COLUMBIA | SIPA Center for Environmental Economics and Policy

CEEP WORKING PAPER SERIES Working Paper Number 2

February 2019 This Version: September 2020

Market Expectations of a Warming Climate

Wolfram Schlenker and Charles Taylor

https://ceep.columbia.edu/sites/default/files/content/papers/n2_v2.pdf Online Appendix

Market Expectations of a Warming Climate*

Wolfram Schlenker^{1,2,*}, Charles A. Taylor¹

Abstract

We compare prices of financial derivatives whose payouts are based on future weather outcomes to CMIP5 climate model predictions as well as observed weather station data across eight cities in the US from 2001 through 2020. Derivative prices respond both to short-term weather forecasts for the next two weeks and longer-term warming trends. We show that the long-term trends in derivative prices are comparable to station-level data and climate model output. The one exception is February in the northeastern US, where financial markets price in a polar vortex-induced cooling effect, a recent scientific finding that was not present in the older CMIP5 climate output. When looking at the spatial and temporal heterogeneity in trends, future prices are more aligned with climate model output than observed weather station trends, suggesting that market participants closely align their expectations with scientific forecasts rather than recent observations.

Keywords: Market expectations, belief formation, weather markets, climate change *JEL:* Q54, Q02, D84

Scientists overwhelmingly agree that the climate is changing because of human activity. The American Association for the Advancement of Science (December 9, 2006) reported that "the scientific evidence is clear: global

^{*}We are grateful to Siqi He for extracting the Futures data from Bloomberg Terminals. This work was supported by the U.S. Department of Energy, Office of Science, Biological and Environmental Research Program, Earth and Environmental Systems Modeling, MultiSector Dynamics, Contract No. DE-SC0016162.

^{*}Corresponding author. E-mail address: wolfram.schlenker@columbia.edu

¹School of International and Public Affairs, Columbia University, 420 West 118th St., New York, NY 10027.

 $^{^2 \}rm National Bureau of Economic Research (NBER), 1050 Massachusetts Ave., Cambridge, MA 02138.$

climate change caused by human activities is occurring now." But public opinion in the US remains mixed. As of 2016, less than half of Americans believed that the Earth is getting warmer due to human activity, a number which has not budged much since the Pew Research Center started asking the question in 2006.³ Views on climate change vary greatly across geography, political affiliation, educational status, and economic sector (Leiserowitz et al., 2017). Politicians in the US have questioned the evidence on climate change, with some famously calling it an "elaborate hoax."

Given the divergent beliefs about climate change, debate persists about the accuracy of global climate models and the extent to which agents incorporate these forecasts into their actions. We address these issues by examining how market participants update their expectations about weather over time. The Chicago Mercantile Exchange (CME) offers futures contracts for eight cities on two main weather products: cooling degree days, which measure how much cooling is necessary during hot temperatures in summer, and heating degree days, which measure how much heating is required during cold temperatures in winter. The payoffs from these contracts depend on observed temperatures over the course of a month. The contracts are traded before the month in which the weather is realized, and thus provide a direct measure of the market's view on the expected climate.

First, we show that the futures market capitalizes weather shocks, i.e., deviations from climate averages, in the two weeks before an unexpected weather deviation occurs. This is consistent with Dorfleitner and Wimmer (2010) and the more general finding that for horizons beyond 8-10 days, "the nature of temperature dynamics simply makes any point forecast of temperature unlikely to beat the climatological forecast at long horizons, because all point forecasts revert fairly quickly to the climatological forecast" (Campbell and Diebold, 2005). Futures prices several weeks before the start of a month should reflect expectations about a month's weather before the outcomes can be known.

Second, we find that market expectations, as measured by futures prices when weather outcomes are unknown, have been changing at the same annual rate as temperature forecasts in the CMIP5 archive, the latest repository in which various climate modeling groups made predictions for 2006

 $^{^{3}} https://www.pewresearch.org/science/2016/10/04/public-views-on-climate-change-and-climate-scientists/.$

onward. The time trend also aligns with the observed annual change from weather station data. All find significant warming, i.e., an increase in cooling degree days in summer and a decrease in heating degree days in winter. Climate models' predictions have materialized, especially on average, validating model forecasts.

Third, the futures market closely follows advances in the climate literature. When we regress the trend in futures prices for each airport and contract-month observed over our sample period on the observed trend at the weather station as well as climate forecasts, the latter has the most explanatory power. Further, the futures market seems to price in recent climatological advances that were not available in the CMIP5 archive and that have not been detectable in weather station observations. Recent research predicts that a shift in the jet stream will reduce late winter temperatures in the northeastern US via an increase in cold air from the Arctic (i.e., a polar vortex). Likewise, the futures market has shown a significant increase in heating degree days in February. Together this suggests that market participants are taking into account both global climate model output and the latest research rather than simply projecting forward past time trends.

Finally, we present evidence in the Appendix how oceanic oscillations like El Niño - Southern Oscillation (ENSO) affect temperatures over the medium term across the eight cities in our sample. Employing LASSO regressions to select relevant oceanic oscillation indices, we find that removing these largescale effects reduces the year-to-year variability in observed weather, but does not change the time trend. The observed warming trend is hence not driven by oceanic drivers of natural variability in temperatures, but rather increased greenhouse gas emissions.

