya et al. 2023, Giglio, Maggiori, Stroebel, Tan, Utkus & Xu 2023). At present, only a small number of instruments are designed to directly hedge various climate risks, most prominently the relatively illiquid catastrophe bonds. However, investors interested in hedging climate risks can still build hedge portfolios using other assets, such as stocks or bonds, that are exposed to climate risks. To do so, investors need to identify which assets would benefit and which would lose from climate risk realizations. A long-short portfolio that buys the former and sells the latter would increase in value when climate risks materialize, and thus provide a valuable climate risk hedge (see the discussion in Engle et al. 2020).

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

In this paper, we propose a new methodology to build portfolios to hedge newly emerging risks such as those from climate change. Our *quantity-based* approach uses cross-sectional information on investors' trading activity to identify which stocks to hold in a hedge portfolio. The approach first identifies "idiosyncratic climate belief shocks," shocks that shift the climate risk beliefs of a small set of investors. While such shocks do not move asset prices—they are, after all, idiosyncratic—they can still influence the affected investors' trading activity. Based on this insight, the quantity-based approach explores how investors *trade* in response to these idiosyncratic climate belief shocks to learn how their demand for each asset is shifted by changes in perceived aggregate climate risks. Concretely, our approach identifies which stocks investors tend to buy and sell after becoming more concerned about climate risks; the hedge portfolio is then built by going long the former and short the latter. This portfolio is expected to rise in *price* when aggregate climate risks materialize. The reason is that, while idiosyncratic shocks only move quantities and not prices, the occurrence of an aggregate climate shock affects many investors. As long as investor demand for different assets responds to aggregate climate risk shocks in similar ways to how it responds to idiosyncratic shocks, the correlated shift in demand of many investors will move prices.

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

We propose two ways to identify idiosyncratic shocks to the climate risk beliefs of mutual fund managers. The first exploits geographically localized extreme heat events that have been shown in prior work to affect beliefs about aggregate climate risk (see, e.g., Egan & Mullin 2012, Deryugina 2013, Joireman et al. 2010, Li et al. 2011, Fownes & Allred 2019, Sisco et al. 2017, Constantino et al. 2022, Sisco & Weber 2022). Our baseline measure of local extreme heat shocks is based on the presence of fatalities and injuries due to extreme heat in a county. We confirm that the occurrence of this shock, which only affects a small number of investors, leads to local increases in Google searches for the term "climate change," consistent with the prior literature's finding that such local heat shocks induce updates in climate beliefs.

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

We then study how U.S. active non-sector mutual funds change their portfolio allocations across industries when their managers experience one of these idiosyncratic climate belief shocks.[1] The two belief shocks are essentially uncorrelated with each other, suggesting that mutual funds' disclosures are updated in response to many events other than heat shocks that also affect managers' climate beliefs in idiosyncratic ways. Despite this finding, and despite the presence of sizable estimation error in measuring funds' trading responses, we observe a significant correlation in the industries that funds buy and sell in response to our two idiosyncratic climate belief shocks: fund managers that experience extreme heat events generally buy and sell similar industries as managers that report increased climate risk concerns in their investor disclosures. This finding suggests that the portfolio adjustments contain useful and consistent signals about different industries' climate risk exposures.

Several interesting patterns emerge by studying which industries are bought or sold in response to climate belief shocks. For example, by the end of 2019, the auto industry had one of the strongest positive quantity responses; that is, mutual funds tended to buy auto stocks after increasing concerns about climate risks. While this finding may appear surprising at first glance—automobiles are, after all, an important source of current carbon emissions—the managers' reaction may reflect their beliefs that the transition to electric vehicles provides substantial opportunities for incumbent auto makers to sell more vehicles over the coming years. We also find that investors tended to buy firms in the semiconductor sector, many of which build products, such as solar panels and chips for the management of smart grids, that are important for decarbonizing the economy. On the flip side, after become more concerned about climate risks, investors generally sold equities in the real estate sector, whose assets are tied to particular locations that might be exposed to climate risks.

Of course, we cannot know with certainty what factors determine fund managers' assessments of the various industries' climate risk exposures. Indeed, if we had a good understanding of different industries' exposures ex-ante, the construction of hedge portfolios would not require elaborate approaches. The key benefit of our approach is that it can be effective *even*

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

without a full understanding of the economic determinants of each industry's climate risk exposure. Instead, to build hedge portfolios based on observed quantities, we rely on the consistency of fund managers' asset demand response to idiosyncratic and aggregate shocks.

We validate this assumption in several ways, for example by showing that the trading activity in response to idiosyncratic shocks is similar across investors and belief shocks. Most importantly, we show that the cross-sectional quantity information is useful for predicting industry price responses to aggregate climate risk shocks. To do this, we use the trading of mutual funds in response to managers' idiosyncratic belief shocks to build long-short industry portfolios and study their out-of-sample hedging performance with respect to various measures of *aggregate* climate risk. We build separate hedge portfolios using our two idiosyncratic belief shocks (based on heat events and investor disclosures) as well as a hedge portfolio based on combining these shocks. We then evaluate their performance against alternative approaches for constructing hedge portfolios that have been proposed in the literature.

The first alternative approach—which we call the "narrative" approach—chooses long and short positions based on economic reasoning. For example, such an approach might suggest buying clean energy companies, selling coal companies, or buying companies with high ESG scores as in Engle et al. (2020), Pástor et al. (2021), and Hoepner et al. (2018). This approach will hedge climate risk if the underlying economic intuition is aligned with that of the average investor determining how prices move when aggregate risk materializes. Similar to the quantity-based approach, the narrative approach has the advantage that it does not require long time series to be implemented. Instead, it requires having correct priors about investors' perceptions of each industry's exposure to climate risk.

The second alternative approach is the "mimicking portfolio" approach of Lamont (2001), which projects a climate risk series onto a set of base asset returns using time-series information. The mimicking portfolio approach relies on the availability of substantial time-series data, frequent risk realizations, and a substantial time-series stability of risk exposures: 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 asset returns during prior 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 hedge portfolio returns and measures of aggregate climate shocks between 2015 and 2019. For the mimicking portfolio approach and the quantity-based approach, we construct the hedge portfolios using rolling five-year windows of price and quantity data, respectively.[2] To evaluate the hedging performance with respect to aggregate climate 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 time series of news about physical and regulatory climate risks. Rather than choosing a preferred climate risk series, we evaluate how the portfolios perform in hedging various series constructed by Engle et al. (2020), Faccini et al. (2021), and Ardia et al. (2020) as well as national temperature shocks and attention to

[2] Prior to 2010, climate risks were unlikely to be incorporated into market prices and unlikely to affect investor behavior, making all approaches difficult to implement. Consistent with this assessment, we show that none of the approaches can meaningfully hedge climate news between 2000 and 2010.

climate risk as measured through Google searches.

We document several patterns. First, at a broad level, hedging climate risks is difficult, and few approaches manage to achieve more than a 20% out-of-sample correlation with the climate shock series, confirming and extending this finding from Engle et al. (2020). Second, both the mimicking portfolio approach and the narrative approach provide mixed results at best: 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 performances compared to the alternatives. For example, our quantity-based portfolio based on the heat shock yields positive out-of-sample correlations with *all* of our aggregate climate shock series, with average correlation of about 17%. The disclosure-based portfolio does almost as well, while the quantity portfolio built using the pooled shock has the highest average correlation with the various climate risk series. This validates the idea that the cross-sectional information on which the quantity portfolios are based is useful for constructing hedges of aggregate climate news shocks.[3]

In addition to documenting the strengths of our quantity-based methodology, our empirical results highlight some important downsides of the traditional approaches. The mimicking portfolio approach is very sensitive to the availability of time-series data, and suffers when the time series is short. As an illustration, consider a mimicking portfolio that *only* uses the S&P 500. While this portfolio is composed of only one asset, historical data is still required to establish whether to take a long or short position: will the broader stock market increase or decrease upon the realization of climate risks? This relationship turns out to be unstable over time: during 2010-2014, the S&P 500 comoved positively with climate risk realizations, while during 2015-2019, it comoved negatively, highlighting the challenges of the mimicking portfolio approach for constructing successful climate hedges. Adding more base assets in the construction of the mimicking portfolio can help to better target the hedge, but requires estimating more parameters, again a problem in short samples.

Narrative-based portfolios are immune to such short-sample issues, since historical data is not used to determine the positions of different assets. However, deciding on positions in an *a priori* way is challenging: as an example, for many industries, the different co-authors of this paper would have picked different holdings for their hedge portfolios. In the data, we find that various seemingly plausible narrative portfolios have very different out-of-sample hedging properties. For example, buying clean energy stocks appears to provide a solid hedge against negative climate news, but shorting traditional oil and gas firms does not.

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 method-

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

ology, since climate risks have only recently attracted investors' attention. As a result, there is not enough time-series data to allow researchers to precisely estimate the climate risk exposures of different assets based on price data alone. However, our approach can, in principle, be applied to hedging any macro risk series for which similar idiosyncratic shocks (e.g., stemming from local events or measurable through investor disclosures) affect investors' beliefs about aggregate risks. For example, in recent work, Kuchler & Zafar (2019) show that locally experienced house price movements affect expectations about future U.S.-wide house price changes; they also show that personally experienced unemployment affects beliefs about future national unemployment rates. 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—or discussions of these topics in investor disclosures—allow us to construct portfolios that perform well at hedging innovations in the corresponding national series.

Our work contributes to a growing literature that studies the interaction between climate change and asset markets (see Giglio, Kelly & Stroebel 2021, for a recent review). In equity markets, Bolton & Kacperczyk (2021*a*) and Hsu et al. (2022) document that high-pollution firms are valued at a discount, and Giglio, Maggiori, Stroebel, Tan, Utkus & Xu (2023) find that retail investors expect high-ESG firms to underperform the market. Engle et al. (2020) find that stocks of firms with lower exposure to regulatory climate risk experience higher returns when there is negative news about climate change, and Barnett (2020) shows that increases in the likelihood of future climate policy action lead to decreased equity prices for firms with high exposure to climate policy risk. Choi et al. (2020) document that the stocks of carbon-intensive firms underperform during periods of abnormally warm weather, where investors' attention to climate risks is likely heightened. Climate risk has also been shown to affect prices in other asset classes such as real estate markets (Baldauf et al. 2020, Bakkensen & Barrage 2022, Bernstein et al. 2019, Giglio, Maggiori, Rao, Stroebel & Weber 2021, Murfin & Spiegel 2020) and municipal bond markets (Painter 2020, Goldsmith-Pinkham et al. 2021, Acharya, Johnson, Sundaresan & Tomunen 2022).

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

1 Quantity-Based Portfolios: A Simple Model

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

Setup. Consider a continuum of investors $i \in [0,1]$ who choose a portfolio of securities A and B. There are no other assets and no consumption available to investors at the time of this asset allocation decision. Investor i's demand for security A is given by $q_A\big(p_A, \epsilon_A(i), \epsilon_B(i)\big)$, where p_A is the price of security A relative to that of security B, $\epsilon_A(i)$ gives investor i's beliefs about the future payoffs of security A, and $\epsilon_B(i)$ gives investor i's beliefs about the future payoffs of security B (all expressed in units of security B's price). For simplicity, assume that $q_A\big(p_A, \epsilon_A(i), \epsilon_B(i)\big) = f(p_A) + g\big(\epsilon_A(i)\big) + h\big(\epsilon_B(i)\big)$, with f, g and h continuously differentiable, and $\frac{\partial f}{\partial p_A} < 0$. The market-clearing condition is:

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

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 about stock A's future payoff can be decomposed into a common component ν_A and an investor-specific idiosyncratic component $\omega_A(i)$, such that $\epsilon_A(i) = \nu_A + \omega_A(i)$. The common belief ν_A is driven by shocks or news that are observed by all investors, and that correspond to the types of news that investors might want to hedge (e.g., well-publicized news about accelerating global warming that shifts all investors' beliefs about physical climate risks). The idiosyncratic belief component $\omega_A(i)$ can instead be affected by "local" events that are only observed or experienced by investor i (e.g., a localized heat wave in the location of investor i 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 on carbon emissions will reduce the future profitability of stock A (formally, we also assume that the shock to beliefs about stock A does not change beliefs about stock B). 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,

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

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 move equilibrium prices. Thus, investor i's change to her *equilibrium* allocation q^* is directly informative about her demand sensitivity to changes in beliefs, $\frac{\partial q}{\partial \epsilon_A(i)}$.

From Quantities to Prices. Suppose now there is news about stock A that affects *all* investors' beliefs: a change in ν_A. For example, an announcement by a senior politician might cause all investors to believe that climate change regulation has become more likely, reducing the expected profitability of heavily-emitting firm A. By the implicit function theorem and the chain rule, equilibrium price responses are given by:

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

In words, the sensitivity of prices to national news is directly proportional to average quantity sensitivities, $\int_{i=0}^{i=1} \frac{\partial q_A}{\partial \epsilon_A(i)} di$.[4] Together with the earlier result, this shows how idiosyncratic quantity responses can help predict national price responses. Intuitively, by studying how investors react to local shocks that have no effect on the equilibrium price, we can predict how their demand shifts in response to news that affects all investors. Aggregate news then moves the demand functions of many investors simultaneously, leading to price movements in response to aggregate shocks. Therefore, a hedging portfolio built using idiosyncratic quantity data can hedge aggregate climate shocks.

2 Idiosyncratic Belief Shocks & Portfolio Changes

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

We developed two approaches to identifying such idiosyncratic belief shocks, one built on the presence of extreme heat events in the locations of the investors, and one based on investor disclosures. We next provide details on the construction of these two types of idiosyncratic belief shocks, before exploring investor trading responses to them.

2.1 Idiosyncratic Belief Shocks: Extreme Heat

Our first idiosyncratic climate belief shock is motivated by an extensive literature that identifies local extreme heat events as important drivers of climate change attention and beliefs in affected populations (e.g., Egan & Mullin 2012, Deryugina 2013, Joireman et al. 2010, Li et al. 2011, Fownes & Allred 2019, Sisco et al. 2017, Herrnstadt & Muehlegger 2014, Constantino et al. 2022, Sisco & Weber 2022, Zaval et al. 2014). Following this work, we aim to identify instances of geographically concentrated extreme heat events, based on the premise (which we validate below) that investors experiencing these events will, on average, update their beliefs or attention related to climate change.

[4] 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).

The observed tendency of individuals to adjust their climate change beliefs or attention based on locally experienced extreme heat events is unlikely to be fully rational—after all, if such events were truly informative about aggregate climate risk, the relevant information could be easily obtained and incorporated by all investors, not just by those experiencing the extreme heat event. Nevertheless, the importance of such 'local experience effects' is highly consistent with a large literature that has documented how the beliefs and behaviors of sophisticated investors, including mutual fund managers and equity analysts, are affected by local and personal experiences.[5] Most directly relevant is recent work by Reggiani (2022), who shows that equity analysts experiencing warmer temperatures tend to issue more pessimistic forecasts for firms exposed to regulatory and physical climate risks.

It is also useful to note that the construction of climate hedge portfolios using investor trading responses to extreme heat does not require that climate beliefs are the *only* factor through which extreme heat events influence investors' trading. Instead, all we require is that there is some signal in the trading in response to local heat shock that is informative about trading in response to aggregate climate risk realizations, leading to successful hedging of these climate risk realizations out of sample.

Measuring Extreme Heat Events. There are several plausible ways to define extreme heat events as potential shifters of climate change attention and beliefs of the affected populations. Our baseline measure identifies events that involve fatalities or injuries due to extreme heat. In the Appendix, we show that our results are robust to considering several other definitions of extreme heat events, such as those built on temperature anomalies relative to local historical averages.

We construct our baseline measure of extreme heat events using monthly information from NOAA's National Center for Environmental Information, as collected in the *Spatial Hazard Events and Losses Database for the United States* (SHELDUS) database. We use the reported per-capita injuries and per-capita fatalities attributed to extreme heat to construct local heat shocks. For our baseline shock, we define a county-month as experiencing a heat shock if there were positive numbers reported for either per-capita injuries or for per-capita fatalities (or both).[6] Appendix Figure A.1 shows the spatial distribution of the resulting heat shock. Panel A of Table 1 shows the frequency of the occurrence of the shock in our data, highlighting that about 0.1% of all county-months in the U.S. between 2010 and 2019 had reported fatalities or injuries due to extreme heat.

---

[5]Personal experiences have been shown to affect the beliefs and behaviors of sophisticated agents across a wide range of other settings. For example, Malmendier et al. (2021) show that personal lifetime experiences of inflation impact the inflation forecasts and voting patterns of FOMC members. Alok et al. (2020) show that fund managers located in a major disaster region tend to invest less in stocks from the disaster-stricken area compared to managers who are situated further away. Chang (2022) shows that mutual fund managers extrapolate from local industry conditions in their investment strategies. More generally, professional investors' decisions have been shown to deviate from rationality in a range of other settings (e.g. Kuchler et al. 2022, Hirshleifer et al. 2021, Kaustia et al. 2008, Haigh & List 2005, and many others).

[6]In the baseline SHELDUS data, the threshold for reporting a non-zero value for fatalities or injuries per capita requires at least 1 fatality or injury per 200k residents, respectively; we examine robustness to this threshold in Section 3.5.

Table 1: Summary Statistics on Idiosyncratic Climate Belief Shocks

Panel A: Local Shocks: Summary

Climate Shock	Event Description	Frequency Monthly	Frequency Sample
Heat Shock	Injuries or fatalities	0.10%	1.32%
Disclosure Shock	Change in fund disclosures about transition risk	-	0.15%
Pooled Shock	Pool of heat and disclosure shock	-	1.76%

Panel B: Local Shocks: Sample Jaccard Correlations across Fund-Quarters

	Heat Shock	Disclosure Shock
Heat Shock	1.00	
Disclosure Shock	0.00	1.00

Note: Panel A provides an overview of the idiosyncratic belief shock measures. The "monthly" frequency shows the share of county-month observations in the U.S. from 2010 to 2019 that experience an extreme heat event. The "sample" frequency shows the share of observations in our final sample that are assigned a climate belief shock, with the unit of observation being a pair of consecutive (3 months apart) filings of a given mutual funds' holdings (see Section 2.3). Panel B shows the Jaccard correlations among the shock measures across all fund-quarters in our sample. Jaccard correlations quantify the intersection's size relative to the union of two sets. We applied Jaccard correlation to address the inherent binary nature of data.