In addition to contributing to the literature on the impact of climate change on firms and financial markets, our findings have relevance to climate adaptation. Economists have estimated the benefits and costs from a changing climate (Auffhammer, 2018). Many of the recent micro-level estimates relate outcomes of interest to random exogenous year-to-year weather fluctuations to obtain unbiased damage estimates (Dell et al., 2014). While random and exogenous year-to-year variation is preferable from a statistical perspective, adaptation to a permanent change in climate might mitigate some of the weather sensitivity that is observed in response to unknown random weather shocks. Agents should undertake adaptation investments in response to anticipated permanent shifts in the climate that are either unprofitable or infeasible for a one-time unknown weather shock. However, before agents can adapt, they first must form a belief about whether the climate is changing and the extent to which it is. This paper suggests that agents, at least those participating in weather markets, have been actively updating their beliefs about the extent and geography of warming.

Our paper adds to several strands of literature. The first examines the impact of weather fluctuations and climate change on the corporate sector and financial markets. Corporate earnings of several economic sectors are sensitive to temperature fluctuations (Addoum et al., 2020), and understanding the extent to which financial markets are pricing in climate change risks has implications for financial stability (Carney, 2015). Some papers find that the stock market underreacts to the impact of predictable climatic trends on firms' profitability and valuation (Hong et al., 2019), while others show that real estate market and municipal bonds do price in sea level rise (Bernstein et al., 2019) and that agricultural land markets capitalize climate change expectations (Severen et al., 2018). Weather derivates can provide a useful hedge against such fluctuations as well as a direct measure of the market's expectation of future climate.

Second, studies have emphasized how climate change *policies* designed to limit emissions can affect firm profitability using financial data. Anttila-Hughes (2016) finds that energy company valuations respond to extreme events that may be evidence of climate change. Meng (2017) shows how the stock market incorporates changes in the likelihood of US carbon regulation as measured by betting markets. Limiting emissions may render a company's marginal reserves, i.e., the most costly ones, worthless if they can no longer be extracted (McGlade and Ekins, 2015). Thus, future climate expectations are key to the energy sector's profitability and will be reflected in financial markets.

Third, another strand of the literature focuses on how agents adjust their behavior in response to environmental forecasts (Rosenzweig and Udry, 2014; Neidell, 2009). Shrader (2020) finds that fishermen update their beliefs using El Niño medium-range weather forecasts in order to make optimal fishing decisions. Before El Niño forecasts were available, the cost of weather shocks was much higher because fisheries could not adapt. On the other hand, Burke and Emerick (2016) find that changes in agricultural yields in response to observable long-term temperature trends are not significantly different from yield changes in response to random weather shocks. Some authors have modeled how market participants learn about and adapt to changing weather conditions. For example, Kala (2019) examines how Indian farmers dependent on monsoon precipitation update their beliefs. Twitter reactions show that people become habituated to extreme weather events as they become more frequent over time (Moore et al., 2019). Similarly, public opinion surveys ask respondents to self-report their beliefs. Public opinion also seems driven by recent weather events, especially extremes. Many studies have shown that people's beliefs about climate change are strongly influenced by recent local weather conditions (Myers et al., 2013; Deryugina, 2013; Akerlof et al., 2013; Li et al., 2011; Zaval et al., 2014). Observed periods of cooling can translate into climate skepticism (Kaufmann et al., 2017). It is also possible that agents hold differing private and public beliefs about climate change, especially if certain views on climate change are perceived as more expedient. Significant variation in public opinion about climate change exists across the US, varying by location and demographic characteristics (Howe et al., 2015).

What is common across most of the existing studies about climate change expectations is that researchers infer climate beliefs indirectly, i.e., by backing them out from observed indirect actions or by relying on stated responses. We add to this literature by using a different revealed preference approach to measure beliefs about climate change by analyzing financial derivatives whose value directly depends on expected weather. This allows us to observe the evolution of market expectations on warming by looking at the price of futures contracts that are linked to future weather outcomes.

1. Data

1.1. Financial Data

Weather futures contracts are traded on the Chicago Mercantile Exchange (CME). The products were first launched in the fall of 2001 and became fully operational for the first full year in 2002. Contracts are available for eight geographically-distributed cities across the US over our sample period 2001-2020. Each city is linked to a specific weather station in the city at one of the airports. These are: Atlanta (ATL), Chicago O'Hare (ORD), Cincinnati - Northern Kentucky (CVG), Dallas-Fort Worth (DFW), Las Vegas (LAS), Minneapolis - Saint Paul (MSP), New York LaGuardia (LGA), and Sacramento (SAC). The location across the US is displayed in Appendix Figure A1. In the past more cities had weather markets, but trading in several cities was halted due to a lack of liquidity, while at the same time new cities like Portland and Tokyo were launched as recently as 2019. Therefore, we focus on

the eight US cities for which contracts were consistently available through spring 2020.

The main participants in the weather market are insurance companies and firms seeking to offset weather risk. For example, an energy company may sell an HDD contract to mitigate the risk of lower demand for heating oil due to a mild winter. Likewise, a citrus company may purchase an HDD contract to mitigate the risk of a winter freeze. The other market participants are speculators who take contract positions based on their expectations of future weather. More generally, volumes in this market decreased in recent years due to the entry of reinsurance firms offering bespoke weather-based hedging services to market participants.

The final settlement price of the futures contract is based on the respective weather station HDD or CDD index for the month as reported by MDA Federal Information Systems, Inc. Each degree day in a contract has a payout multiplier of \$20. For example, if a customer buys one July CDD contract for 300 cooling degree days, the cost would be \$6,000. If the realized cumulative cooling degree days for the month of July settled at 330 degree days, the clearance value would be \$6,600, and the trader would reap a profit of \$600 (\$20 times the increase of 30 degree days). Trading volume generally increases in the two weeks prior to the start of a contract month, with lower trade volume more than two weeks before the start of the contract month.