Heat Shocks and Climate Change Attention and Beliefs. We next explore whether this measure of local heat shocks affects local climate change attention or beliefs, as measured by Google searches related to climate change (see Stephens-Davidowitz 2014, Choi et al. 2020, for similar approaches). Since Google Trends data is not available at the county level and is often missing at the MSA level, we conduct this analysis at the state-month level.

The Google search series measures relative interest in a topic, such as the fraction of all Google searches in a region for "climate change." For each geographic unit, Google scales the relative search interest for a topic to be between 1 and 100. This means that, for each state, the month with the most relative searches for a given term receives a score of 100. All other months' scores represent their relative searches as a fraction of the relative searches of the highest-ranked month. For example, if for California, July 2008 is the month with the most relative searches for "climate change" and February 2009 has half as many relative such searches, then California's score in February 2009 would be 50. Given this multiplicative scaling factor, we explore how local climate shocks affect the logarithm of the Google search index using the following specification:[7]

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

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 indicator for a local extreme heat event, set equal to one if at least

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

one county within the state experienced a heat shock during the month.[8] State and time fixed effects are given by δ_s and γ_t. Consistent with the work cited above, Table 2 shows that extreme heat shocks are associated with more attention paid to climate risks, though the aggregation to the state level leads to only marginally statistically significant results.

Table 2: Heat Shocks and Climate Attention

	Log(Google Search Volume)
Heat Shock	0.016*
	(0.008)
R^2	0.75
State & Month FE	Y
N	5,823

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: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

2.2 Idiosyncratic Belief Shocks: Investor Disclosures

Our second idiosyncratic climate belief shock is based on changes in the discussion of climate risks in mutual funds' semi-annual shareholder reports. In contrast to the local heat shock, which captures idiosyncratic drivers of changes in climate risk beliefs, this disclosure-based approach directly measures changes in the stated climate beliefs of different investors, without identifying why a particular investor may have changed her beliefs.

To construct this belief shock, we work with the semi-annual shareholder reports of actively managed mutual funds, which are filed as N-CSR reports with the Securities and Exchange Commission. We first extract sentences from these N-CSR reports that use the following terms related to climate change: climate change, carbon emission(s), greenhouse gas(es), or global warming. These terms were selected based on their cosine similarity to "climate change" in Google's word2vec implementation. We augment this list with the following terms referring directly to transition and physical risks: carbon tax, carbon pricing, extreme weather, extreme temperature, extreme precipitation, flooding, drought, or sea level rise. We further extract one sentence before and after each selected sentence to construct a more complete passage with context. We then feed the extracted passages into GPT4 to classify whether each passage refers to physical and/or transitional climate risk (we provide example classifications in Appendix Table A.1, and our GPT prompt in Appendix A.1).

Overall, we identify 133 reports (corresponding to 348 fund-quarters, as a report can cover multiple funds; we use quarterly frequency to describe the data as our final matched dataset will be at this frequency) as expressing concerns about or considerations related to climate risk.[9] For our baseline disclosure shock, we exclude climate-focused investment

[8]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.

[9]One might be concerned that this relatively small number of N-CSR reports that describe concerns

funds' disclosures, which leaves us with a total of 231 identified fund-quarters.[10]

2.3 Holdings and Location Data

There are several reasons why mutual fund managers are a natural focus of our study of trading responses to idiosyncratic climate belief shocks. First, mutual funds make up a substantial share of the investor universe and their portfolio holdings are observable 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). And second, data on the location of mutual fund advisors allows us to construct idiosyncratic belief shocks based on locally experienced heat events, while the content of N-CSR reports allows us to construct an alternative measure based on the extent of self-reported concerns about climate risks.

For our quantity-based approach to identify hedge portfolios, mutual fund managers must adjust their portfolios in response to perceived changes in climate risk. Such portfolio adjustments could occur for various reasons.[11] Managers may believe that equilibrium valuations do not fully account for climate risks and that they can thus earn alpha by reducing their holdings of more exposed stocks (Krueger et al. 2020). Alternatively, managers may view climate change as an additional risk to hedge, either because of their investors' preferences (Ceccarelli et al. 2021, Giglio, Maggiori, Stroebel, Tan, Utkus & Xu 2023) or to manage flows in response to aggregate climate disasters (Dou et al. 2022). For some investors, several of these motivations may be relevant at the same time. For example, in a recent article exploring CalPERS' motivations for adjusting portfolios based on climate risk exposures, it was described that CalPERS *"plan[s] to generate alpha from climate investments"*, but also that *"another objective of the focus on climate investments is also to build more resilience into the portfolio"* (Cashion 2024). Importantly for our purpose, these different motivations generate similar predictions: when managers become more concerned about climate risk, they sell stocks that are more exposed to that risk, and buy stocks that are less exposed.

about climate risks suggests that those risks are not a first order consideration of mutual fund investors. However, it is likely that not all investors who consider climate risks talk about these risks in the N-CSR filings. Indeed, more direct evidence on the relevance of climate risks for mutual fund investors comes from Krueger et al. (2020), who survey institutional investors about the perceived financial materiality of three sources of climate risk: physical, regulatory, and technological. They report that, on average, *"respondents regard the financial materiality of climate risks to be somewhere between 'important' and 'fairly important.'"* Similar evidence can be found in a related survey by Stroebel & Wurgler (2021). Our sample will likely select a (small) subset of those funds that consider climate change relevant.

[10] We identify these climate-focused funds based on the fund name including one of the following: climate, esg, sustainable, carbon, and/or green. Such funds represent 4.5% of fund quarters in our sample. Given the higher general propensity of these funds to discuss climate risks in investor disclosures, changes in disclosures that we flag could be less informative about changes in fund managers' beliefs.

[11] Consistent with this discussion, Giglio, Maggiori, Stroebel, Tan, Utkus & Xu (2023) find that retail investors mention several possible motivations for investing in ESG assets, including their expected performance, the possibility that such assets hedge climate risk realizations, and ethical reasons.

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.[12] We combine the holdings data with fund characteristics from CRSP.[13] Since we hope to identify deliberate fund manager asset reallocations in response to idiosyncratic climate belief shocks, we restrict our analysis to actively managed funds, i.e., those with an Investment Objective Code of 2 ("Aggressive Growth"), 3 ("Growth"), 4 ("Growth & Income"), or missing, and that have "Equity Domestic Non-Sector" as their CRSP Objective Code (see Song 2020).

We obtain stock-level characteristics from CRSP and Compustat and assign end-of-month prices from CRSP to the holdings. We restrict holdings to assets with share codes 10, 11, 12, and 18, and exchange codes 1, 2, and 3, focusing the assets for our hedging portfolio on North American common stocks. We obtain firm GICS industry codes from Compustat by merging the stocks on their CUSIP identifiers. The first four GICS digits determine the stock's classification into the 24 "industry groups" that are the main focus of our analysis.[14]

Measuring Active Portfolio Changes. In our main analysis, we explore how idiosyncratic climate belief shocks induce changes in the portfolio share of fund f in industry I *through active trading* between consecutive holdings reports. Since holdings are usually reported at three-month intervals—often, though not always, at the end of a quarter—we measure fund composition changes over such intervals.[15] We perform our baseline analysis at the industry level, since the sparsity of the stock-level holding matrix would lead to very noisy estimates of stock-level exposures. For every fund f and month t with a holdings report, we define the active change in industry I holdings as:

$$ActiveChanges^I_{f,t} = 100 * \left[\left(\frac{\sum_{j \in I} P_{j,t-3} S_{f,j,t}}{\sum_{j} P_{j,t-3} S_{f,j,t}} \right) - \left(\frac{\sum_{j \in I} P_{j,t-3} S_{f,j,t-3}}{\sum_{j} P_{j,t-3} S_{f,j,t-3}} \right) \right] \frac{1}{\left(Share^I_{t-3}\right)}. \qquad (2)$$

$P_{j,t-3}$ is the price for stock j at the end of month t-3, the time of the prior report. $S_{f,j,t}$ is the number of shares of stock j held by fund f at the end of month t, and $Share^I_{t-3}$ is the market capitalization of industry I as a share of the U.S. stock market at the beginning of the interval. The term in square brackets thus captures the active three-month change of the share of industry I in fund f's portfolio.[16] The reason for scaling by industry size is that a

[12] We restrict our sample to end-of-month reports, that is reports that have an as-of date on the 27th of the month or later. The S12 database occasionally contains multiple records for the same fund-CUSIP-report date tuple, corresponding to multiple SEC filings by the same fund; for our analysis, we use data from the first filing submitted to the SEC to ensure real-time implementability.

[13] We link mutual funds across Thomson Reuters and CRSP using their Wharton Financial Institution Center Number (WFICN) as reported in WRDS MFLINKS. Our observations are thus set to the WFICN-quarter-CUSIP level, which corresponds to quarterly, stock-level fund holdings.

[14] 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.

[15] For the disclosure based measure, since the N-CSR report is filed semi-annually, we instead measure fund composition changes over six month intervals. We provide more detail in Section 2.4

[16] We verify below that our approach is robust to considering a separate variable, $TotalChanges^I_{f,t}$, where the first fraction uses $P_{j,t}$ instead of $P_{j,t-3}$, i.e., current period holdings are valued at current period prices.

given change in the portfolio share of an industry—i.e., a shift of a given dollar amount—is likely to induce larger price movements in smaller industries. We winsorize $ActiveChanges$ at the 1% level to mitigate the effect of outliers due to, e.g., fund mandate changes.

Investor Location Data. We obtain data on the locations of mutual fund advisers, which are primarily responsible for making investment decisions (see Chang 2022). Specifically, we parse adviser locations from funds' SEC filings (N-SAR filings until 2017, N-CEN filings from 2018 onward). Since these 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, allowing us to match 84.1% of fund-quarter observations. We drop all observations that do not match, leading to a sample with 2,496 unique funds, making up 57,961 fund-quarter observations (23.2 observations per fund on average) between 2010 and 2019.[17]

For 64.0% of funds, all advisers reside in the same county. For the extreme heat shock, whenever funds have multiple advisers who are not all located in the same county, we assign fund-level climate shock exposure as an average of fund adviser shock exposures. For example, if a fund has two advisers in county A and one adviser in county B, and county A is affected by a local extreme heat shock, we assume the fund is affected by $2/3$ of a local extreme heat shock; our results are robust to alternative aggregation choices.

Figure 1 shows the geographic distribution of fund advisers for the subset of funds where all advisers reside in the same county. While some areas have a larger concentration of advisers, advisers are generally spread throughout the country. The table below the figure shows that about a quarter of advisers are located in New York (most of them in New York City), 13.5% are located in Massachusetts (most of them in Boston), and 9.5% are located in California. This gives us important geographic variation and therefore differential of fund managers exposure to local extreme heat shocks.

Summary Statistics. Panels A and B of Table A.3 present summary statistics on the GICS industries and the portfolio holdings of the mutual funds in our final sample. For example, in an average quarter, funds in our sample held 245 unique companies in the Energy sector (GICS code 1010). In the average sample-quarter, the energy sector market share held by funds in our sample was 7.6%. The smallest industry by average share held by funds in our sample was "Auto & Components", comprising of an average of 44 held firms and an average market capitalization share of 0.9%. Panel B shows that, on average, mutual funds held 211 unique firms across 19.6 unique industries.

Panel C of Table A.3 shows summary statistics on $ActiveChanges^{I}$. Intuitively, active changes of 0 imply that there were no active changes in industry I's relative weight within fund f's portfolio, while active changes of 1 imply that I's weight in the portfolio increased by a percentage equal to one percent of the industry's market share (e.g., if industry I has

This alternative approach takes price changes into account, and would be a more suitable model if we assume that funds constantly rebalanced their portfolios.

[17] 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

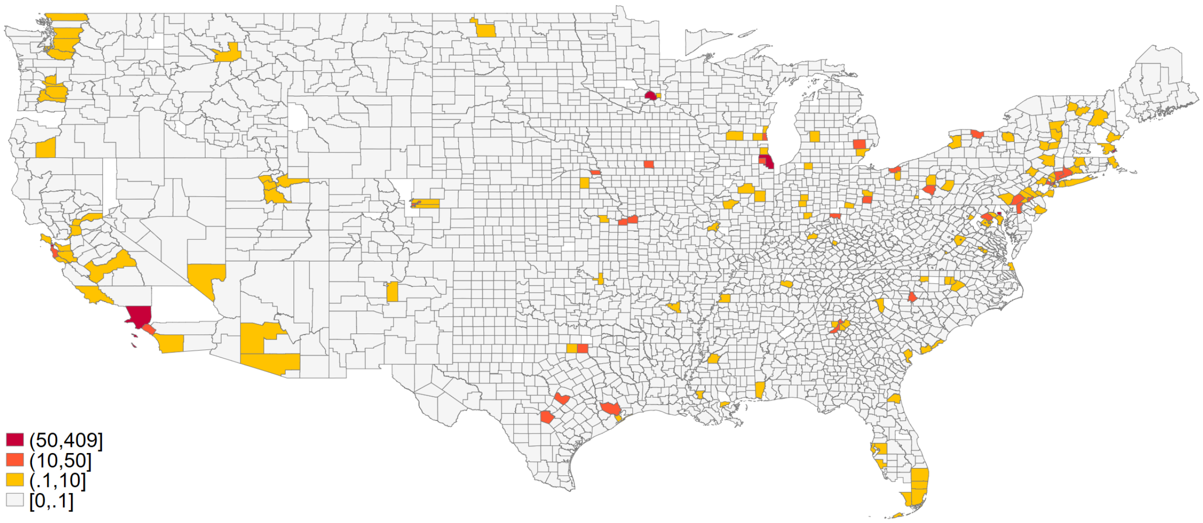

Panel A: Adviser Locations - Largest Counties				
FIPS	County	State	% Funds	% Fund-Quarters
36061	New York	NY	20.5	21.2
25025	Suffolk (Boston)	MA	13.0	10.3
17031	Cook (Chicago)	IL	4.9	4.8
06075	San Francisco	CA	3.8	3.0
06037	Los Angeles	CA	3.0	3.6

Panel B: Adviser Locations - Largest States			
State name	State	% Funds	% Fund-Quarters
New York	NY	24.0	24.9
Massachusetts	MA	13.5	10.4
California	CA	9.6	9.0
Illinois	IL	6.3	6.5
Pennsylvania	PA	5.3	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.

a 10% market share, and a fund increased I holdings from 5% to 15% of its portfolio, then $ActiveChanges^I$ would be 100). The average and median active changes in our sample are close to zero.

2.4 Estimating the Response to Idiosyncratic Climate Shocks

To understand how mutual funds' portfolios change with managers' 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 and δ^I_t represents year-month fixed effects. $S_{f,t}$ captures the presence of an idiosyncratic climate belief shocks.

For the local extreme heat shock, $S_{f,t}$ is set to "1" whenever there was a heat shock in the fund advisers' county in at least one of the three months during which we measure active portfolio changes (i.e., months t, t-1, and t-2). Table 1 shows that about 1.32% of fund-holdings changes have a positive value for this shock. The data is at the quarterly frequency, i.e., $ActiveChanges_{f,t}$ is the change in holdings of fund f between the end of quarter $t-1$ to end of quarter t, and $S_{f,t}$ refers to a heat shock a fund experiences during quarter t.

For the disclosure-based shocks, the N-CSR report is only available at a semi-annual frequency, so we instead construct the sample semi-annually. Our sample construction follows a three step process. First, to match disclosure reports to holdings data while ensuring real-time implementabilty, we match each report to the next available mutual fund holdings data, up to two months in the future. Second, to measure active portfolio changes, we use holdings data from 6 months prior to the matched holdings data. Finally, we construct our disclosure shock based on the *change* in discussion of physical and/or transition risks from the prior, semi-annual N-CSR report.[18] Specifically, we code the shock as '1' if the current report expresses concern about climate change while the previous report did not express concern, and '0' otherwise. Table 1 shows that about 0.15% of fund-holdings changes have a non-zero value for this shock.

In addition to these two individual shocks, we also consider a pooled shock, constructed as the sum across the two individual shocks. Combining the shocks raises the issue of matching the time windows, since the heat shock can be constructed quarterly, but the disclosure shock is only available semi-annually. To address this issue, we construct a version of the disclosure shock based on a similar logic to our heat shock, and construct the pooled sample with three month active holdings changes. Specifically, we set the disclosure portion of the pooled shock to 1 if any month of the three-month holdings change window lies between two report release dates that indicate an increase in climate concern, and zero otherwise.[19]

The main objects of interest are the estimates of $\beta^{I,S}$. For each industry I, these represent the differential active change in fund holdings of that industry for funds affected by

[18] To ensure that we are capturing regular, semi-annual reports, we only prior consider N-CSR reports filed between five and seven months after the matched, prior report.

[19] Note that this increases the effective sample frequency of the disclosure shock, explaining why the sample frequency of the pooled shock shown in Table 1 is greater than the sum of the two component shocks.

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

2.4.1 Baseline Estimates

Table 3 reports the estimated $\beta^{I,S}$ coefficients for the idiosyncratic climate belief shocks described above, sorting industries by the estimate obtained using the "Pooled" shock.[20] The table shows estimates obtained over the 2010 to 2019 period, together with stars representing statistical significance for standard errors clustered at the fund level to account for possible correlation of holdings changes within funds over time. To interpret the magnitudes, recall that active changes are defined as 100 times the portfolio percentage change, scaled by the industry relative market cap. So, for example, the Auto & Components industry has a relative market cap of around 1.1% on average in our sample. A coefficient of 4.21, for the pooled shock, shows that on average funds that are treated by the shock increase their allocation by $1.1\% \times 4.21 = 4.5$bp; for the disclosure shock, where the coefficient is 24.42, the increase in allocation is 27bp. Appendix Table A.4 reports the t-statistics for each $\beta^{I,S}$ estimate and the p-values of their joint significance tests. Industries towards the top of Table 3 are those that mutual fund managers disproportionately buy after receiving idiosyncratic climate belief shocks, while industries towards the bottom are those that investors disproportionately sell (relative to managers who do not receive a shock).

A first key observation is that estimates of $\beta^{I,S}$ are generally noisy. This conclusion is perhaps unsurprising given the relative rarity of the idiosyncratic belief shocks and the fact that many considerations beyond climate risks drive funds' investment decisions. However, despite the substantial estimation noise, we note that many of the industry estimates in the top fifth and bottom quintiles (according to the pooled shock, which has more power) are statistically significant at standard confidence levels.

The ordering of industry-specific climate quantity betas appears broadly correlated across the "Heat" and "Disclosure" based idiosyncratic climate belief shocks, despite the fact that the correlation of the two belief shocks in the panel is zero, and the corresponding quantity betas are thus based on independent information. Panels A and B of Table 4 formally show this by reporting the correlation and rank-correlation, respectively, of the industry-specific climate quantity betas obtained from running the regression in Equation 3 for the period of 2010-2019. It appears that mutual funds change their portfolios in broadly consistent ways in response to the heat shock and the disclosure-based shock (correlation of 0.20).

Our interpretation of the estimates in Table 3 is that they indicate the relative climate risk exposures of industries as perceived by fund managers. Analzying whether this inferred ranking of climate risk exposures is reasonable is difficult, in part because the estimates are naturally very noisy, and because each industry's exposure to climate risk is determined by a variety of economic mechanisms, not all of which might be immediately apparent. The

[20] Appendix Table A.5 reports the industry betas estimated without scaling *ActiveChanges* in Equation (2) by industry share. The correlation between the scaled and unscaled industry betas is 0.87. Appendix Figure A.2 shows the scatterplot of industry rankings based on scaled and unscaled betas.

Table 3: Industry-Specific Climate Quantity Betas

GICS	Description	Pooled Shock	Disclosure Shock	Heat Shock
2030	Transportation	4.79*	24.42**	0.95
2510	Auto & Components	4.21*	24.58*	2.88
4530	Semiconductors & Equip.	2.46	3.80	4.60**
2010	Capital Goods	2.38*	13.07**	0.53
1510	Materials	1.69	6.45	1.34
4010	Banks	1.60*	1.58	2.46**
3030	Household & Pers. Prod.	1.34	6.18	-0.14
1010	Energy	1.32*	4.50	1.77*
4520	Tech. Hardw. & Equip.	0.96	-8.40	3.81***
2530	Consumer Services	0.20	-2.06	0.27
4020	Diversified Financials.	-0.12	2.15	0.44
4510	Software & Services	-0.19	2.46	0.91
3010	Food & Staples Retailing	-0.21	0.14	0.94
3020	Food, Bev. & Tobacco	-0.69	5.65	-1.91*
2520	Consum. Durables & Apparel	-0.69	1.27	3.67
5020	Media & Entertainment	-0.85	3.85	-1.48
5010	Communication Services	-0.94	1.32	-0.77
5510	Utilities	-1.08	7.60	-2.43*
3520	Pharma., Biotech., & Life Sc.	-1.19	5.05	-1.84**
3510	Health Care Equip. & Serv.	-1.78	-10.06	-1.11
4030	Insurance	-1.90	-4.13	-2.13
2020	Commercial & Prof. Serv.	-2.20	-10.89	-3.52
6010	Real Estate	-2.72**	-4.15	-3.60**
2550	Retailing	-3.52**	5.84	-6.37***

Note: Industry-specific climate quantity betas as in Equation (3). The coefficients are estimated based on pooled data from 2010 to 2019 inclusive. Industries are sorted by the "Pooled Shock". The standard errors are clustered at the fund level. Appendix Table A.4 reports the t-statistics of each climate quantity beta and the P-values of the joint significance tests. Significance levels: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

ultimate test for whether the estimated climate quantity betas meaningfully correspond to investors' perceived industry-level climate risk exposures is the ability of the portfolios built on these betas to hedge climate news. Before moving to that assessment, the remainder of this section explores in greater detail the stability and interpretation of the industry betas.[21]

2.4.2 Industry-Specific Quantity Betas: Possible Economic Mechanisms

Since the precise ranking of industries by their quantity beta is noisy, it is important to not over-interpret differences between industries with similar estimates for climate quantity

[21] Table 3 shows different industries' climate quantity betas estimated over the full horizon between 2010 and 2019. However, the climate risk exposures of an industry are not guaranteed to be stable. For example, the focus of government policies changes frequently, which could affect industries differentially due to shifting tax and subsidy policies. Similarly, industries' exposures might change due to strategic adaptation along the transition path. For example, many traditional fossil fuel firms now have substantial investments in renewable energies (van Benthem et al. 2022), reducing their transition risk exposure. The quantity-based approach can quickly learn about changes in underlying risk exposures since each period delivers multiple data points from the cross-sectional information. To explore such changes, Appendix Figure A.3 shows the industry ranking based on data for each (rolling) five-year window between 2010 to 2019; low numbers correspond to the biggest "long" positions in the implied hedge portfolio. Despite the presence of sizable estimation noise, there is quite some stability in rankings over time, though there are notable exceptions.

Table 4: Across-Shock Correlations of Industry-Specific Climate Quantity Betas

Panel A: Pearson Industry Climate Beta Correlations

	Disclosure Shock	Heat Shock	Pooled Shock
Disclosure Shock	1.00		
Heat Shock	0.20	1.00	
Pooled Shock	0.69	0.74	1.00

Panel B: Spearman (Rank) Industry Climate Beta Correlations

	Disclosure Shock	Heat Shock	Pooled Shock
Disclosure Shock	1.00		
Heat Shock	0.12	1.00	
Pooled Shock	0.51	0.81	1.00

Note: Panel A shows the Pearson correlations among the industry-specific climate quantity betas. Similarly, Panel B shows the Spearman *rank* correlation among the industry-specific climate quantity betas. The coefficients are based on estimating Equation 3 using pooled data from 2010 to 2019 inclusive.

betas. Nevertheless, it can be instructive to try and understand what economic forces might drive mutual fund managers to buy and sell some industries in response to climate belief shocks. We discuss some of the top ("Semiconductors and Semiconductor Equipment" and "Automobiles and Components") and bottom sectors ("Real Estate", "Commercial and Professional Services", and "Retail") by estimated beta, and also discuss "Energy" as a particularly interesting industry.

Semiconductors & Semiconductor Equipment. Table 3 suggests that the semiconductor sector is positively exposed to climate risk realizations. There are several mechanisms that can explain such exposures. First, many firms in this sector produce solar panels, which are generally based on silicone, a semiconductor material. Both regulatory and technological transition risk realizations to avoid carbon emissions thus benefit this sector directly:

> [O]ur module technology displaces up to 98% of greenhouse gas emissions and other air pollutants when replacing traditional forms of energy generation. [...] Other technological developments in the industry, such as the advancement of energy storage capabilities, have further enhanced the prospects of solar energy as an alternative to traditional forms of energy generation. [First Solar, 2020Q2]

Semiconductors and chips are also an integral component of other applications that contribute to (and benefit from) the green transition, from the development of smart grid systems, to electric vehicles, to smart homes and buildings.

> We also focus on increasing our [carbon] "handprint"—the ways in which Intel technologies can help others reduce their footprints, including Internet of Things solutions that enable intelligence in machines, buildings, supply chains, and factories, and make electrical grids smarter, safer, and more efficient. [Intel, 2019Q4]

Automobiles & Components. Some of the largest positive climate quantity betas are found among firms in the automotive sector, suggesting that these firms would benefit from climate risk realizations. Since cars produce a large share of current carbon emissions (Ritchie et al. 2020), this might at first appear counter-intuitive. However, it is important to realize that the auto sector is at the forefront of the technological transition towards a green economy, with electric vehicles playing an important role in plans to decarbonize the economy. Indeed, shifting consumer preferences and increasing tax benefits for producers and consumers of electric vehicles have the potential to lead to a faster-than-usual turnover of the current vehicle fleet, raising the sales and profits of automotive firms. In a recent 10-k filing, Tesla described the effects of such programs as follows:[22]

> We and our customers currently benefit from certain government and economic incentives supporting the development and adoption of electric vehicles. In the U.S. and abroad, such incentives include tax credits or rebates that [...] allow us to lower our costs and encourage customers to buy our products. [Tesla, 2019Q4]

Importantly, these subsidies will not only raise the sales of firms exclusively focused on electric cars, such as Tesla, but also of incumbents such as Ford and General Motors, which have dramatically expanded their electric vehicle production capacities. This sentiment is reflected in numerous equity analyst reports that we reviewed, 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", and in 10-K reports such as:

> In 2020, we announced the commitment of $27 billion in investments in electric and autonomous vehicle technologies through 2025, with plans to launch 30 new electric vehicle models globally in that timeframe. [General Motors, 2020Q4]

Retailing. The retail sector is negatively exposed to physical climate risk. The escalating frequency and severity of natural hazards have the potential to devastate physical assets such as factories and warehouses, and disrupt supply chains and distribution networks. As a result, retail businesses may face increased operational disruptions, reduced inventory availability, and increased costs associated with damage repairs and recovery efforts.

> Natural disasters, such as hurricanes and tropical storms, fires, floods, tornadoes, and earthquakes; unseasonable, or unexpected or extreme weather conditions; or similar disruptions and catastrophic events can [. . .] also disrupt or disable operations of stores, support centers, and portions of our supply chain and distribution network, including causing reductions in the availability of inventory and disruption of utility services. [Lowe's Companies Inc, 2020Q1]

[22]For example, Tesla benefits from the California Alternative Energy and Advanced Transportation Financing Authority Tax Incentive that provides multi-year sales tax exclusions on purchases of manufacturing equipment and Nevada Tax Incentives that provide abatements for sales, use, real property, personal property and employer excise taxes, discounts to the base tariff energy rates and transferable tax credits. Similarly, the Inflation Reduction Act provides a tax credit of up to $7,500 for consumers who purchase electric vehicles.

Commercial & Professional Services. This relatively small sector has substantially negative climate quantity betas. To understand why this sector could be negatively exposed, consider one of the sub-sectors, Environmental & Facilities Services, which has multiple firms focused on waste management services. These entities are susceptible to regulatory risks as a result of more stringent regulations governing air quality, waste management, and water conservation.

> Stricter environmental regulation of air emissions, solid waste handling or combustion, residual ash handling and disposal, and wastewater discharge could materially affect our cash flow and profitability. [Covanta Holding Corporation, 2019Q2]

Other firms negatively exposed to climate risks are those in the Research & Consulting Services subsector that service carbon-intensive sectors:

> Legislation, international protocols, regulation or other restrictions on emissions could also affect our clients, including those who are involved in the exploration, production or refining of fossil fuels, emit greenhouse gases through the combustion of fossil fuels or emit greenhouse gases through the mining, manufacture, utilization or production of materials or goods. Such policy changes could increase the costs of projects for our clients or, in some cases, prevent a project from going forward, thereby potentially reducing the need for our services, which would in turn have a material adverse impact on our business, financial condition, and results of operations. [Jacobs Engineering Group Inc, 2016Q3]

Real Estate. It is perhaps not surprising that investors view real estate, and in particular REITs, as negatively exposed to climate risks. While the most natural exposure for the sector is physical climate risks (since real estate values are inherently tied to geography), there are also transition risks associated with it (for example, regulation, building requirements, preventive costs to mitigate climate exposures, higher energy costs due to cooling buildings).

> We have significant operations and properties in Northwest Florida that could be materially and adversely affected by natural disasters, manmade disasters, severe weather conditions or other significant disruptions. [...] especially our coastal properties, could experience significant, if not catastrophic, damage. Such damage could materially delay sales or lessen demand for our residential or commercial real estate in affected communities and lessen demand for our hospitality operations and leasing operations. [...] Furthermore, an increase in sea levels due to long-term global warming could have a material adverse effect on our coastal properties and forestry business. [The St. Joe Company]

Energy. The energy sector, which is focused on oil & gas exploration and production, displays positive climate quantity betas. This is perhaps surprising, given the sectors' substantial carbon emissions (van Benthem et al. 2022). However, several mechanisms are consistent with a positive exposure of publicly traded firms in this sector, in particular since concerns about "stranded assets" have likely been reflected in prices for some time. First, fear of tighter future regulation can discourage new entry. The increased market power of incumbents thus raises their profits from selling hydrocarbons while renewable alternatives remain unreliable, even if a faster transition might reduce industry-wide profits in a calculation that included potential entrants (see Ryan 2012, Elliott 2022, Magolin & Santino 2022, Acharya, Giglio, Pastore, Stroebel & Tan 2022). Second, large energy companies play an important role in innovation in the clean energy space (see Pickl 2019, Cohen et al. 2020).

Summary. While the individual quantity-beta estimates for each industry are certainly noisy, many of the estimates can be supported by plausible narratives about the various industries' climate risk exposures. Industries with positive climate quantity betas have business models that should benefit from the decarbonization of the economy (see Fuchs et al. 2024), while industries with negative climate quantity betas are often substantially hurt by both physical climate risk and regulatory efforts to reduce carbon emissions.

At the same time, while many of the estimated climate risk exposures can potentially be rationalized *ex post*, a ranking based purely on an *ex ante* narrative approach might have looked quite different; for example, it is certainly true that the different members of our research team had a range of ex ante priors about different industries' exposures. The quantity-based approach proposed in this paper removes the need for researchers to take a strong *ex ante* stand on different industries' exposures, while also providing position sizing in addition to position direction for how to include each industry in a climate risk hedge portfolio. Taken together, these findings illustrate the complexity of determining sectors' and firms' exposure to climate risk, and help illustrate how approaches that rely on the 'wisdom of the crowd' help distill high-dimensional and possibly conflicting intuitions and sources of information into an optimal hedge portfolio.

2.4.3 Subsample stability of climate quantity betas

To further probe the precision of our estimates, we study the stability of β^I coefficients estimated across subsamples. Specifically, we split the universe of mutual funds into two equally sized random subsamples, estimate β^I in each subsample, and obtain the Spearman and Pearson correlations across the two sets of estimates. We repeat this 250 times, and report the average correlations in Table 5. Quantity betas estimated across different subsamples are substantially correlated, indicating that despite the sizable noise in each estimate, there is a common signal across the climate risk exposures estimated across the different samples.

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

Climate Shock	Fund Split Spearman	Fund Split Pearson
Heat Shock	0.34	0.30
Disclosure Shock	0.15	0.16
Pooled Shock	0.32	0.26

Note: This table shows the average Spearman (rank) and Pearson correlations of industry-specific climate quantity betas of a fund split robustness test.

3 Quantity-Based Climate Hedge Portfolios

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

3.1 Portfolio Construction and Description

We build our hedge portfolio for each month t by estimating $\hat{\beta}_t^{I,S}$ as described in the previous section, using data from the five years prior to t.[23] We compute excess returns of each of the 24 industries as in Equation 4, where R_t^I is the value-weighted industry return and R_t^f denotes the risk-free rate. We use the estimated $\hat{\beta}_t^{I,S}$ as the portfolio weights (since each component of the portfolio is an excess return, the portfolio is a net-zero investment and we do not require $\sum_i \hat{\beta}_t^{I,S} = 1$). The excess return of the quantity-based hedge portfolio is:

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

Table 6 shows the correlations of monthly returns among the quantity-based hedge portfolios based on the different idiosyncratic climate belief shocks. Stars correspond to significance levels. Given that the idiosyncratic belief shocks are largely uncorrelated (see Table 1), the high correlations in the return series provide additional evidence that our various shocks are picking up a common signal.

Table 6: Portfolio Return Correlations

Panel A: Pearson Portfolio Return Correlations

	Heat Shock	Disclosure Shock	Pooled Shock
Heat Shock	1.00		
Disclosure Shock	0.44***	1.00	
Pooled Shock	0.72***	0.73***	1.00

Note: Monthly return correlations (constructed as in Equation 4) for the period 2015 to 2019 among our quantity-based hedge portfolios. Significance levels: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$.

[23] To ensure that the portfolios are implementable in real time, for each month, we limit the regression sample to reports that were filed with the SEC as of the end of the previous month.

3.2 Climate Hedge Targets

One challenge with designing portfolios that hedge climate risk is that there is no unique way of choosing the relevant 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 threats to certain business models from regulations to reduce emissions. Different risks may be relevant for different investors, and these risks are imperfectly correlated. In addition, climate change is a long-run threat, and we would thus ideally build portfolios that hedge the long-run realizations of climate risk, something difficult to do in practice.

To overcome these challenges, Engle et al. (2020) argue that the objective of hedging long-run realizations of a given climate risk can be achieved by constructing a sequence of short-lived hedges against *news* (one-period innovations in expectations) about future realizations of these risks. Following the initial work in Engle et al. (2020), researchers have developed a number of climate news series capturing a variety of different climate risks.

In this paper, we do not take a stand on which of these news series represents the right hedge target—in part because the right target will vary across investors based on their background risk exposure[24]—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 climate news measures proposed in the recent literature, which we describe below. Each measure is signed such that a larger number corresponds to negative news. We aggregate daily series to the monthly level by taking the average of the daily news series. Building on Engle et al. (2020), we use monthly AR(1) innovations of each climate news series as the hedge targets. For a given climate news series c, we denote these AR(1) innovations in month t as $CC_{c,t}$.[25]

We consider the following series:

Engle et al. (2020). The Wall Street Journal (WSJ) and Crimson Hexagon Negative News (CHNEG) climate news indices created by Engle et al. (2020) were, to our knowledge, the first climate news series used as hedge targets. The first one captures the number of news articles in the WSJ dedicated to climate change (broadly assuming that "no news is good news"), the second one builds upon proprietary news aggregations from Crimson Hexagon combined with sentiment analysis that allows the separation of good news and bad news. Both indices capture a mix of news about physical and transition risks. These news indices are available at a monthly frequency. The WSJ index covers the period of February 1984 to June 2017. The CHNEG index covers the period of July 2008 to May 2018.