The weather contracts are based on cumulative heating and cooling degree days in a given month. These are indexed to 65° F (18°C), the temperature considered the most comfortable for humans on average, and a common standard for utility companies because cooling and heating systems tend to be turned on above and below that level, respectively. For example, a mean daily temperature of 85° F would count as 20 cooling degree days. These daily degree days are then summed over the course of the contract month.

Cooling degree days (CDD) measure by how much daily average temperatures T_{ad} at airport *a* on day *d* exceed 65°F and thus require cooling, hence the name cooling degree days. The exact formula to derive CDD_{*am*} for month *m* is obtained by summing over all days d(m) of the month

$$CDD_{am} = \sum_{d(m)} \max\{T_{ad} - 65, 0\}$$
 (1)

Likewise, heating degree days (HDD) measure by how much and for how long temperature fall below 65°F and thus require heating. The exact formula to

derive HDD_{am} is

$$HDD_{am} = \sum_{d(m)} \max\{65 - T_{ad}, 0\}$$
(2)

For our baseline analyses, we use end-of-day daily futures prices obtained from Bloomberg terminals. Prices are carried forward in the absence of market activity. For example, if there is a recorded trade on June 17 at a price of 300 cooling degree days for the July contract, followed by no trade on June 18, the Bloomberg data will show a price of 300 again. Unfortunately, the volume data only includes contracts traded via the exchange and not private over-the-counter block trades (Dorfleitner and Wimmer, 2010)⁴, and it is missing for most days. Some data cleaning was necessary because of "sticky fingers," e.g., sudden price jumps by a factor of 10. The exact adjustments are listed in Appendix Section A1.

The raw daily data we downloaded from the Bloomberg terminals is displayed in Figure 1 for the two airports with the highest volume in cooling degree days: LGA and DFW. We pick two representative months: The left column shows cooling degree days contracts in July, while the right column shows heating degree day contracts for December. Contracts for the remaining airports and months are shown in Appendix Figure A2. Each graph displays the annual prices series for roughly two and a half months. Day 0 is the last day of the month on which the contract is based. Both the end of the month and the beginning of the month are indicated by vertical dashed black lines. The temporal extent ranges from 70 days prior to the end of the contract month (roughly 40 days prior to start of the contract month) to 10 days past the end of the contract month. Years are color-coded from blue (2001) to red (2020). Prices generally do not move past the end of the contract month (day 0) as all information has been revealed. Most price volatility occurs one to two weeks prior to the start of the contract month and within the contract month. There are limited price changes more than two weeks before the start of the contract month, as limited information on weather shocks is revealed that the market could incorporate that far out. These flat prices depict market expectations of the climate before annual

⁴Due to the illiquidity of the weather market, we cannot guarantee that contracts were actually traded on days where the settlement price provided by CME does not change. In order to ensure that only traded prices were considered, we sometimes exclude time periods where the settlement price never changes, but the results are robust to the inclusion/exclusion of these days.

weather shocks are realized.

The main finding of our paper is clearly visible in the raw data: looking at futures prices a month before the start of the contract month (i.e., the left side of each graph), we see how prices for cooling degree day contracts in the left column are generally drifting upward over the years (color coded from blue to red), indicating an upward shift in the required amount of cooling as it gets hotter. By the same token, the right column shows prices for heating degree day contracts drifting downward over the years, indicating a downward shift in the expected amount of heating required.

While we do not have reliable volume data for the Bloomberg terminal time series, Appendix Figure A3 displays the fraction of days there has been a price change for the two-months period ranging from one month prior to the contract month to the contract month itself. It shows how the number of day-to-day price changes increase from 2001 to 2010, a likely indication that trading volume is picking up, before declining again until 2020. The decrease in volume is the reason that some of the original contract cities are no longer offered.

We contacted the Chicago Mercantile Exchange and obtained volume data for the subset of the contracts shown in Appendix Figure A4. Note the reduction in the number of lines representing contracts relative to Figure 1, our baseline dataset from Bloomberg. We display volume data for this subset in Appendix Table A1. Panel A shows volume by year. It is increasing from the start of weather derivatives in 2002 to 2008, when sales for winter and summer contracts combined topped US\$ 2 billion per year. Volume declines between 2008 and 2016, before another uptick in activity since 2017. Panel B aggregates the volume data by airport. Volume is highest in both cooling and heating degree days for New York (LGA) with a combined trading value of US\$3.9 billion. The second largest value for cooling degree days is for Dallas-Fort Worth (DFW), and for heating degree days at Chicago (ORD). The smallest value is for Sacramento at US\$ 0.2 billion. The combined traded value over all airports and years for this subset of the data (and hence a lower bound) exceeds US\$10 billion, a large enough amount to ensure that the market should efficiently incorporate weather information.

1.2. Weather Data

We pair the futures data with weather data: both weather station observations at the location associated with each contract as well as gridded climate model forecasts. For station data, we obtained the ID of the airport weather station underlying each contract and downloaded daily minimum and maximum temperatures from the National Oceanic and Atmospheric Administration's FTP server. We then computed the daily mean by averaging the minimum and maximum temperature before calculating the degree days for the 65°F bound as given in equations (1) and (2) above.