Ardia et al. (2021). Ardia et al. (2020) build on the WSJ index of Engle et al. (2020) by

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

[25] We construct hedge targets at the monthly frequency because, for many events, it is hard to pin down the news occurrence to a specific day. For example, news coverage of heat waves and similar natural disasters can stretch over several days or even weeks; similarly, news coverage can sometimes predate policy announcements by writing in anticipation of international summits.

including additional media outlets and differentiating between positive and negative news. Their daily Media Climate Change Concerns index is available January 2003 to June 2018.

Faccini et al. (2021). We include four of Faccini et al. (2021)'s climate news indices: international climate summits, global warming, natural disasters, and narrative indices. The international climate summits, global warming, and natural disasters indices measure news coverage of the respective topics; the narrative index is constructed by manually reading and classifying 3,500 articles. The international climate summits and narrative indices capture news about transition risk, while the global warming and natural disasters indices are more likely to capture news about physical risk (though these risk categorizations are not always easy to separate, as bad news about realizations of physical risks may make subsequent regulation more likely). These news measures are available at the daily frequency between January 2000 and November 2019.

National Google searches. This climate news series is the national Google search interest for "climate change", capturing attention paid to climate change and its associated risks by the general population. This monthly 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. Therefore, we include the AR(1) innovation of nationwide monthly maximum temperature, controlling for the month fixed effects, as a climate news series.

Appendix Table A.6 shows the correlations of the climate news series as well as the correlations of the AR(1) innovations. While most measures are positively correlated, many of the correlations are somewhat small in magnitude (and some are actually negative), highlighting the fact that these measures capture different aspects of climate risk.

3.3 Alternative Approaches to Building Hedge Portfolios

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

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

We build several portfolios using such a narrative-based approach. Our first narrative-based portfolio takes positions in all U.S.-listed stocks covered by the Sustainalytics ESG scores: the portfolio's position in each stock is the stock's ESG score percentile in each period, minus 50. For example, the portfolio takes a long position of 50 in the company with the highest ESG score and a short position of -50 in the company with the lowest score in each month. Stocks with the median ESG score are not held. This portfolio corresponds to the Sustainalytics hedging portfolio proposed in Engle et al. (2020).

Our second narrative-based strategy uses industries to take a directional view. We build portfolios using four ETFs: the Invesco Global Clean Energy ETF (Ticker: PBD), which invests in firms focused on the development of cleaner energy and conservation; the iShares Global Clean Energy ETF (Ticker: ICLN), which tracks the S&P Global Clean Energy Index; 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; and the iShares U.S. Energy ETF (Ticker: IYE), which tracks the Dow Jones US Energy Sector Index. Based on the narrative that transition risk realizations would benefit renewable energy firms at the expense of incumbant fossil-fuel based firms, we propose that realizations of climate change news should increase PBD and ICLN's returns and decrease XLE and IYE's returns. Therefore, the hedge portfolio would go long PBD and ICLN and short XLE and IYE.

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

Our fourth narrative-based portfolio is constructed as a carbon-emissions-sorted portfolio, which involves taking positions in all U.S.-listed stocks covered by the Trucost carbon emission data. We calculate the sum of Scope 1 and 2 greenhouse gas emissions divided by a company's revenue to measure a firm's carbon intensity. For each period, we rank companies based on their carbon intensity and assign the portfolio's position in each firm as the percentile of the firm's carbon intensity within the distribution, and remove that from 50. For instance, in a given month, the portfolio takes a long position of 50 in the company with the lowest carbon intensity, while taking a short position of -50 in the company with the highest carbon intensity.

The final narrative approach attempts to measure a firm's climate risk exposure based on textual and sentiment analysis of its 10-K reports (see Baz et al. 2023). First, we identify climate-related sentences in 10-K statements using regular expression searches, using the same climate dictionary that we use to construct the disclosure-based measure based on shareholder reports. Given that climate mentions can encompass both risks and opportunities, we employ sentiment analysis to separate such mentions. Specifically, we use the BERT model to classify each climate-related sentence into positive, neutral, and negative sentiments. For each firm-year, we count the number of positive and negative sentences and compute a 10K-Climate-Negative Score as the number of negative climate sentences minus the number of positive climate sentences (see Giglio, Kuchler, Stroebel & Zeng 2023, for a related approach). Each year, companies are ranked according to their 10K-Climate-Negative Scores, and the portfolio's position in each firm is assigned as 50 minus the percentile of the firm's score. For example, the portfolios take a long position of 50 in the company with the lowest score and short a position of -50 in the company with the highest score.

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

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

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

where $CC_{c,t}$ denotes the (mean zero) climate hedge target of type c in month t, w_c is a vector of N portfolio weights, and R_t is a vector of demeaned excess returns. The portfolio weights are estimated each month using five-year rolling windows.

We consider different sets of excess returns (base assets) to build mimicking portfolios. First, we use the market alone (the SPY ETF). A mimicking portfolio built using only one asset is effectively equivalent to studying whether a long or a short position in that asset was historically correlated with climate risk. Second, we use the three Fama-French factors (Market, SMB, and HML).[26] Third, we use two of the ETFs described above, PBD and XLE, in combination with the Fama-French factors. Fourth, we add to the Fama-French factors the excess returns of the 24 GICS industry portfolios. Fifth, we use the 207 firm characteristics obtained from Chen & Zimmermann (2022) to construct a 'factor zoo' portfolio using the returns of portfolios sorted on those 207 characteristics. Given the short time series available for estimation, we regularize the estimation for the two latter hedge portfolios using LASSO, choosing the tuning parameter by cross-validation, to minimize in-sample overfitting.

3.4 Hedging Climate Shocks: Evaluation of Hedge Portfolios

In this section, we evaluate the hedge performance of the different proposed portfolios. For the quantity-based and mimicking portfolio approaches, for every month in our testing period of 2015-2019, we construct the portfolios as described above using five-year rolling windows of data. The portfolio weights for the narrative hedge portfolios do not vary over time (e.g., the PBD-based portfolio is always 100% long PBD). We focus on the post-2010 period to

[26] By construction, this portfolio, like the one based on the market alone, wouldn't be able to hedge the component of climate risk that is not spanned by the factors; therefore, it cannot provide investors with a way to hedge the part of climate risk that is orthogonal to them, which limits its usefulness as a hedge portfolio. We include it here for illustrative purposes, and because it does contain information about how much climate risk is already reflected in the market and other aggregate factors.

train our models, as investors likely paid very little attention to climate risks before 2010. As a result, we do not expect information on prices and quantities from before 2010 to be useful in building hedge portfolios today. As a criterion to evaluate the various hedging approaches, we compare the out-of-sample correlations between the hedging portfolio returns and the AR(1) innovations to the various climate news series in the same month, $CC_{c,t}$.[27]

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

Table 7: Climate Hedge Performance of Various Portfolios

		Faccini et al.				Engle et al.		Ardia et al.	Google	Temp.
	Avg.	IntSummit	GlobWarm	NatDis	Narrative	WSJ	CHNEG	MCCC	National	National
Pooled Shock	0.18***	0.23	0.26**	0.12	0.02	0.08	0.13	0.31*	0.38***	0.11
Heat Shock	0.17***	0.38***	0.18	0.05	0.06	0.05	0.10	0.25**	0.34*	0.14
Disclosure Shock	0.11**	0.10	0.15	0.26**	0.01	0.06	0.06	0.23	0.06	0.09
Emission Portfolio	0.08*	-0.03	0.13	0.06	-0.03	0.20	0.31**	0.18	0.05	-0.15
Long PBD ETF	0.07	0.06	0.09	0.19	0.04	-0.02	-0.03	0.23	0.02	0.05
Short Stranded Asset	0.02	-0.06	0.05	0.25***	0.12	-0.01	0.14	-0.11	0.04	-0.20
Long ICLN ETF	0.01	0.02	-0.02	0.22**	-0.06	-0.06	-0.11	0.12	-0.04	0.00
Short IYE ETF	-0.08*	-0.08	-0.16	-0.10	0.13	-0.12	0.06	-0.27**	-0.05	-0.09
Short XLE ETF	-0.08*	-0.09	-0.16	-0.12	0.14	-0.12	0.05	-0.27**	-0.05	-0.07
10-K Negative Portfolio	-0.09**	0.04	-0.01	0.02	0.08	-0.16	-0.17	-0.26**	-0.30***	-0.09
Sustainalytics Portfolio	-0.10**	0.13	-0.08	0.06	0.07	-0.25**	-0.23	-0.20	-0.27**	-0.14
Lasso: All Industry+FF	-0.02	0.15*	0.08	0.14*	0.04	0.01	-0.06	-0.36***	-0.19*	0.00
Lasso: Factor Zoo	-0.03	0.09	-0.03	-0.06	-0.12	-0.10	0.00	-0.04	0.00	0.00
Reg: ETFs+FF	-0.03	0.01	0.05	0.05	-0.04	0.11	-0.17	-0.32**	0.03	-0.02
Reg: FF 3-Factors	-0.04	-0.02	0.06	0.06	-0.11	0.20	-0.27**	-0.26	0.00	0.01
Reg: SPY ETF	-0.09**	-0.09	-0.03	-0.12	-0.08	-0.17	-0.14	-0.13	-0.00	-0.07

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 different climate news series. All climate news series are coded such that high numbers indicate negative climate news. Therefore, positive correlation coefficients indicate successful hedges. While the narrative and quantity portfolios stay identical along the rows, the mimicking portfolios in each cell was specifically trained on the respective climate news series. Significance levels: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. P-values are calculated using a bootstrap of 1000 iterations.

Table 7 reports out-of-sample correlations between the hedge portfolio returns and innovations in each of the news series. Each row in the table corresponds to a different way of forming a hedge portfolio, whereas each column corresponds to a different climate news series (the first column reports the average correlation across all news series). All climate news series are coded such that high numbers are indicative of negative climate news. Therefore, positive correlations imply successful hedges. The table also reports stars representing statistical significance.[28] Similar information is displayed in Figure 2, which reports out-of-sample correlations on the horizontal axis, and has one row for each hedge portfolio. Each point in the dot plot is the out-of-sample correlation coefficient of a hedge portfolio return with one of the climate news series. The different colors represent the different climate news series (hedge targets) described above. The red rhombus shows the average among all correlations, and portfolios are sorted top-to-bottom by this value.

Figure 2: Climate Hedge Performance of Various Portfolios

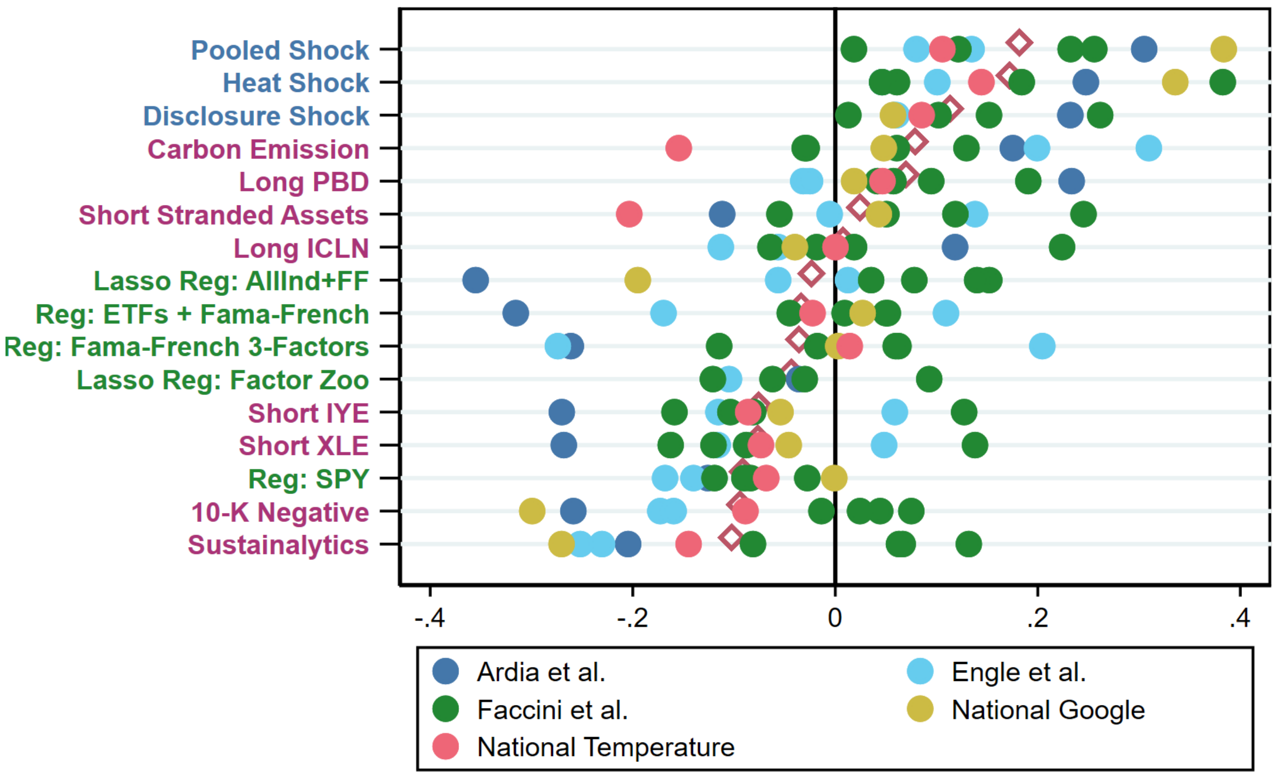

Note: Dot plot of monthly out-of-sample return correlations for various climate hedge portfolios with various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Different colors represent different groups of climate news series. The red rhombus shows the unweighted average among all correlations, and portfolios are sorted top-to-bottom by this value.

[28] We use bootstrap standard errors to compute the significance; the bootstrap resamples monthly observations with replacement from the 2015-2019 period—the period over which the correlations are calculated—sampling jointly the climate news innovations and the returns of all hedging portfolios. Table A.7 reports the bootstrap standard errors for each correlation.

The first three rows of Table 7 show the hedging performance of the quantity-based climate hedge portfolios. These portfolios tend to produce relatively high out-of-sample correlations for a large variety of climate news series (the blue rows of Figure 2 show the same results).[29] The heat-based, disclosure-based, and pooled portfolio returns correlate positively with *essentially all* climate news innovations. Only some of the correlations are individually statistically significant, but the average correlation is strongly significant for both the heat and disclosure hedge portfolios, as well as the pooled portfolio. These findings suggest that our various quantity-based portfolios perform well in terms of hedging a range of climate risks, spanning both physical and transition risks. Given that our quantity-based approaches are not tailored to hedge specific climate targets, their good performance against a variety of targets suggests that they are providing a hedge against some common component of climate risks that is shared by the measures we consider.

Rows 4-11 of Table 7 (and red rows of Figure 2) show the performance of the narrative-based portfolios. The main advantage of these portfolios is that they do not require estimating the portfolio weights from historical data, since the direction of the trades is based on ex-ante information and beliefs. For example, the narrative portfolio that has a short position in the stranded asset portfolio is motivated by the fact that this portfolio is dominated by polluting companies in the coal sector. A similar logic applies to the narrative portfolio that shorts high-carbon-emissions firms. The narrative portfolio featuring a long position in the ETF PBD is motivated by a belief that a clean energy fund should gain upon transition risk realizations. All three of these portfolios have substantially positive average hedge performances, though not quite at the same level as the quantity-based portfolios. More broadly, however, the hedging performance of narrative portfolios is mixed, with some of the narrative portfolios yielding among the most negative correlations with climate shocks. The uneven hedging performance highlights just how difficult it is to predict, based only on economic intuition alone, which stocks will gain or lose in response to climate shocks.

The remaining rows of Table 7 report the hedging performance of mimicking portfolios based on aggregate time-series information (see also the green rows in Figure 2). The performance of these portfolios varies substantially across climate news series, but is poor on average. For example, the portfolio built using the three Fama-French factors has a relatively high correlation of 0.2 with the WSJ index from Engle et al. (2020). 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 the MCCC index from Ardia et al. (2020). All of the other correlations are close to zero. Note that the mimicking portfolios have a relatively weak hedging performance despite the fact that they are estimated separately for each hedge

[29] When evaluating this out-of-sample hedging performance, it is worth keeping in mind that hedging macroeconomic shocks using stocks is generally difficult. As a reference, when building mimicking portfolios for macro risks using a regularized projection method, Giglio & Xiu (2021) report in-sample R^2s of 2.25% for industrial production growth and 4.07% for consumption growth at the monthly frequency, corresponding to *in-sample* correlations between the target and the hedging portfolio return of 0.15 and 0.20 respectively. Appendix Table A.9 reports the *in-sample* R^2 and adjusted-R^2 for each ex post maximally correlated portfolio of industries. While the in-sample R^2s are high, this is due to the large number of covariates relative to the number of observations; thus, the adjusted-R^2s are, on average, approximately zero.

target, giving them additional flexibility compared to the other methodologies (which instead build a single hedge portfolio for all climate news series).

To further compare the hedging performance of the different approaches, we report in Table A.8 the p-value for a test of the difference between the hedging correlation obtained by the pooled portfolio and each other hedging portfolio, calculated using bootstrapped standard errors as above. The table shows that the two quantity-based portfolios (heat and disclosure-based) produce correlations that are statistically indistinguishable from the pooled shock. But in almost all cases, the pooled quantity portfolio yields significantly higher correlations than the narrative and mimicking portfolios (the exceptions are the few relatively well performing narrative portfolios).

Overall, the results show that our quantity-based approach to forming hedge portfolios consistently delivers the best out-of-sample climate hedging performance. Among the alternative approaches, with only little historical data available for periods when climate risk was potentially priced, the mimicking portfolio approaches do not deliver successful climate hedges. In contrast, the narrative-based approach to building hedging portfolios is potentially promising—for example, PBD has positive correlations with all but two 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. Consistent with this, other portfolios using the narrative approach do not perform well across measures. For example, the short XLE position—an intuitive trade ex ante—does not perform well as a climate hedge.

3.5 Robustness to Specification Choices

In this section we study the robustness of our results with respect to a variety of choices that affect the construction of the portfolios and the implementation of the analysis.

Construction of the portfolios. We begin by varying several of the choices made in the construction of the quantity-based portfolios. Figure 3 shows average out-of-sample correlations similar to Figure 2 for variations of the pooled quantity-based portfolio. Appendix Figure A.4 shows similar robustness checks for hedge portfolios built separately using the heat and disclosure-based idiosyncratic belief shocks. We consider the following variations to our baseline portfolio construction choices:

(i) Add the interaction of time fixed effects with fund type or fund family fixed effects in the regression in Equation 3 ("Fund Type $\times$ Time FE"; "Fund Family $\times$ Time FE");

(ii) Measure changes in investors' industry-level portfolio holdings using current prices, defined as *TotalChanges* in footnote 16. This allows for portfolio changes to be driven by price changes in addition to active trading ("Total Changes");

(iii) Change how we handle extreme changes in investors' industry-level portfolio holdings from the the baseline procedure of winsorizing at the 1% level ("1% Trimming"; "0.5% Winsorizing"; "2% Winsorizing"; "No Winsorizing");

(iv) Do not weight changes in investors' industry-level portfolio holdings by the industries' relative market size ("No Industry Weighting");

(v) Change the relevant universe of funds to be defined using only CRSP or Thomson Reuters IOC data ("CRSP Equity Domestic"; "IOC in 2, 3, or 4"), and to exclude ESG funds that are likely to have a large climate risk focus even before receiving idiosyncratic belief shocks ("No ESG");

(vi) Only keep funds where all advisers reside in the same county, and for which the geographic allocation of heat shocks is thus easier ("Unique County"), though at the cost of having fewer shocks; only keep funds that have been hit by at least one of the idiosyncratic belief shocks ("Shocks hit fund ≥ 1");

(vii) Measure extreme local heat shocks and investor locations at the commuting zone level instead of the county level ("Commuting Zone Shocks");

(viii) Use three-year and one-year rolling windows of trading activity to identify industry-level climate quantity betas, instead of five-year rolling windows as in the baseline ("Three-Year Rolling Window"; "One-Year Rolling Window");

(ix) Orthogonalize each hedge portfolio with respect to the market factor ("Orthogonal to MktRF") or the Fama-French market, size and value factors ("Orthogonal to MktRf, HML, and SMB");

(x) For news indices at the daily frequency, first calculate daily AR(1) innovations and then aggregate to the monthly level instead of the baseline approach of aggregating to the monthly level before taking innovations ("Aggregate daily innovations");

(xi) The SHELDUS per-capita variables we use to construct the heat shock are rounded to the nearest 10^{-5}, which means that a heat shock requires at least 1 fatality or injury per 200k residents; we re-construct the heat shock using alternative cutoffs of one death/injury per 100k residents or one per 300k residents (">1 per 100k Inj/Fat", >1 per 300k Inj/Fat");

Most of the variation in how we construct hedge portfolios has little effect on the overall hedge performance, and multiple changes appear to even improve the hedge performance relative to our baseline data construction choices. One notable exception is excessively shortening the estimation period for the quantity coefficients (regression 3); moving from five-year rolling windows to one-year rolling windows reduces by 80% the amount of quantity data that our approach can use to identify the industries' quantity betas. This loss in data appears to outweigh the potential gains from being able to better detect time-varying exposures with one-year rolling windows. Performance also declines, though by much less, when going from a 5-year window to a 3-year window to estimate climate quantity betas.

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

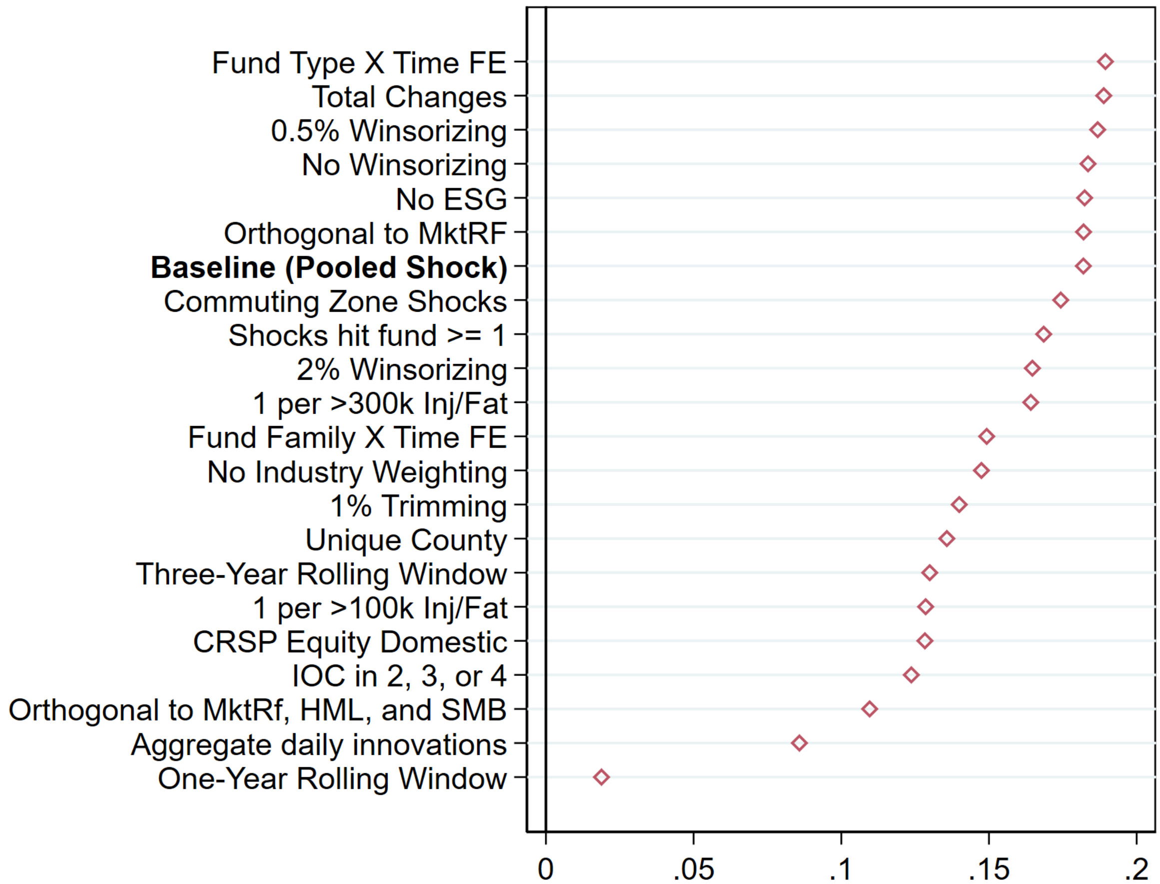

Note: Dot plots of monthly out-of-sample return correlations for variations of the quantity-based portfolio constructed using all idiosyncratic belief shocks. Each row corresponds to a different way to build the hedge portfolio, described in the text. The red rhombus shows the unweighted average among all correlations, and portfolios are sorted top-to-bottom by this value.

Definition of heat shock. As described above, there are several plausible ways to construct local extreme heat shocks. Our baseline construction identifies local extreme heat shocks that produce fatalities and injuries; here we explore two alternative definitions.

The first alternative definition uses temperature data from the National Oceanic and Atmospheric Administration (NOAA) to identify extreme temperatures. We flag county-months with a maximum temperature of at least 10 degrees Fahrenheit above the county's five-year historical average maximum for that month. We enforce that this maximum temperature is above 90 degrees Fahrenheit, the threshold for "extreme caution" by the U.S. National Weather Service. About 0.70% of county-months have such an extreme heat event.

The second alternative heat shock is based on crop indemnity payments. The underlying data are collected by the U.S. Department of Agriculture, and we use a version maintained by SHELDUS.[30] We normalize the crop indemnity payments by the number of acres reported as being planted adjusted by the insured's share in the commodity. An extreme heat event

[30] Crop indemnity payments are insurance payments to farmers, which are paid when external disruptions

is identified when the monthly normalized heat-related crop indemnity payments in a given county exceed the 99.9th percentile across all county-months in the preceding 10 years; about 0.08% of county-months between 2010 and 2019 had such an event.

The correlation between the three versions of heat shocks in the panel is low (they are all positive but less than 25% in magnitude). The reason is that different types of shocks tend to affect different areas: for example, crop indemnity payments tend to affect rural areas.

Appendix Figure A.5 shows the out-of-sample correlations between the various heat-based quantity portfolios and the climate news series; the figure also includes the baseline quantity portfolios for comparison. The figure shows that all heat-based quantity portfolios are positively correlated with the various climate news shocks on average, with average correlations ranging from 0.08 to 0.18. The heat-based quantity portfolios outperform almost all of the narrative and mimicking portfolios.

Granularity of Quantity Base Assets. In our baseline analysis, we focus on industry-level changes in holdings because the large number of stocks and the relatively sparse holding matrix would imply a large estimation error in each stock's climate risk exposure (climate quantity beta). To explore how our results depend on the granularity of the base assets for the quantity-based approach data, Appendix Figure A.6 shows the hedging performance of the pooled portfolio constructed with (i) 11 GICS 2-digit industries, 72 GICS 6-digit industries and 175 GICS 8-digit industries; (ii) the top 10 and top 25 largest stocks within each industry measured by the market cap in December 2019; and (iii) all stocks held by Dow Jones 30 Index, NASDAQ 100 Index, and S&P 500 Index according to the index constitution in December 2019 downloaded from Bloomberg. In portfolios formed with individual stocks, we represent the active changes using indicators. Specifically, we assign a value of 1 if the active change in a given shock is positive and -1 if the active change is negative. We find that, while the hedge portfolios continue to work relatively well using a lower level of aggregation, the performance does deteriorate as more assets are used, reflecting the increased noise in the quantity-beta estimates. Hedge portfolios based on individual stocks are not able to provide stable out-of-sample hedges against climate risk realizations.

A placebo test: hedging before 2010. Finally, as mentioned in Section 3.4, we would not expect any of our hedging approaches to produce successful hedges before 2010, when investors likely did not price in climate risks. We test this conjecture by computing the hedge performance of the various approaches for the period 2000-2010, and report the results in Appendix Figure A.7; among the quantity-based approaches, we only include the heat-based one, since it is possible expand it backwards in time without substantial computational cost. The figure confirms that the hedging ability is substantially lower for the pre-2010 period, consistent with findings from Acharya, Johnson, Sundaresan & Tomunen (2022) that physical climate risks only started being priced in municipal bond markets after 2010.

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.

4 Tradeoffs Between Different Hedging Approaches

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

Short time series and time-varying exposures. The mimicking portfolio methodology takes a purely statistical approach to constructing hedge portfolios. It requires little input from the researcher beyond the choice of the base assets used in the projection, and instead relies on the time series information (the historical covariance between the hedge target and asset returns) to choose the portfolio weights. This approach works well as long as the time series is sufficiently long ($T \to \infty$) and asset risk exposures are stable over time: in that case, a sufficiently large set of base assets (i.e., as $N \to \infty$) asymptotically generates the optimal hedge portfolio (see Giglio & Xiu 2021). The mimicking portfolio approach suffers particularly when T is small, as is the case with newly emerging risks such as climate risk, since a small T lead to noisy estimates of the covariance of prices with the hedging target. For similar reasons, the mimicking portfolio approach also struggles when asset exposures are time-varying, for example because firms' strategic decisions affect their risk exposures over time.[31] In fact, one can think of a change in exposure over time as having the same effect as reducing the time periods T that can be used to learn about the new exposure.

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

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

Data Requirements. While the quantity-based approach has important benefits relative to the mimicking portfolio approach, it has stronger data requirements. First, the quantity-

[31] 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. Over time, Orsted's exposure to transition risk would thus have shifted from negative to positive.

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

The quantity-based approach also requires the researcher to observe portfolio holdings or trading data. Such quantity data is generally less available than price data, in particular for assets other than equities. Expanding the base assets to include, for example, commodities or derivatives on emission allowances in the European Emissions Trading System (ETS) would require researchers to also observe investor holdings in these assets. While various regulators have access to portfolio holdings data that would allow them to implement our approach with a wider range of base assets, those data sets are often not publicly accessible.

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

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

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

5 Hedging Macro Factors

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

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

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

$$\Delta Unemp_{t,t-3,c} = \theta_c \Delta Unemp_{t-3,t-6,c} + \delta_{m(t)} + \epsilon_{t,c}, \tag{5}$$

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

$$\Delta Unemp_{t,t-1} = \theta_{national} \Delta Unemp_{t-1,t-2} + \delta_{m(t)} + \epsilon_t, \tag{6}$$

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

$$\Delta Log(ZHVI_{t,t-3,c}) = \theta_c \Delta Log(ZHVI_{t-3,t-6,c}) + \delta_{m(t)} + \epsilon_{t,c}. \tag{7}$$

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

$$\Delta Log(ZHVI_{t,t-1}) = \theta_{national} \Delta Log(ZHVI_{t-1,t-2}) + \delta_{m(t)} + \epsilon_t. \tag{8}$$

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

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

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

We then use the Bidirectional Encoder Representations from Transformers (BERT) model, a state-of-art language model proposed by Devlin et al. (2018), to classify each of these sentences to determine whether it expresses a positive or negative sentiment about the labor or the housing market.[32] The pre-trained model we used is developed by Araci (2019). Positive sentences get a score of "1", and negative sentences get a score of "-1". Table A.2 presents examples of sentences relating to these risks, alongside their BERT sentiment classification.

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

Hedge Portfolio Construction. The construction of the hedge portfolios closely follows the approach described in Section 2. We first estimate the regression in equation 3 with

[32] While it is theoretically possible to use GPT-4 for classification as above, this imposes a significant financial cost, so for this section, we employ the (free) BERT model.

different fund-specific measures of idiosyncratic belief changes to obtain industry-specific quantity betas for each of the two macro risks. For the location-based belief shocks, we use the county-level innovations from the AR(1) processes estimated in equations 5 and 8 as proxies for each fund's idiosyncratic belief shock $S_{f,t}$.

For the disclosure-based measures of changes in unemployment beliefs, we assign $S_{f,t}$ a value of "1" if the fund's overall unemployment sentiment deteriorates between reports (e.g., when it changes from positive to negative or neutral and when it changes from neutral to negative), a value of "0" if the sentiment is unchanged, and a value of "-1" if the fund's sentiment about unemployment rates improves (e.g., when it changes from negative to neutral or positive and when it changes from neutral to positive). Positive quantity betas, which correspond to long positions in the hedge portfolio, therefore describe industries that investors disproportionately buy when they become more pessimistic about national unemployment rates. We would thus expect the hedge portfolio to outperform when the national unemployment rate increases unexpectedly. Positive innovations in house price growth constitute *good news*.[33] Therefore, for the disclosure-based measure of changes in house price beliefs, we assign $S_{f,t}$ a value of "1" if the fund's expressed sentiment regarding the housing market improves between reports, and a value of "-1" when the sentiment deteriorates.

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

Hedge Performance. Table 8 shows the monthly out-of-sample correlations of various hedge portfolios with AR(1) innovations in the national unemployment rate and the national growth rate of house prices. To align with the approach to hedging climate news, we estimate the out-of-sample performance using data from 2015 to 2019, and estimate the hedge portfolios using five-year rolling windows of holdings and price data.

Table 8: Macro Hedge Performance

	Hedge Target	
	Growth in House Prices	Δ Unemployment Rate
Mimicking Portfolio Approaches		
Reg: Fama-French Three-Factors	**0.11**	-0.03
Reg: SPY	**0.13**	-0.01
Lasso Reg: All-Industries + Fama-French	**0.01**	-0.13
Quantity-based Approaches		
Quantity: Local Shocks	**0.18**	**0.20**
Quantity: Disclosure	**0.14**	**0.10**

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

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

The mimicking portfolio approach produces hedge portfolios with positive out-of-sample correlations with house price shocks, and zero or even negative correlations with the unemployment shock. The quantity-based portfolios (both the ones based on local shocks and those based on disclosures) perform well at hedging the two macro series, in both cases outperforming the best mimicking portfolios. These results highlight the fact that the quantity-based methodology can be applied to a variety of other settings beyond hedging climate risks.

6 Conclusions and Directions for Future Research

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

Despite the initial success of the quantity-based approach, we believe that investors interested in operationalizing this approach can further improve upon the resulting hedge performance by introducing portfolio holdings data from a wider range of investors, including retail investors, and by expanding the set of base assets beyond industry equity portfolios. For example, including positions in commodity or carbon futures may further improve the hedge portfolios' ability to hedge the arrival of aggregate physical or transition risk news. Similarly, a fruitful direction for further work would be to explore whether other severe weather events beyond extreme heat (e.g., hurricanes and wildfires) can also be used as shifters of investors' climate risk beliefs, thus potentially expanding the set of trading activities that can inform the construction of quantity-based hedge portfolios.[34]

Our work has focused on exploring the hedging ability of various portfolios in terms of the correlations of their returns with realizations of climate news. Future work should consider both the expected returns of these hedge portfolios as well as on the average portfolio volatilities, both of which are informative about the overall costs of hedging climate risks.

The focus of our application is to allow for an optimal allocation of climate risks 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 large changes in the cost of capital are, and the extent to which they can reduce overall emissions (see Pedersen et al. 2021, Pástor et al. 2021, Goldstein et al. 2022, Berk & van Binsbergen 2021, Bolton & Kacperczyk 2021*b*). More generally, even the direction of the overall amount of climate risk is not clear. For example, hedging climate risk decreases its economic cost and could lessen the incentives to mitigate these risks. Ultimately, the effect of individual investors' hedging of climate risk on the aggregate amount of risk is ambiguous, and understanding the quantitative importance of the various channels is an important area for research.