Climate projections were taken from the Coupled Model Comparison Project (CMIP) repository, which asks various modeling groups to simulate changing temperatures under comparable assumptions. We rely on the 5th round, i.e., the CMIP5 archive where these groups predicted climate trends from 2006 onwards. We obtain daily values from NASA NEX-GDDP, a dataset of 21 models that were spatially downscaled to a common 0.5° grid and select the grid cell in which the weather station is located. NASA NEX-GDDP has data for two scenarios. Representative Concentration Pathway (RCP) 4.5 assumes an additional energy flux of 4.5 Watts per meter square. This is a moderate warming scenario in which greenhouse gas emissions are reduced and radiative forcing stabilizes such that the global mean temperature increases by $1.8^{\circ}C$ ($3.2^{\circ}F$) by 2100. Note there is large spatial heterogeneity, and warming in the US is usually projected to be higher than the global average by a factor of roughly two. RCP 8.5, on the other hand, simulates major warming where emissions continue to rise such that there will be additional radiative forcing of 8.5 Watts per square meter resulting in a global mean temperature increase of 3.7° C by 2100. In the short term of our study period (2001-2020), however, both models give similar projections. The models are predicted to diverge further towards the end of the century as carbon emissions accumulate over time.

Appendix Figure A5 shows box plots for the number of cooling and heating degree days by month for the eight cities with weather futures contracts. The red line displays the weather station data and the blue line shows the climate model data. Both use data from 1950-2005, which was the historical baseline period in the CMIP5 archive. There is close alignment in the mean values as well variance around the means in both datasets. Recall that the climate models predict average temperature over the entire grid, and hence might differ from the observed temperature at any given point (i.e., weather station) if there is spatial heterogeneity. For example, a city's airport located close to a mountain might have a different temperature than that of the surrounding area when averaged over the entire grid.

We observe strong seasonality: more cooling degree days in the summer,

and more heating degree days in the winter. As expected, northerly cities (Chicago, Minneapolis, New York) have relatively more HDD and less CDD, while southerly cities (Atlanta, Dallas, Las Vegas) have less HDD and more CDD. Across the eight cities, there are very few occurrences of HDD in the summer months and CDD in winter months, which is why HDD futures contracts are not traded in summer and CDD contracts are not traded in winter.

Appendix Figure A6 plots the price of each weather derivative at the end of the contract month against the realized weather at the underlying weather station. The output closely follows the 45-degree line, demonstrating that the market is active enough to ensure that weather outcomes are fully priced in by contract close and that there are no arbitrage opportunities.

2. Empirical Analysis

We start by analyzing the timing of when futures prices capitalize weather shocks in Section 2.1. Forecasting and prediction skill of weather (shortterm) and climate (medium to long-term) are closely connected (Auffhammer et al., 2013). Climate models build on a foundation of short-term weather dynamics and the same underlying physical laws apply to the predictions of both weather and climate models. If market participants are accurately updating their longer-term beliefs based on climate warming trends, it would be expected that they also accurately update their short-term beliefs based on weather forecasts. The long-term trends are examined in Section 2.2.

2.1. Capitalization of Short-term Weather Shocks

Weather forecasts are widespread and freely available. There has been a sustained improvement in weather forecasting across all prediction ranges over recent decades. Bauer et al. (2015) present forecasting skill over time for weather anomalies, defined as deviations from the average climate, e.g., it is 10°F hotter today than what it is normally this time of the year. A score of 1 indicates that the forecasting model explains 100% of the year-to-year anomaly, while a score of 0 implies it cannot explain anything more than what is expected from the average conditions for the season.⁵ A 3-day forecasts

 $^{^{5}}$ The score is defined as 1 minus the ratio of the root mean squared error in the full weather forecast model relative to the root mean squared error of a baseline model that just predicts the average climatology. The authors state that "Values greater than 60%"

has improved from a skill of 80% in 1981 to 98% in 2014. On the other hand, a 10-day forecasts (not offered in 1981), increased from 30% in 1995 to 45% in 2014. Thus we would expect an inverted U-shape in terms of the impact of weather shocks on futures prices since long-term forecasts beyond 10 days have quickly diminishing value and since very short-term forecasts should have already been incorporated into prices given their certainty, aligning with Dorfleitner and Wimmer (2010) who find that weather forecasts only influence futures prices up to 11 days into the future. After this point, using the average outcome as prediction is just as good. As such, anticipated changes in weather around one week out should have the largest impact on current prices in an efficient market.

To test this, we estimate when weather shocks capitalize into futures prices for the eight airports in our sample. In a first step, we remove the seasonality to obtain weather shocks (anomalies), i.e., deviations from the average value that a rational market participant should expect. Specifically, we regress daily average temperature T_{ad} at airport a on day d on a constant α_a as well as flexible spline that is a function f of the day of the year.⁶ We also include a linear time trend γ_a in the year y(d) as the weather might be warming over time. The regression equation is

$$T_{ad} = \alpha_a + \beta_a f(d) + \gamma_a y(d) + \epsilon_{ad} \tag{3}$$

The estimated seasonality for each airport $\widehat{\beta}_a f(d)$ is shown in Appendix Figure A7. Years are color-coded to show the linear trend over time. The annual increase has not been uniform, e.g., Las Vegas warmed faster than Sacramento as there is a large distance between the red line (2020) and the blue line (2001). The weather shock on day d is simply the observed number of degree days $D(T_{ad})$ minus the degree days that would be expected at the predicted average climate according to the seasonality regression $D\left(\widehat{T_{ad}}\right)$.⁷

In a second step, we then regress the change in future prices Δp_{cd} for

indicate useful forecasts, while those greater than 80% represent a high degree of accuracy." ⁶To address leap years, we normalize the start of the year on January 1st to equal 0

and the end of the year on December 31st to equal 1. The five knots of the restricted cubic spline are at 0.05, 0.27, 0.50, 0.72, 0.95. This will give us four variables for the phase of the year f(d). We force the seasonality on December 31st to equal January 1st to guarantee continuity by running a constraint regression.