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

Lastly, while our focus on this paper has been on hedging climate risks, investors are also increasingly focusing on other emerging risks, such as biodiversity risks (Giglio, Kuchler, Stroebel & Zeng 2023, Giglio et al. 2024), cybersecurity risks (Florackis et al. 2023), or pandemic risks (Gormsen & Koijen 2023). An interesting avenue for future work would be to explore the extent to which our new quantity-based approach can allow investors to also improve investors' ability to hedge these and other risks.

References

Acharya, V., Berner, R., Engle, R., Jung, H., Stroebel, J., Zeng, X. & Zhao, Y. (2023), Climate stress testing, Technical report, National Bureau of Economic Research.

Acharya, V. V., Giglio, s., Pastore, S., Stroebel, J. & Tan, J. (2022), 'Climate transition Risk and Energy Prices', *Working Paper, NYU Stern School of Business* .

Acharya, V. V., Johnson, T., Sundaresan, S. & Tomunen, T. (2022), Is physical climate risk priced? evidence from regional variation in exposure to heat stress, Technical report, National Bureau of Economic Research.

Alok, S., Kumar, N. & Wermers, R. (2020), 'Do Fund Managers Misestimate Climatic Disaster Risk', *The Review of Financial Studies* **33**(3), 1146–1183.

Araci, D. (2019), 'Finbert: Financial sentiment analysis with pre-trained language models', *arXiv preprint arXiv:1908.10063* .

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

Armantier, O., Bruine de Bruin, W., Topa, G., van der Klaauw, W. & Zafar, B. (2015), 'Inflation expectations and behavior: Do survey respondents act on their beliefs?', *International Economic Review* **56**(2), 505–536.

Armona, L., Fuster, A. & Zafar, B. (2019), 'Home Price Expectations and Behaviour: Evidence from a Randomized Information Experiment', *Review of Economic Studies* **86**(4), 1371–1410.

Bachmann, R., Berg, T. O. & Sims, E. R. (2015), 'Inflation expectations and readiness to spend: Cross-sectional evidence', *American Economic Journal: Economic Policy* **7**(1), 1–35.

Bailey, M., Cao, R., Kuchler, T. & Stroebel, J. (2018), 'The economic effects of social networks: Evidence from the housing market', *Journal of Political Economy* **126**(6), 2224–2276.

Bailey, M., Dávila, E., Kuchler, T. & Stroebel, J. (2019), 'House price beliefs and mortgage leverage choice', *The Review of Economic Studies* **86**(6), 2403–2452.

Bailey, M., Johnston, D., Koenen, M., Kuchler, T., Russel, D. & Stroebel, J. (2020), 'Social networks shape beliefs and behaviors: Evidence from social distancing during the covid-19 pandemic'.

Bakkensen, L. A. & Barrage, L. (2022), 'Going Underwater? Flood Risk Belief Heterogeneity and Coastal Home Price Dynamics', *The Review of Financial Studies* **35**(8), 3666–3709.

Baldauf, M., Garlappi, L. & Yannelis, C. (2020), 'Does climate change affect real estate prices? only if you believe in it', *The Review of Financial Studies* **33**(3), 1256–1295.

Barnett, J. (2020), 'Global environmental change ii: Political economies of vulnerability to climate change', *Progress in Human Geography* **44**(6), 1172–1184.

Baz, S., Cathcart, L., Michaelides, A. & Zhang, Y. (2023), 'Firm-level climate regulatory exposure', *Available at SSRN 3873886* .

Berk, J. B. & van Binsbergen, J. H. (2016), 'Assessing asset pricing models using revealed preference', *Journal of Financial Economics* **119**(1), 1–23.

Berk, J. & van Binsbergen, J. H. (2021), 'The impact of impact investing', *Available at SSRN 3909166* .

Bernstein, A., Gustafson, M. T. & Lewis, R. (2019), 'Disaster on the horizon: The price effect of sea level rise', *Journal of Financial Economics* **134**(2), 253–272.

Bolton, P. & Kacperczyk, M. (2021*a*), 'Do investors care about carbon risk?', *Journal of Financial Economics* **142**(2), 517–549.

Bolton, P. & Kacperczyk, M. (2021*b*), Global pricing of carbon-transition risk, Technical report, National Bureau of Economic Research.

Busse, M. R., Pope, D. G., Pope, J. C. & Silva-Risso, J. (2015), ' The Psychological Effect of Weather on Car Purchases *', *The Quarterly Journal of Economics* **130**(1), 371–414.

Cashion, P. (2024), 'Calpers' plan to generate alpha from climate investments', https://www.top1000funds.com/2024/07/calpers-plans-to-generate-alpha-from-climate-investments/.

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

Chang, S. (2022), 'Local industry bias in investor behavior: Evidence from mutual funds', *Working Paper* .

Chang, T. Y., Huang, W. & Wang, Y. (2018), 'Something in the Air: Pollution and the Demand for Health Insurance', *The Review of Economic Studies* **85**(3), 1609–1634.

Chen, A. Y. & Zimmermann, T. (2022), 'Open source cross-sectional asset pricing', *Critical Finance Review* **27**(2), 207–264.

Chen, Q., Goldstein, I. & Jiang, W. (2010), 'Payoff complementarities and financial fragility: Evidence from mutual fund outflows', *Journal of Financial Economics* **97**(2), 239–262.

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

Cohen, L., Gurun, U. G. & Nguyen, Q. H. (2020), The esg-innovation disconnect: Evidence from green patenting, Technical report, National Bureau of Economic Research.

Constantino, S. M., Cooperman, A. D., Keohane, R. O. & Weber, E. U. (2022), 'Personal hardship narrows the partisan gap in covid-19 and climate change responses', *Proceedings of the National Academy of Sciences* **119**(46), e2120653119.

D'Acunto, F., Malmendier, U. & Weber, M. (2022), 'What Do the Data Tell Us About Inflation Expectations?', *Working Paper* .

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

Devlin, J., Chang, M.-W., Lee, K. & Toutanova, K. (2018), 'Bert: Pre-training of deep bidirectional transformers for language understanding', *arXiv preprint arXiv:1810.04805* .

Dou, W. W., Kogan, L. & Wu, W. (2022), Common fund flows: Flow hedging and factor pricing, Technical report, National Bureau of Economic Research.

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

Elliott, J. T. (2022), 'Investment, emissions, and reliability in electricity markets', *Working Paper* .

Engle, R. F., Giglio, S., Kelly, B., Lee, H. & 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* .

Florackis, C., Louca, C., Michaely, R. & Weber, M. (2023), 'Cybersecurity risk', *The Review of Financial Studies* **36**(1), 351–407.

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.

Fuchs, M., Stroebel, J. & Terstegge, J. (2024), Carbon vix: Carbon price uncertainty and decarbonization investments, Technical report, Copenhagen Business School.
URL: *https://www.carbonvix.org/project/paper/Carbon$_P$rice$_U$ncertainty.pdf*

Gennaioli, N., Ma, Y. & Shleifer, A. (2016), 'Expectations and investment', *NBER Macroeconomics Annual* **30**(1), 379–431.

Giglio, S., Kelly, B. & Stroebel, J. (2021), 'Climate finance', *Annual Review of Financial Economics* **13**, 15–36.

Giglio, S., Kuchler, T., Stroebel, J. & Wang, O. (2024), The economics of biodiversity loss, Technical report, National Bureau of Economic Research.

Giglio, S., Kuchler, T., Stroebel, J. & Zeng, X. (2023), 'Biodiversity risk', *NBER Working Paper 31137* .

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* **34**(8), 3527–3571.

Giglio, S., Maggiori, M., Stroebel, J., Tan, Z., Utkus, S. & Xu, X. (2023), Four facts about esg beliefs and investor portfolios, Technical report, National Bureau of Economic Research.

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

Giglio, S., Maggiori, M., Stroebel, J. & Utkus, S. (2021*b*), 'The joint dynamics of investor beliefs and trading during the covid-19 crash', *Proceedings of the National Academy of Sciences* **118**(4), e2010316118.

Giglio, S. & Xiu, D. (2021), 'Asset pricing with omitted factors', *Journal of Political Economy* **129**(7), 1947–1990.

Goldsmith-Pinkham, P. S., Gustafson, M., Lewis, R. & Schwert, M. (2021), 'Sea level rise exposure and municipal bond yields', *Working Paper* .

Goldstein, I., Kopytov, A., Shen, L. & Xiang, H. (2022), On esg investing: Heterogeneous preferences, information, and asset prices, Technical report, National Bureau of Economic Research.

Gormsen, N. J. & Koijen, R. S. (2023), 'Financial markets and the covid-19 pandemic', *Annual Review of Financial Economics* **15**(1), 69–89.

Grinblatt, M. & Titman, S. (1989), 'Mutual fund performance: An analysis of quarterly portfolio holdings', *The Journal of Business* (3), 393–416.

Haigh, M. S. & List, J. A. (2005), 'Do professional traders exhibit myopic loss aversion? an experimental analysis', *The Journal of Finance* **60**(1), 523–534.

Herrnstadt, E. & Muehlegger, E. (2014), 'Weather, salience of climate change and congressional voting', *Journal of Environmental Economics and Management* **68**(3), 435–448.

Hirshleifer, D., Lourie, B., Ruchti, T. G. & Truong, P. (2021), 'First impression bias: Evidence from analyst forecasts', *Review of Finance* **25**(2), 325–364.

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

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

Huberman, G. (2001), 'Familiarity breeds investment', *The Review of Financial Studies* **14**(3), 659–680.

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

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

Kaustia, M., Alho, E. & Puttonen, V. (2008), 'How much does expertise reduce behavioral biases? the case of anchoring effects in stock return estimates', *Financial Management* **37**(3), 391–412.

Koijen, R. S., Richmond, R. J. & Yogo, M. (2020), Which investors matter for equity valuations and expected returns?, Technical report, National Bureau of Economic Research.

Koijen, R. S. & Yogo, M. (2019), 'A demand system approach to asset pricing', *Journal of Political Economy* **127**(4), 1475–1515.

Krueger, P., Sautner, Z. & Starks, L. T. (2020), 'The Importance of Climate Risks for Institutional Investors', *The Review of Financial Studies* **33**(3), 1067–1111.

Kuchler, T., Li, Y., Peng, L., Stroebel, J. & Zhou, D. (2022), 'Social proximity to capital: Implications for investors and firms', *The Review of Financial Studies* **35**(6), 2743–2789.

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

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

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

Magolin, S. & Santino, K. (2022), Policy will engage with reality: Re-focus on secular gas theme as oil madness abates, Technical report, Wolfe Research.

Malmendier, U. & Nagel, S. (2011), 'Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?', *The quarterly journal of economics* **126**(1), 373–416.

Malmendier, U., Nagel, S. & Yan, Z. (2021), 'The making of hawks and doves', *Journal of Monetary Economics* **117**, 19–42.

Murfin, J. & Spiegel, M. (2020), 'Is the risk of sea level rise capitalized in residential real estate?', *The Review of Financial Studies* **33**(3), 1217–1255.

Painter, M. (2020), 'An inconvenient cost: The effects of climate change on municipal bonds', *Journal of Financial Economics* **135**(2), 468–482.

Pástor, L., Stambaugh, R. F. & Taylor, L. A. (2021), 'Sustainable investing in equilibrium', *Journal of Financial Economics* **142**(2), 550–571.

Pedersen, L. H., Fitzgibbons, S. & Pomorski, L. (2021), 'Responsible investing: The esg-efficient frontier', *Journal of Financial Economics* **142**(2), 572–597.

Pickl, M. J. (2019), 'The renewable energy strategies of oil majors: From oil to energy?', *Energy Strategy Reviews* **26**, 100370.

Reggiani, P. (2022), 'Climate change expectations: Evidence from earnings forecasts'.

Ritchie, H., Roser, M. & Rosado, P. (2020), 'Co2 and greenhouse gas emissions', *Our world in data* .

Roth, C. & Wohlfart, J. (2020), 'How do expectations about the macroeconomy affect personal expectations and behavior?', *The Review of Economics and Statistics* **102**(4), 731–748.

Ryan, S. P. (2012), 'The costs of environmental regulation in a concentrated industry', *Econometrica* **80**(3), 1019–1061.

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

Sisco, M. R. & Weber, E. U. (2022), 'Local temperature anomalies increase climate policy interest and support: Analysis of internet searches and us congressional vote shares', *Global Environmental Change* **76**, 102572.

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

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

Stroebel, J. & Wurgler, J. (2021), 'What do you think about climate finance?', *Journal of Financial Economics* **142**(2), 487–498.

van Benthem, A., Crooks, E., Giglio, S., Schowb, E. & Stroebel, J. (2022), 'The effect of climate risks on the interactions between financial markets and energy companies', *Nature Energy* **7**(8), 690–697.

Wermers, R., Yao, T. & Zhao, J. (2012), 'Forecasting Stock Returns Through an Efficient Aggregation of Mutual Fund Holdings', *The Review of Financial Studies* **25**(12), 3490–3529.

Zaval, L., Keenan, E. A., Johnson, E. J. & Weber, E. U. (2014), 'How warm days increase belief in global warming', *Nature Climate Change* **4**(2), 143–147.

A Appendix

A.1 GPT Prompt

The prompt provided to GPT-4 to analyze the investor disclosures was:

> Given a passage from Form N-CSR, evaluate if it discusses climate physical risks and climate transition risks. Fully analyze the passage first, then answer the questions. Physical climate risks can be defined as risks related to the physical impacts/damages due to climate change (e.g. heat waves, sea level rises, global warming, etc.). Transition climate risks can be defined as risks related to the transition to a lower-carbon economy (e.g. risks stemming from regulatory or governmental responses to climate change, risks from legal actions to force carbon-emitting firms to provide compensation to those harmed by climate change, or risks from shifts of consumer demand towards lower-carbon products). This does not include those passages only describing fund strategies related to climate change.
>
> ——
>
> Please answer the following questions and present your findings as a single JSON object, conforming to the following structure:
>
> 'Question1': '(choice id)';
>
> 'Question2': '(choice id)';
>
> 'Question3': Provide detailed explanations on Question1 and Question2. The explanation should be concise and precise, directly relating to the aspects mentioned in the article. (less than 100 words);
>
> Question1: Does the passage discuss physical climate risks?
>
> (a) No, the passage does not discuss the presence or effects of physical climate risks
>
> (b) Yes, the passage does discuss the presence or effects of physical climate risks
>
> Question2: Does the passage discuss transition climate risks?
>
> (a) No, the passage does not discuss transition climate risks
>
> (b) Yes, the passage does discuss transition climate risks
>
> ——
>
> Here are some examples for reference: Example 1:
>
> Passage: 'Santa Barbara employs disciplined, rigorous fundamental research combined with an objective, proprietary EcoFilter to construct the Fund's portfolio. This proprietary EcoFilter is a positive screen that focuses on environmental

and climate change practices and scores candidate stocks using a proprietary algorithm.' Question1: 'a' Question2: 'a'

Example 2: Passage: 'The agreement is intended to put an end to the dominance of fossil fuels as the primary engine of economic growth and demonstrate that governments across the planet are serious about climate change.' Question1: 'a' Question2: 'b'

Example 3: Passage: 'On the ESG front, climate change will likely remain a dominant theme, as the world witnesses more climate change-related calamities and governments introduce climate-related regulations.' Question1: 'b' Question2: 'b'

——

Now, Analyze this passage and answer the questions: <passage>

, where <passage> is replaced with the relevant text selection.

A.2 Appendix Tables

Table A.1: GPT classification examples

Risk	Passage
Neither	Commodities were once again the leaders, driven by the precious metals, which were up on safe-haven demand. Energy was slightly negative, with carbon emissions giving back a little ground. Softs were mixed, but short positions in cotton and sugar both gained nicely.
Neither	Specifically, capturing measures of quality related to sustainability, or so-called ESG (environmental, social and governance) criteria, were highly additive during market volatility. These insights, including employee sentiment and greenhouse gas emissions, provided much needed defense during periods of volatility. These contributions, however, were offset by weaker performance from traditional quality insights such as balance sheet and efficiency measures, which surprisingly failed to cushion against volatility.
Physical	Trade-related issues with China and other trading partners are not yet fully resolved. Climate change remains a concern in the form of more severe weather-related events. The Fed's policy reversed this year in the face of a slowing economy.
Transition	There were numerous policy developments in the first half of 2014 impacting upon our markets. The European Union (EU) unveiled its 2020 Climate and Energy Framework, proposing a reduction in the region's greenhouse gas emissions by 40% in 2030, but does not extend as far as legally binding renewables targets for individual member states beyond 2020. However, it sets an EU-wide goal to boost the share of renewable energy and also includes an indicative goal to boost energy efficiency by 25%.
Both	While some adjustment to our weightings can be implemented over time, in the short-to-intermediate-term our animal welfare investment strategy leads us to these over-and under-weightings, which could also have a material effect on the Funds performance. Looking forward, climate change and its impact on the planet and its inhabitants, especially animals, requires changes in the behaviors of all stakeholders, especially corporations. Were already seeing rapid and disruptive change in many industries, including fossil fuels and food and transportation, due to a growing understanding of the threats posed by a warming climate.

Note: This table contains examples of sentences we selected and input into GPT-4, along with its resulting classification as explained in Section 2.1. The sentences could be indicated as not discussing climate risk ("neither"), only discussing physical risk ("physical"), only discussing transitional risk ("transitional"), or discussing both risks ("both").

Table A.2: BERT classification examples

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

Note: This table contains examples of sentences selected for the macro hedging disclosure shock in Section 5, along with their classification by the BERT model. Sentences classified as positive are assigned a score of 1 and sentences classified as negative are assigned a score of -1.

Table A.3: Sample Summary Statistics

Panel A: Industry Summary Statistics

		Number of Companies			Share of Fund Holdings (%)		
GICS	Industry	Avg.	Min	Max	Avg.	Min	Max
1010	Energy	245	217	265	7.6	4.2	11.6
1510	Materials	211	180	233	3.7	2.1	4.8
2010	Capital Goods	330	308	351	7.6	5.0	8.6
2020	Commercial & Prof. Serv.	129	118	144	1.5	1.3	1.7
2030	Transportation	71	61	86	2.6	1.8	3.2
2510	Auto & Components	44	40	46	0.9	0.6	1.2
2520	Consum. Durables & Apparel	122	112	140	2.0	1.3	2.5
2530	Consumer Services	142	129	153	2.9	2.7	3.3
2550	Retailing	154	143	161	5.5	2.7	6.5
3010	Food & Staples Retailing	27	23	32	1.4	1.1	1.6
3020	Food, Bev. & Tobacco	94	83	105	4.3	3.3	5.2
3030	Household & Pers. Prod.	38	35	46	1.6	1.2	1.9
3510	Health Care Equip. & Serv.	256	233	291	6.1	5.4	6.9
3520	Pharma., Biotech., & Life Sc.	382	271	535	8.2	6.5	10.1
4010	Banks	433	400	503	6.4	5.3	8.4
4020	Diversified Financials.	153	145	160	4.9	4.2	6.1
4030	Insurance	109	95	131	3.0	2.4	3.5
4510	Software & Services	284	262	309	9.4	7.8	13.2
4520	Tech. Hardw. & Equip.	222	174	275	5.3	2.1	7.3
4530	Semiconductors & Equip.	108	81	135	3.5	2.8	4.4
5010	Communication Services	42	31	53	1.7	1.3	2.5
5020	Media & Entertainment	106	83	133	5.2	2.2	12.1
5510	Utilities	92	77	105	2.7	2.4	3.2
6010	Real Estate	152	113	185	2.3	1.2	4.4

Panel B: Mutual Fund Summary Statistics

	Number of Companies			Number of Industries		
	Avg.	p10	p90	Avg.	p10	p90
Mutual Fund Holdings	211	33	469	19.6	14.0	24.0

Panel C: Active Changes Summary Statistics

	Avg.	p1	p25	p50	p75	p99
Active Industry Change	-0.12	-121.79	-6.32	0.00	5.99	124.97

Note: Panel A shows, for each industry, (i) the average, minimum, and maximum of the number of unique companies held by at least one fund in our sample, and (ii) the average, minimum, and maximum industry market share, all calculated across the 40 quarters between 2010 and 2019. Panel B shows the average and the 10th and 90th percentiles of unique companies and industries held by individual funds across the 72,732 observed fund-quarter observations. Panel C shows summary statistics for the active industry changes measure in pp as defined in Equation (2). The unit of observation is a fund-quarter-industry change and the sample size is 1,391,064. Note that we require two consecutive holding reports to observe an active change, which are not always available.

Table A.4: Industry-Specific Climate Quantity Betas Significance

GICS	Description	Pooled Shock	Disclosure Shock	Heat Shock
2030	Transportation	1.95	1.98	0.46
2510	Auto & Components	1.74	1.83	0.91
4530	Semiconductors & Equip.	1.39	0.53	2.14
2010	Capital Goods	1.79	2.22	0.33
1510	Materials	1.31	1.18	0.82
4010	Banks	1.93	0.50	2.28
3030	Household & Pers. Prod.	1.18	1.17	−0.10
1010	Energy	1.69	1.64	1.78
4520	Tech. Hardw. & Equip.	0.65	−1.62	2.72
2530	Consumer Services	0.09	−0.21	0.09
4020	Diversified Financials.	−0.11	0.61	0.31
4510	Software & Services	−0.20	0.52	0.77
3010	Food & Staples Retailing	−0.16	0.02	0.58
3020	Food, Bev. & Tobacco	−0.81	1.32	−1.79
2520	Consum. Durables & Apparel	−0.19	0.06	1.06
5020	Media & Entertainment	−0.84	1.07	−1.12
5010	Communication Services	−0.81	0.26	−0.48
5510	Utilities	−1.00	1.45	−1.82
3520	Pharma., Biotech., & Life Sc.	−1.51	1.16	−2.09
3510	Health Care Equip. & Serv.	−1.17	−1.61	−0.59
4030	Insurance	−1.34	−0.61	−1.19
2020	Commercial & Prof. Serv.	−0.78	−0.80	−1.05
6010	Real Estate	−1.98	−0.66	−2.12
2550	Retailing	−2.25	1.13	−2.97
	Joint Significance Test	0.0264	0.0001	0.0038

Note: This table shows the t-stats of industry-specific climate quantity betas as reported in Table 3, and P-values of the joint significance test in the last row. We test if the climate quantity betas for all industries deviate from zero jointly, controlling the month by industry fixed effect and the standard errors are clustered at the fund level. The coefficients are based on pooled data from 2010 to 2019 inclusive.

Table A.5: Industry-Specific Climate Quantity Betas (unscaled)

GICS	Description	Pooled Shock	Disclosure Shock	Heat Shock
2010	Capital Goods	16.1073*	82.0678**	4.4132
1010	Energy	15.9075**	38.6244	18.0780*
4010	Banks	12.3007**	7.3828	18.9470**
2030	Transportation	9.7578*	48.0269*	2.0660
1510	Materials	9.5616*	26.8613	9.4658
4530	Semiconductors & Equip.	6.1142	4.9171	14.4070**
4520	Tech. Hardw. & Equip.	3.6722	-47.6637*	20.4623***
2510	Auto & Components	3.5100	23.3109*	2.3746
3030	Household & Pers. Prod.	2.7243	13.6251	-0.6114
4020	Diversified Financials.	2.6979	10.4414	7.7727
2530	Consumer Services	0.6662	-5.2088	0.2067
5510	Utilities	-1.8847	22.2410	-4.8670
5010	Communication Services	-2.1177	3.5863	-1.0000
4510	Software & Services	-2.1428	15.1465	7.1090
2520	Consum. Durables & Apparel	-2.5238	-2.1896	4.5504
2020	Commercial & Prof. Serv.	-3.8104	-10.4609	-5.9619
4030	Insurance	-3.9739	-11.3642	-4.2553
3010	Food & Staples Retailing	-4.2850	0.1823	-4.1413
5020	Media & Entertainment	-4.7539	17.3065	-7.6636
3020	Food, Bev. & Tobacco	-6.1528	30.8282	-15.2827**
3520	Pharma., Biotech., & Life Sc.	-7.9134	49.4110	-14.8916**
6010	Real Estate	-7.9772*	-10.0917	-9.6117*
3510	Health Care Equip. & Serv.	-10.0095	-49.0774	-6.2651
2550	Retailing	-16.7284**	25.1143	-30.7865***

Note: This table reports industry-specific climate quantity betas, multiplied by 100 for easy readability, as in Equation (3), but without scaling ActiveChanges by industry size in Equation (2). The coefficients are based on pooled data from 2010 to 2019 inclusive. Industries are sorted by the "Pooled Shock".

Table A.6: Climate News Series Correlations

Panel A: Climate News Correlations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Faccini et al.									
(1) IntSummit	1.00								
(2) GlobWarm	0.28	1.00							
(3) NatDis	0.43	0.37	1.00						
(4) Narrative	−0.26	−0.23	−0.31	1.00					
Engle et al.									
(5) WSJ	0.04	0.64	0.46	−0.30	1.00				
(6) CHNEG	0.16	0.46	0.69	−0.08	0.23	1.00			
Ardia et al.									
(7) MCCC	−0.13	0.35	0.07	0.43	0.38	0.19	1.00		
Google									
(8) National	0.38	0.20	0.34	0.10	0.07	0.18	0.52	1.00	
Temperature									
(9) National	−0.08	−0.20	−0.10	−0.11	0.08	−0.31	0.01	−0.32	1.00

Panel B: Climate News AR(1) Innovations Correlations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Faccini et al.									
(1) IntSummit	1.00								
(2) GlobWarm	0.22	1.00							
(3) NatDis	0.09	0.36	1.00						
(4) Narrative	−0.20	−0.25	−0.25	1.00					
Engle et al.									
(5) WSJ	0.02	0.46	0.37	−0.06	1.00				
(6) CHNEG	0.09	0.39	0.43	−0.07	0.35	1.00			
Ardia et al.									
(7) MCCC	−0.02	0.46	0.05	0.21	0.60	0.49	1.00		
Google									
(8) National	0.20	0.30	0.12	−0.05	0.38	0.52	0.45	1.00	
Temperature									
(9) National	0.03	−0.05	−0.13	0.07	0.01	0.04	−0.06	0.11	1.00

Note: This table reports the correlation across climate news series, using monthly data from 2015-2020. Panel A shows the correlation between news indices. Panel B shows the correlation of their AR(1) innovations.

Table A.7: Climate Hedge Performance of Various Portfolios (Bootstrap Standard Errors)

		Faccini et al.				Engle et al.		Ardia et al.	Google	Temp.
	Avg.	IntSummit	GlobWarm	NatDis	Narrative	WSJ	CHNEG	MCCC	National	National
Pooled Shock	0.05	0.14	0.11	0.13	0.14	0.15	0.17	0.17	0.14	0.14
Heat Shock	0.05	0.14	0.12	0.10	0.15	0.14	0.16	0.13	0.19	0.14
Disclosure Shock	0.05	0.16	0.11	0.13	0.12	0.15	0.17	0.20	0.15	0.13
Emission Portfolio	0.05	0.13	0.14	0.11	0.12	0.17	0.14	0.17	0.13	0.14
Long PBD ETF	0.05	0.13	0.12	0.12	0.13	0.25	0.19	0.20	0.13	0.12
Short Stranded Asset	0.04	0.13	0.14	0.09	0.11	0.18	0.14	0.15	0.10	0.13
Long ICLN ETF	0.05	0.13	0.13	0.11	0.13	0.22	0.19	0.19	0.13	0.13
Short IYE ETF	0.04	0.14	0.12	0.12	0.11	0.16	0.17	0.14	0.13	0.13
Short XLE ETF	0.04	0.14	0.11	0.11	0.11	0.17	0.16	0.14	0.13	0.12
10-K Negative Portfolio	0.04	0.13	0.12	0.11	0.14	0.13	0.13	0.12	0.10	0.16
Sustainalytics Portfolio	0.04	0.12	0.10	0.09	0.15	0.12	0.14	0.14	0.11	0.15
Lasso: All Industry+FF	0.04	0.08	0.13	0.08	0.14	0.18	0.10	0.12	0.10	0.00
Lasso: Factor Zoo	0.04	0.10	0.11	0.08	0.11	0.16	0.00	0.18	0.00	0.00
Reg: ETFs+FF	0.04	0.12	0.11	0.09	0.14	0.13	0.13	0.14	0.20	0.14
Reg: FF 3-Factors	0.05	0.11	0.11	0.09	0.14	0.14	0.13	0.17	0.20	0.13
Reg: SPY ETF	0.04	0.09	0.12	0.15	0.09	0.18	0.17	0.15	0.12	0.11

Note: Standard errors of monthly correlations for various climate hedge portfolios' returns with various climate news series AR(1) innovations, calculated using a bootstrap. The bootstrap process is replicated 1000 times, with each iteration involving resampling the full-time period with replacement. In each iteration, we calculate the correlations between the returns of various climate hedge portfolios and climate news, as well as the average correlation across all climate news series using the resampled datasets. The standard errors are calculated as the standard deviations of the correlation values obtained from the 1000 bootstrap iterations. Each row represents a hedge portfolio, whereas each column corresponds to the innovations of a climate news series.

Table A.8: Climate Hedge Performance of Various Portfolios (two-sided p-val of difference)

		Faccini et al.				Engle et al.		Ardia et al.	Google	Temp.
	Avg.	IntSummit	GlobWarm	NatDis	Narrative	WSJ	CHNEG	MCCC	National	National
Heat Shock	0.88	0.47	0.65	0.66	0.83	0.87	0.88	0.82	0.81	0.86
Disclosure Shock	0.35	0.61	0.51	0.45	0.98	0.89	0.76	0.76	0.14	0.93
Emission Portfolio	0.16	0.21	0.40	0.74	0.77	0.63	0.42	0.61	0.11	0.21
Long PBD ETF	0.14	0.41	0.29	0.72	0.91	0.76	0.52	0.75	0.06	0.74
Short Stranded Asset	0.02	0.18	0.20	0.42	0.62	0.70	0.96	0.07	0.05	0.11
Long ICLN ETF	0.02	0.29	0.09	0.55	0.65	0.64	0.34	0.45	0.03	0.60
Short IYE ETF	0.00	0.12	0.01	0.19	0.56	0.35	0.72	0.01	0.03	0.32
Short XLE ETF	0.00	0.12	0.01	0.15	0.51	0.36	0.68	0.01	0.03	0.35
10-K Negative Portfolio	0.00	0.34	0.08	0.52	0.77	0.20	0.17	0.01	0.00	0.37
Sustainalytics Portfolio	0.00	0.61	0.02	0.69	0.83	0.07	0.10	0.02	0.00	0.24
Lasso: All Industry+FF	0.00	0.63	0.26	0.91	0.96	0.81	0.36	0.00	0.00	0.47
Lasso: Factor Zoo	0.00	0.44	0.06	0.22	0.45	0.39	0.44	0.14	0.01	0.47
Reg: ETFs+FF	0.00	0.24	0.17	0.62	0.75	0.81	0.17	0.01	0.13	0.53
Reg: FF 3-Factors	0.00	0.17	0.19	0.67	0.50	0.41	0.07	0.02	0.10	0.66
Reg: SPY ETF	0.00	0.05	0.09	0.22	0.52	0.28	0.27	0.06	0.04	0.32

Note: Bootstrap two-side p-values of the difference between correlation of target news with the "Pooled Shock" and the correlation with the returns of all other portfolios. We conduct 1000 resampling iterations. In each iteration, we randomly select data from the full-time period with replacement. Each row represents a hedge portfolio, whereas each column corresponds to the innovations of a climate news series.

Table A.9: Ex-post Evaluation

		Faccini et al.				Engle et al.		Ardia et al.	Google	Temp.
	Avg.	IntSummit	GlobWarm	NatDis	Narrative	WSJ	CHNEG	MCCC	National	National
R^2	0.48	0.50	0.27	0.29	0.39	0.81	0.74	0.60	0.38	0.32
Adjusted-R^2	-0.03	0.15	-0.24	-0.22	-0.05	-0.11	0.35	0.02	-0.04	-0.13

Note: In-sample R^2 and Adjusted-R^2 from regressing each news series on all 24 industry portfolio returns. The sample consists of monthly observations from 2015-2019.

A.3 Appendix Figures

Figure A.1: Distribution of "Heat Shock"

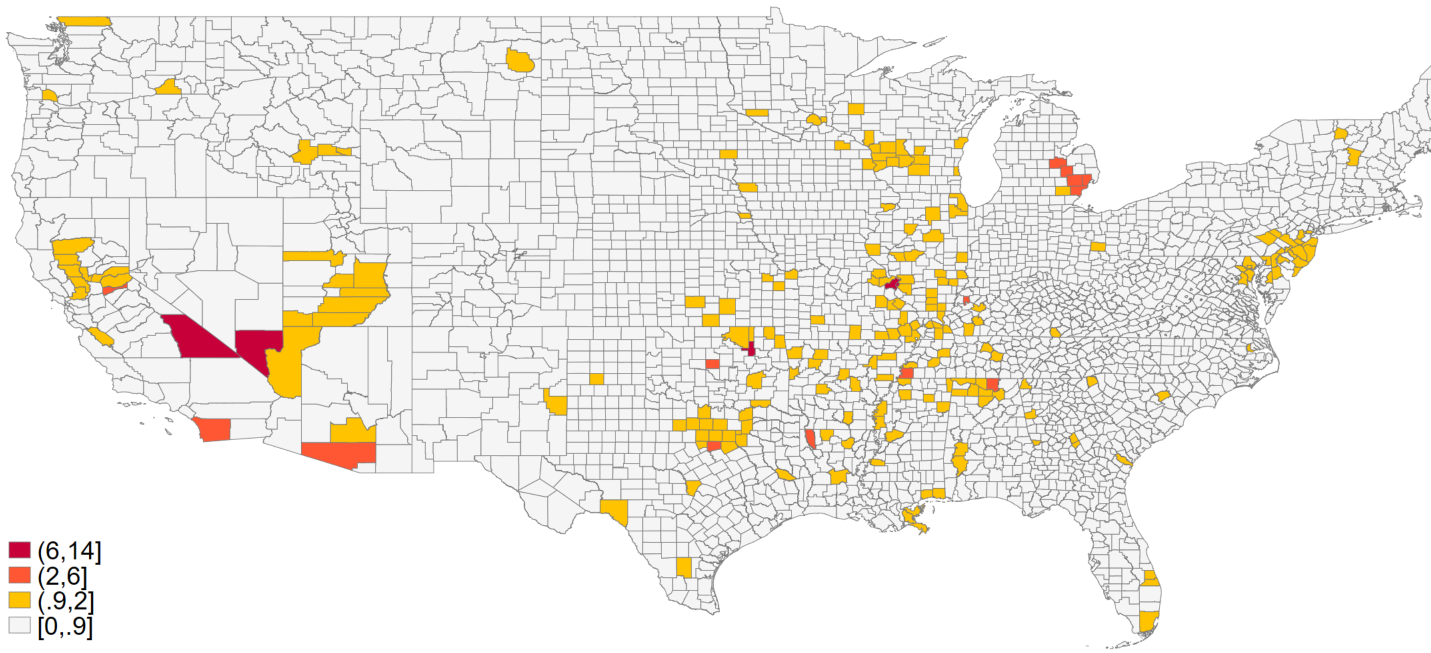

Note: Spatial distribution of the "Heat Shock" from 2010 to 2019. The color-coding shows the number of county-months that experienced the shock during that time interval.