⁷While degree days are a nonlinear transformation when temperatures cross the truncation point at 65°F, the truncation is rarely observed, i.e., average daily temperatures

contract c on day d, i.e., the difference between the closing price to that of the previous close, on lags and leads of daily degree day shocks $\Delta D_{c[d+\tau]} = \left[D(T_{c[d+\tau]}) - D(\widehat{T_{c[d+\tau]}})\right]$ for days that fall within the contract month.⁸

$$\Delta p_{cd} = \alpha_c + \sum_{\tau=-7}^{21} \beta_\tau \left[D(T_{c[d+\tau]}) - D(\widehat{T_{c[d+\tau]}}) \right] + \epsilon_{cd} \tag{4}$$

One particularity about this regression is that while temperature data is available every day, prices are only available on trading days. As a result, the coefficient β_1 is for the sum of all weather shocks after the previous close and today's weather. All other β_{τ} use the weather on a single day, which is $\tau - 1$ days past the current close for leads ($\tau > 0$) and τ days before the previous close for lags ($\tau < 0$).⁹ The coefficient β_0 is normalized to be zero.

In line with the discussion on forecasting skill, future weather shocks should be capitalized into prices when weather forecasts can predict them, so we expect $\hat{\beta}_{\tau} > 0$ for the next two weeks $\tau \in [1, 14]$. After that point, weather forecasts become unreliable and not better than the average climate (Campbell and Diebold, 2005). Past weather is already known to market participants and hence the $\hat{\beta}_{\tau}$ should be zero for $\tau < 0$.

The left panel of Figure 2 shows individual coefficient estimates $\hat{\beta}_{\tau}$ with the expected hump-shaped pattern. The black line shows the point estimates with the 95% confidence band added in grey. As expected, past weather shocks have no effect on futures prices, while coefficients for the next two

are generally above 65° F in the summer and below 65° F in the winter. See Appendix Figure A5 that shows that there are very few heating degree days in the summer and cooling degree days in the winter. Expected degree days are close to degree days at the expected temperature. We obtain similar results whether we fit the seasonality separately for heating and cooling degree days or jointly for average temperature. We focus on the latter to estimate one unique seasonality rather than two separate regressions for summer and winter.

⁸A contract c specifies how many degree days will be observed at airport a in month m of year y, e.g., cooling degree days in June 2015 at LaGuardia airport. For a June contract, the weather shocks for days $d + \tau$ that are outside the month of June are set to zero as the price of a June contract is solely based on weather in June.

⁹For example, if day d is a Monday, β_1 includes the sum of the degree day shocks for Saturday, Sunday, and Monday, β_2 is the degree days shock on Tuesday, β_3 is the degree day shock on Wednesday, etc. On the other hand, β_{-1} is the degree day shock on the previous Friday.

weeks are generally positive as weather shocks get anticipated by the market and priced in prior to realization. Beyond day $\tau = 14$, the coefficients become insignificant again as weather forecasts beyond this time period are generally not better than the average climatology for the location. The right panel of Figure 2 makes this point more visible by plotting the *cumulative* sum of coefficients relative to $\tau = 0$, i.e., $\sum_{k=1}^{\tau} \hat{\beta}_k$ for $\tau > 0$ and $\sum_{k=\tau}^{-1} \hat{\beta}_k$ for $\tau < 0$. on the negative values of days $\tau < 0$ show no trend and the 95% confidence band includes zero. On the contrary, the line increases from 0 to 1 over the next two weeks as 100% of weather shocks get capitalized into the futures price. The curve flattens around 14 days into the future as weather forecasts become unreliable.

Appendix Figure A8 splits the regression into heating and cooling degree days and finds very similar relationships. The one exception is that the coefficient estimate $\widehat{\beta}_{-1}$ is positive for cooling degree days, which measure required cooling on the given day. This is not surprising as the daily maximum, which is crucial for the amount of required cooling, is generally observed in the late afternoon after the market closes, and hence would not get priced in until the next day.

One can invert the estimated relationship to obtain how future prices predict future weather. We can also run the opposite regression for illustrative purposes: do price changes in the futures market predict future weather shocks. In other words, are price changes a reliable weather forecast? We run the following inverse regression problem

$$\sum_{\tau=\tau_0}^{\tau_1} \left[D(T_{c[d+\tau]}) - D(\widehat{T_{c[d+\tau]}}) \right] = \alpha_c + \beta \Delta p_{cd} + \epsilon_{cd}$$
(5)

The regression results are shown in Appendix Table A2. Each entry is from a single regression of the sum of future weather shocks $\tau_0 - \tau_1$ days into the future on today's price change in the weather derivative. Different rows vary the time period $\tau_0 - \tau_1$. The first column pools all airports, the remaining eight columns run the regression by airport. We find that price changes predict weather shocks over the next two weeks, especially days 4-11, the sweet spot of weather forecasts, but cannot predict weather shocks more than two weeks in advance.¹⁰ The coefficient on weather shocks three weeks into the future (15-21 days) is not significant.

¹⁰The regression should be considered with caution as the reverse regression problem

2.2. Capitalization of Long-term Weather Trends

2.2.1. Linear Time Trends

We now turn to our main analysis of market expectations of climate change. With weather futures, we must be careful to separate price changes driven by short-term weather forecasts and those reflecting longer-term market beliefs on warming. Some shocks are partially forecastable over the course of months based on oceanic-atmospheric phenomena like El Niño - Southern Oscillation (ENSO) or the North Atlantic Oscillation (NAO). Ideally we would use futures prices quoted well before the contract's delivery month. However, for the same reason that weather is challenging to forecast far in advance, trading does not pick up until close to the contract delivery month, and early-dated prices may not be representative of the market's true expectation given the illiquidity.

Balancing these two tradeoffs, our baseline model uses average futures prices $\overline{p_{amy}}$ of contract c for airport a in month m of year y. The average price is taken the fourth week (28-22 days) prior to the start of a contract month, e.g., the average price between June 3 and June 9, 2015 for a July 2015 CDD contract in Atlanta. This ensures that prices reflect future expectations and not contemporaneous weather as confirmed in the previous section. In the baseline we pool four summer months (June – September) in the cooling degree days regression, and five winter months (November - March) in the heating degree days regression. We fit a simple linear trend in the year y after including airport-by-month fixed effects α_{am} , e.g., a fixed effect for June contracts in Atlanta. We cluster the error terms for a particular month m as they might be subject to the same common weather shock.

$$\overline{p_{amy}} = \alpha_{am} + \beta y + \epsilon_{amy} \tag{6}$$

Table 1 shows the predicted annual change $\hat{\beta}$ in column (1a). Panel A shows that on average prices increased by \$2.4 per year for each of the four summer months, June to September, or \$10 per year for the combined four-month period. This annual increase is statistically significant at the 1% level. Since

can lead to biased coefficients. In the climate literature, the width of tree rings is often taken as a temperature proxy for past temperatures before weather stations were available. As Auffhammer et al. (2015) point out, weather influences tree rings. Running the inverse regression where temperature is regressed on tree rings will lead to biased coefficients and predictions with artificially low variance.

our dataset spans 20 years, the price for a cooling degree day contract increased by roughly \$50 since 2001 for each of the monthly summer contracts. Recall that the payout of the weather derivatives has a multiple of 20, so a price increase of \$50 implies a change in payout by \$1000 over our sample period. Panel B shows that the price for heating degree day contract declined on average by \$1 per year, or \$5 for the five-month span November to March. It is significant at the 5% level.

Columns (b)-(d) replicate an equivalent analysis using the weather station and climate model data. The dependent variable is no longer the futures price $\overline{p_{amy}}$ but the number of degree days at the weather station or climate grid. Columns (1b)-(1d) hold the set of observations constant, i.e., only include the months where we have futures price data. Column (1b) uses the observed degree days for the contract month from the underlying station data as the dependent variable. The observed trends (annual changes) are larger in magnitude with an increase of 3 cooling degree days per year during the summer and a decrease of 2 heating degree days during the winter. The standard errors are much larger given the greater year-to-year swings stemming from random weather fluctuations. As a result, trends in observed weather are not significantly different from those anticipated by the futures market as shown in column (1a). The smaller standard errors for futures prices relative to the station-level data also suggest that we are correctly measuring longerterm market expectations and not just annual weather realizations, which are much noisier. Columns (1c) and (1d) show trends in the NASA-NEX dataset averaged across the 21 climate models for the RCP4.5 and RCP8.5 scenarios, respectively, as dependent variable for each month.

While columns (1a)-(1d) intentionally keep the set of city-year observations constant, columns (2a)-(2d) replicate the analysis with different subsets of the data. First, to address concerns about market illiquidity, Column (2a) excludes observations where there was no price change in the week over which prices are averaged, i.e., the fourth week prior to the start of the contract month in our baseline specification. This exclusion reduces the sample size by roughly half but results in point estimates of similar magnitude to those in column (1a). The time trends are statistically different from zero, and not statistically different than the estimates in column (1a). The reduction in observations in column (2a) can be explained by the fact that we are taking average prices over the fourth week prior to the start of the contract month, a period when limited information about the eventual weather outcome is available beyond the climate normals. We hence do not expect many price changes, which happen when new information gets incorporated. Nevertheless, it is reassuring that the time trends are similar whether there is a price change (and hence update) or not. Second, to address concerns about the endogeneity of this market, e.g., if contracts are traded more in particularly cold or hot years as firms realize they need a hedge, columns (2b)-(2d) use all available months with weather station and climate model data (even if no future price data existed) and again find very similar annual changes to those in columns (1b)-(1d).

So far we have pooled all months of a season as well as each airport into a single regression. Appendix Tables A3 and Table A4 relax this assumption to examine heterogeneity by geography and month. Each table presents the pooled results from Panels A and B of Table 1 in the top row of the corresponding panel for reference. Appendix Table A3 allows time trends to differ by airport while still pooling all summer or winter months, and Appendix Table A4 allows time trends to differ by month while still pooling all airports. We observe some differences by airport, e.g., in column (1a) the futures market predicts warming in Las Vegas above the national average in both winter and summer, and below average warming in Chicago and Sacramento in the summer, all at the 1% significance level. All significant time trends have the same sign as the national analysis, i.e., more cooling degree days in the summer and fewer heating degree days in the winter, although the winter time trends sometimes become insignificant, especially in the northeastern subset of airports (CVG, LGA, MSP, ORD).¹¹ In column (1b), none of the time trends in weather station data differ significantly by airport, although they are estimated with more noise due to the large year-toyear variability. In columns (1c)-(1d), the climate models show below-average warming in Sacramento in the RCP4.5 data. In summary, while there are small differences, there does not seem to be a systematic significant difference by airport.