Figure A.2: Industry Ranking by Scaled and Unscaled Industry Betas

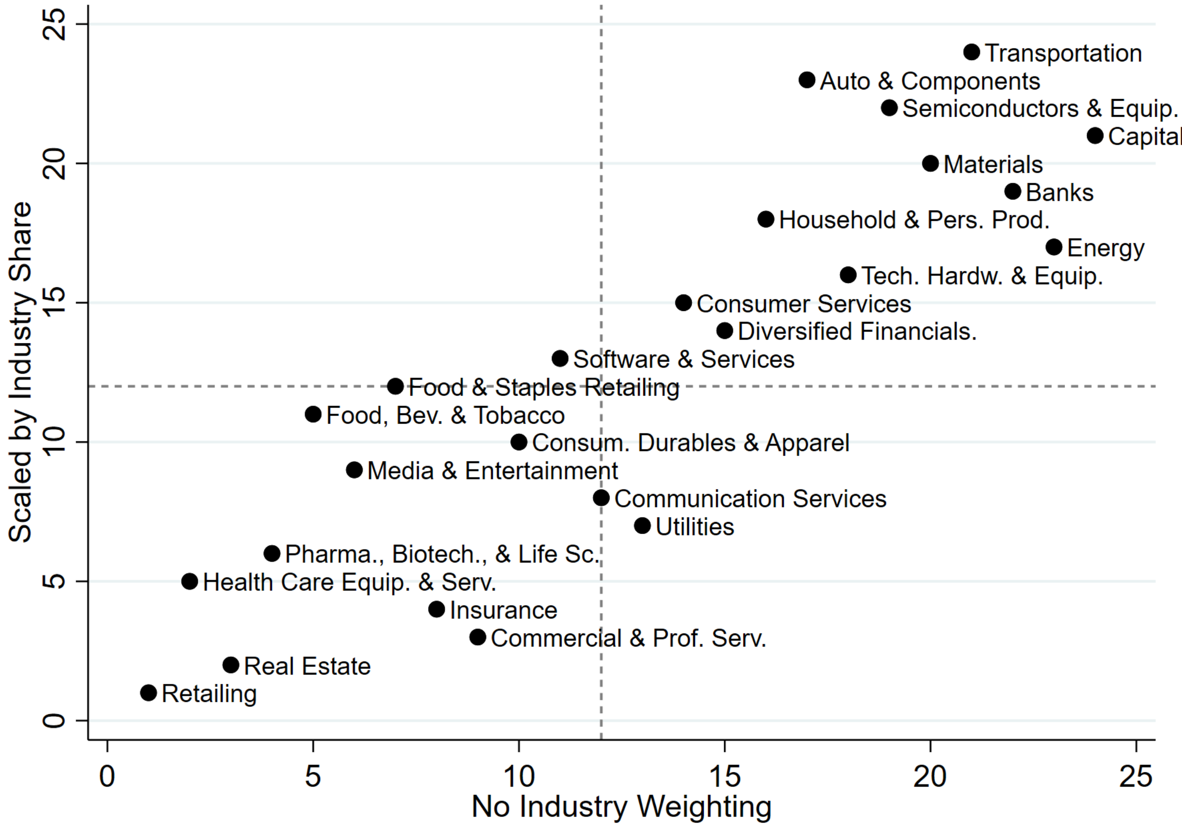

Note: Scatterplot of industry beta ranking estimated with scaled and unscaled ActiveChanges. Both coefficients are based on pooled data from 2010 to 2019 inclusive. Industries are ranked by the "Pooled Shock". Industry that is least exposed (rank=24) has the highest positive quantity beta.

Figure A.3: Industry Betas Rankings Over Time

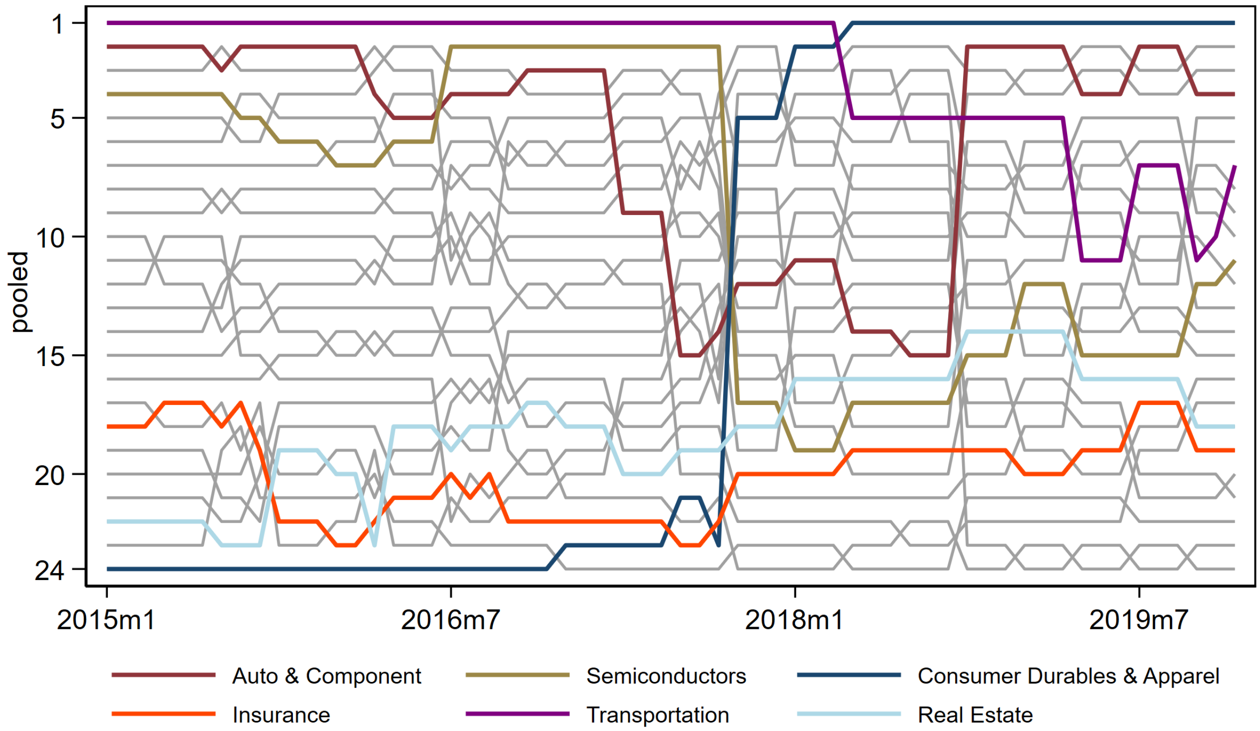

Note: Industry climate beta coefficients (estimated in Equation 3) rankings over time. The rankings are sorted by the "Pooled Shock" and are based on data for every five-year window from 2010 to 2019 inclusive. Each line represents an industry.

Figure A.4: Climate Hedge Performance - Robustness of Portfolio Construction Choices

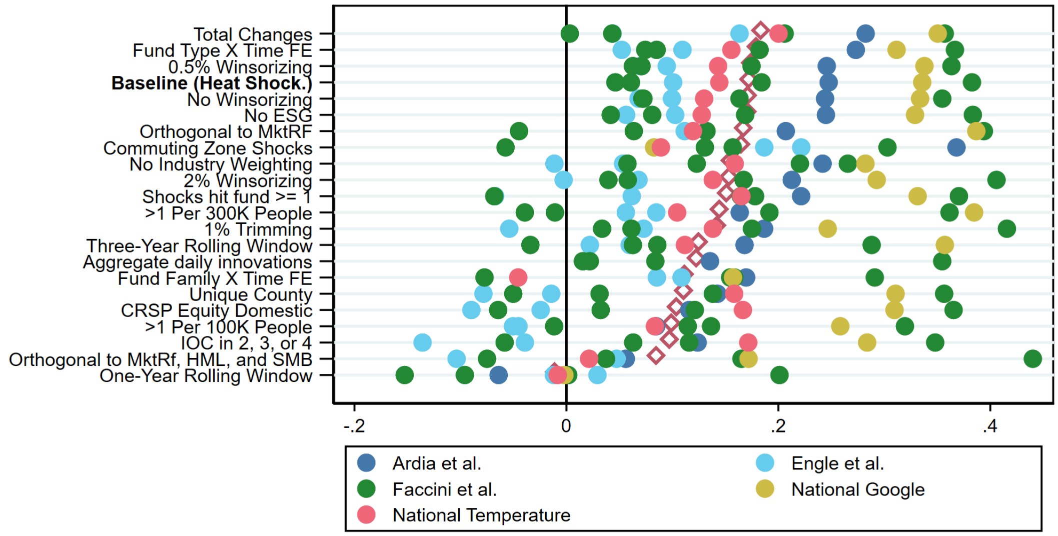

(a) Heat Shock

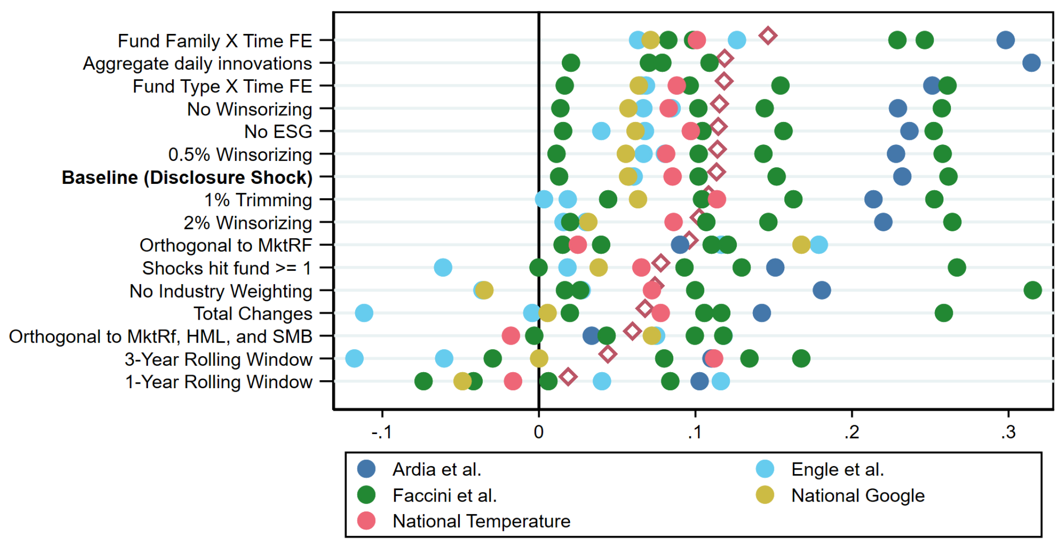

(b) Disclosure Shock

Note: Dot plot of monthly out-of-sample return correlations for the heat shock and disclosure shocks against various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Each row represents a different robustness check. The red rhombus shows the unweighted average among all correlations, and portfolios are sorted top-to-bottom by this value.

Figure A.5: Climate Hedge Performance of Various Portfolios (Alternative Heat Shocks)

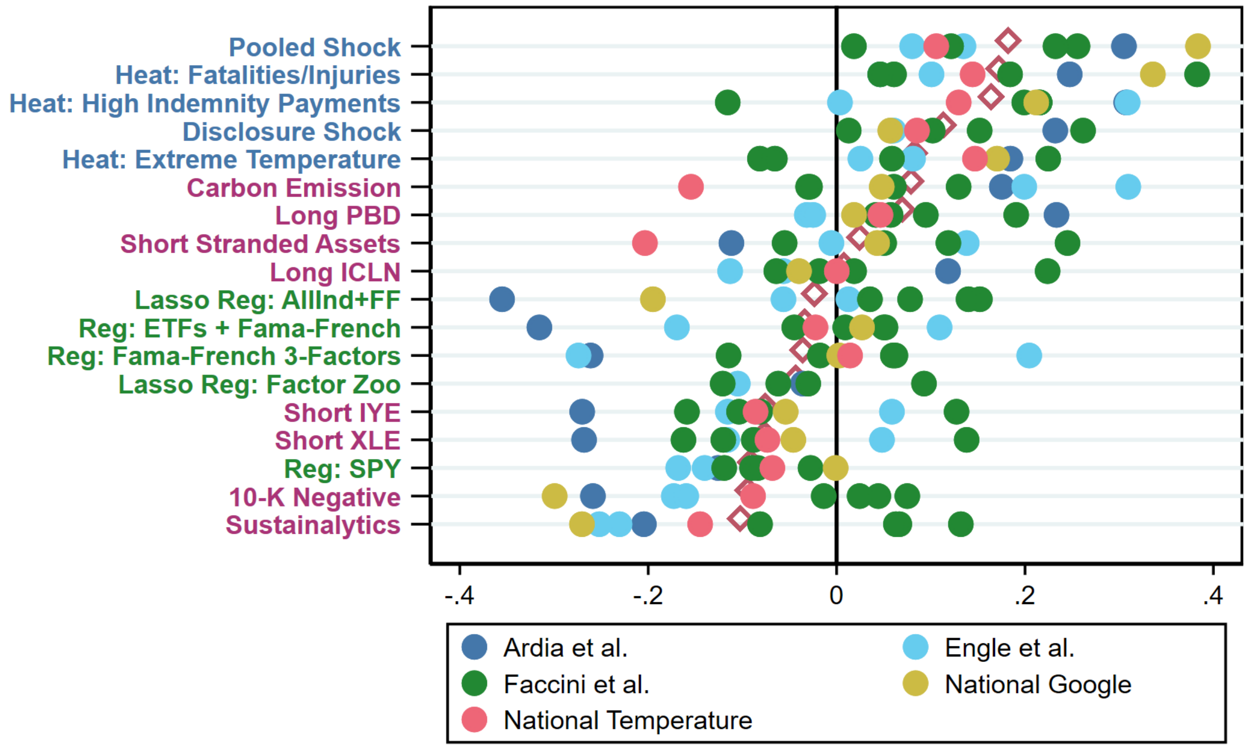

Note: Dot plot of monthly out-of-sample return correlations for various climate hedge portfolios with various climate news series AR(1) innovations. Each dot represents one correlation coefficient. Different colors represent different groups of climate news series. The red rhombus shows the unweighted average among all correlations, and portfolios are sorted top-to-bottom by this value.

Figure A.6: Climate Hedge Performance of Various Granular Portfolios

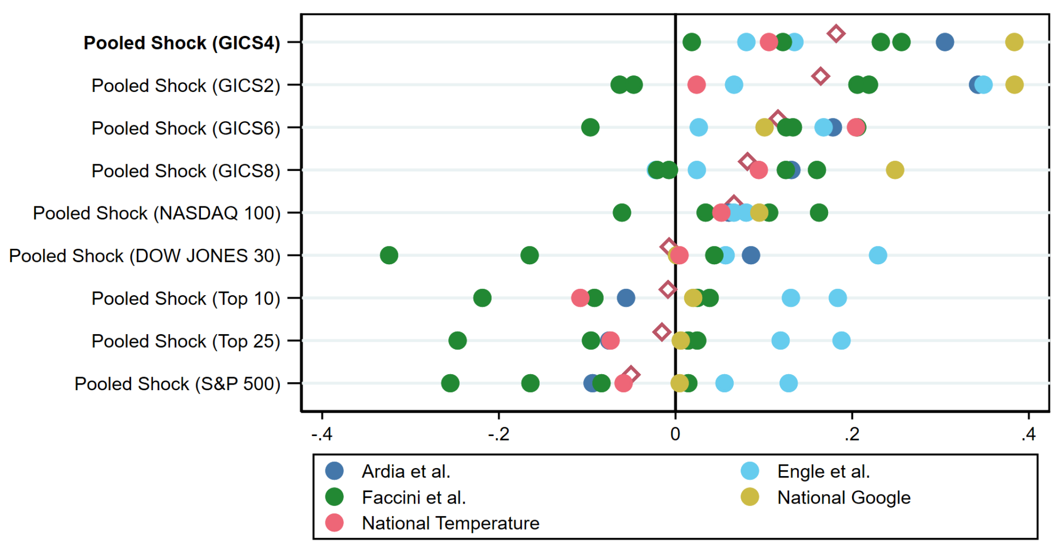

Note: Dot plot of monthly out-of-sample return correlations for various climate hedge portfolios based on the "Pooled Shock" with various climate news series AR(1) innovations. Each row corresponds to a different levels of aggregation used to estimate the climate quantity betas for portfolio formation. Each dot represents one correlation coefficient. Different colors represent different groups of climate news series. The red rhombus shows the unweighted average among all correlations, and portfolios are sorted top-to-bottom by this value.

Figure A.7: Placebo test in earlier period

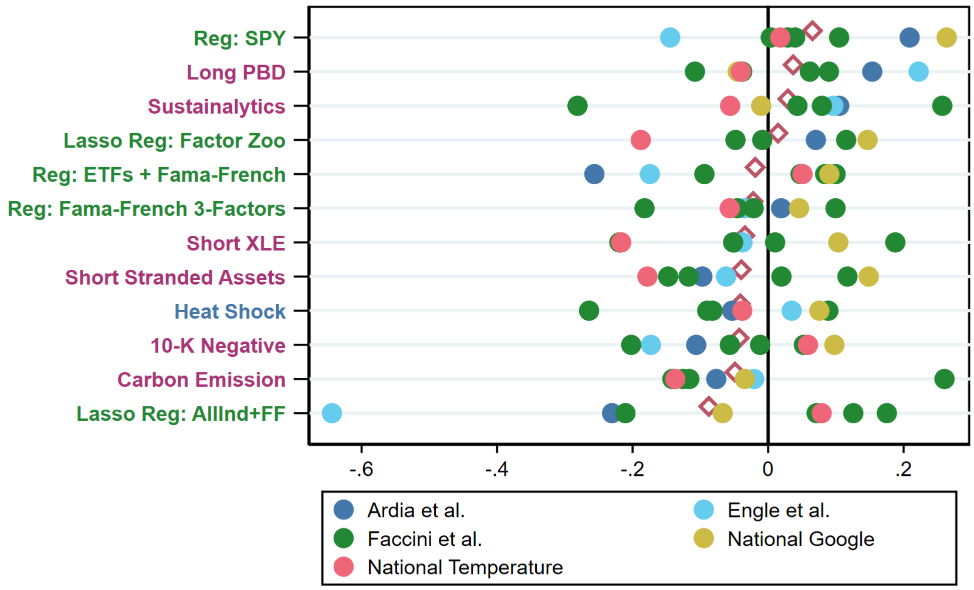

Note: Dot plot of monthly return correlations for the three heat-based hedging portfolios and mimicking portfolios with various climate news series AR(1) innovations, using data from 2000 to 2010. The PBD ETF is only available after July 2007.