The story is different when examining heterogeneity by month in Appendix Table A4. Futures prices show a significant positive annual increase for February heating degree days, suggesting an expectation of colder temperatures. It is highly significant at the 1% level. This finding is primarily driven by regional heterogeneity. Appendix Figure A10 shows time trends

¹¹The winter time trend for Sacramento is also insignificant, although it is less traded than other contracts and the summer time trend was also closer to zero.

per month after separating the eight airports into a northeastern quadrant (CVG, LGA, MSP, ORD) and the rest, mostly in the south and southwest (ATL, DFW, LAS, SAC). The February cooling trend (positive increase in heating degree days) is only observed for the northeastern quadrant in the futures data. Since we are splitting the sample further, the estimated time trends become less precisely estimated, but February cooling is neither supported by recent weather observations nor climate runs in the CMIP5 archive. All other winter months either show a significant negative time trend or an insignificant time trend in heating degree days.

The futures market may be incorporating recent information about a shift of the polar vortex that was not available at the time of CMIP5. Recent studies suggest that melting ice sheets destabilize the jet stream, leading to an increased frequency of stable weather patterns bringing cold arctic air to Europe and North America (Francis and Vavrus, 2015). Zhang et al. (2016) conclude that the "Arctic polar vortex shifted persistently towards the Eurasian continent and away from North America in February over the past three decades. [...] Our analysis reveals that the vortex shift induces cooling over some parts of the Eurasian continent and North America which partly offsets the tropospheric climate warming there in the past three decades." Kim et al. (2014) note that "the mechanism that links sea-ice loss to cold winters remains a subject of debate," so it remains an active topic of research. One crucial paper for our analysis is Charlton and Polvani (2007), who more generally examine a phenomenon called stratospheric sudden warming (SSW) and its relationship to the troposphere, specifically the polar vortex. The authors note that "given the prominent role of SSW events, it is somewhat surprising that relatively few attempts have been made to establish a comprehensive climatology of SSWs." The authors proceed to do so in two accompanying articles in the Journal of Climate in 2007 and operationalize how SSW events in January and February in the stratosphere can influence weather in the troposphere.¹² A fully rational market would incorporate this

¹²The authors write: "A useful analogy might be drawn at this point with the atmosphere–ocean system: in the same way as understanding and successfully modeling the El Niño–Southern Oscillation phenomenon is of primary importance for the atmosphere–ocean system, understanding and successfully modeling stratospheric sudden warming events is of primary importance for the stratosphere–troposphere system." El Niño–Southern Oscillation similarly allows a weather forecast with a lead time of more than four weeks, i.e., the futures data might be picking up relevant information of how a

new finding, an issue we return to in the next subsection where we present nonlinear trends and find an uptick in the 2007-2008 winter, i.e., right after publication.

Before we do so, Panel C row of Table 1 replicates the analysis for heating degree days from Panel B after excluding February contracts for the four northeastern airports. While the exclusion has very limited effect on the estimated annual decrease in monthly heating degree days for the regression using weather station data or climate model outputs in columns (b)-(d), it changes the coefficient on the annual decrease in futures prices in column (a), making it much more closely aligned with the annual changes in observed weather and climate model output.

We present a final sensitivity check of the observed futures price trends to the window over which the prices are averaged in Table 2. Our baseline uses prices that are averaged over the fourth week prior to the start of the contract month. Prices at this point are mostly stable as shown in Figure 1 because new information on the annual shocks are not yet available. The six columns in Table 2 replicates the analysis by averaging anywhere between one to six weeks prior to the start date of the contract month. The time trend on cooling degree days in Panel A is completely insensitive to the chosen time window and very stable around an additional 2.4 cooling degree days per year for each of the summer months. The time trend on heating degree days in Panel B and Panel C are very similar whether we average prices six, five, or four weeks in advance of the start of the contract month. There is a slight uptick as we get closer to the start date of the contract month suggesting an even larger annual decline, although the difference is not significant given the larger standard errors. The overall robustness of the relationship across the time periods supports the idea that markets expected a consistent increase in the need for cooling in the summer and decrease in the need for heating in the winter.

2.2.2. Nonlinear Trends

Figure 3 relaxes the linearity assumption of the time trend and instead plots a semi-parametric regression of the residuals after removing airportby-month fixed effects α_{am} in equation (6) to account for different average

year's weather is shifting. Appendix Section A2 finds that oceanic indices like El Niño are not a major factor of the observed warming trend.

monthly climates (i.e., June in Atlanta is hotter than June in Minneapolis). We use restricted cubic splines to allow for more flexible trends.¹³ The lines in green, red, blue and cyan correspond to the variables listed in columns (1a)-(1d) of Table 1 (Panel A for cooling degree days and Panel C for heating degree days), respectively, i.e., residuals from the weather futures prices, weather station outcomes, and climate forecasts under RCP4.5 and RCP8.5.

The futures prices and climate model output show a steady upward trend in cooling degree days and a downward trend in heating degree days. The trends on the weather station data (red lines) are less smooth for both cooling and heating, partly because of the noisiness inherent in year-to-year swings in weather realizations that are larger than predicted average outcomes in the other datasets. For example, the winter 2017/2018 was especially warm, leading to a sharp drop in heating degree days for that year. There also seems to be a short-term plateau in the observed cooling trend around 2010, but the long-term effects over the 20-year period are similar across datasets. For both cooling and heating, the green lines showing futures price trends closely follow the cyan and blue lines of the climate model projections, and not the red lines. This suggests that beliefs are not myopically updated based on recently observed weather but rather tied to the smooth warming trend projected by climate models and observed in longer-term station data.

In the previous section we found a statistically significant cooling trend in February futures prices for the four northeastern airports. To show this, we again relax the linearity assumption in Figure 4 and plot the residuals of February prices four weeks before the start of the contract month after removing airport fixed effects. We then add a trend line using the same restricted cubic splines in time as well as the 95% confidence band. We observe an almost linear uptick in residuals between 2007 and 2012, which is consistent with the publication of Charlton and Polvani (2007) a study in the premier climatology journal that presents a novel *comprehensive* climatology to predict the "Polar Vortex." While we cannot be sure of when the market

¹³The spline knots are at 2003, 2008, 2013 and 2018. Appendix Figure A11 presents locally-weighted lowess regression of the same residuals. Specifically, we apply STATA's lowess command to the annual average of the residuals. We first average the monthly residuals per year as a locally weighted regression with several observation in the same year would need to arbitrarily pick which of the month to include in the local average. The point estimates are similar to the spline regression, which we utilize going forward because they allow for the construction of confidence bands.

became aware of various findings in the scientific literature, it is striking that starting around 2007 February becomes the only month where the futures markets predicts a cooling in the short-term that eventually diminishes as anthropogenic warming becomes dominant.

2.2.3. Comparing Spatial and Temporal Heterogeneity

The previous section has shown that the market incorporated a unique sub-seasonal cooling dynamic for part of the US. We extend this type of analysis further by examining whether the observed heterogeneity in the time trend mostly aligns with climate model output or observed station-level trends. This allows us to contrast whether futures markets reflect knowledge about climate model forecasts or simply assume the continuation of observed time trends. While all datasets show similar average time trends, the spatial and temporal heterogeneity varies.

Intuitively, if traders rely mostly on recent observed trends, we would expect that airports and/or contract months that show larger than average warming in the station-level data between November 2001 and March 2020 would also have larger than average annual changes in futures prices as well. On the other hand, if market participants mostly respond to climate model forecasts, we would observe the distribution of time trends to more closely align with what is observed in the climate model output.

To test this we estimate time trends β_{am} that are airport and month specific instead of the common trend β used in equation (6):

$$\overline{p_{amy}} = \alpha_{am} + \beta_{am}y + \epsilon_{amy} \tag{7}$$

We run this model with futures price data to obtain $\widehat{\beta_{am}^f}$, observed weather station data to obtain $\widehat{\beta_{am}^s}$, and the climate model output under RCP4.5 to obtain $\widehat{\beta_{am}^{4.5}}$ and RCP8.5 to obtain $\widehat{\beta_{am}^{8.5}}$).¹⁴ In a second step we then regress the estimated time trend in the futures data on the other trends

$$\widehat{\beta_{am}^f} = \alpha_0 + \alpha_s \widehat{\beta_{am}^s} + \alpha_{4.5} \widehat{\beta_{am}^{4.5}} + \alpha_{8.5} \widehat{\beta_{am}^{8.5}} + \epsilon_{am} \tag{8}$$

¹⁴We use all monthly observations from November 2001-March 2020 in the station and climate model data, even if the futures data is not available. Since the weather station data is more variable (it measures actual outcomes versus averages among climate models), we include as many observations as possible in order not to unfairly penalize the station-level data by making the time trend more variable.

If market participants are just incorporating the average for each airport-bymonth, we would only expect the constant α_0 to be significant, as it picks up the common average. On the other hand, if futures prices incorporate the observed heterogeneity in time trends found in the station-level data or climate model output, we would expect α_s , $\alpha_{4.5}$, or $\alpha_{8.5}$ to be significant.

It should be noted that it is much harder to predict spatial heterogeneity in warming than it is to predict average trends because of all the localized feedback loops of the climate system. The average trend is given by a simple balance of energy calculation. For example, if one increases the burner under a pot of water the average temperature will increase, but it is much harder to predict where this extra energy will show up and how it will spread across the volume of water. Similarly, changes in wind patterns might lead to higher warming in some areas while reducing it in others (Hsiang and Kopp, 2018). February cooling due to the polar vortex over eastern North America goes hand-in-hand with higher-than-expected warming in the Arctic. Cooling in East Coast cities does not refute that the globe is warming, which it is in total, but rather reflects the uncertainty on where the extra energy manifests as jet streams shift.

The results are given in Table 3. Columns (a)-(c) include each estimated time trend in the weather/climate data one at a time, while columns (d) jointly include all three. Columns (1a)-(1d) include all 72 airport-month combinations of the 8 airports and 9 months: June - September for cooling degree days in the summer and November - March for heating degree days in the winter. Columns (2a)-(2d) exclude February for the four northeastern airports for a total of 68 observations.

Panel A pools all observations from November 2001-March 2020 in the estimation of the β_{am} . The coefficient on the climate model output in columns (b) and (c) is consistently larger than for the heterogeneity actually observed in the weather station data over the same period. When we include all three in column (d), they are no longer individually significant given the high degree of multicollinearity, but climate model output under the RCP4.5 scenario has the largest point estimate.

Panel B and Panel C limit the observations to 2006-2020 and 2011-2020, respectively, in the calculation of the trends β_{am} . The reason is twofold: first, climate models in the CMIP 5 archive used 1950-2005 as the baseline to calibrate their models. By limiting the data to a period past 2005, the model should predict completely out-of-sample. Note however that we are using the actual observed climate trends from the weather station data $\widehat{\beta}_{am}^s$.

so the climate model would simply incorporate some of the information that is in the station level data. Since it took climate modeling groups several years to run the models before they were posted, Panel C further limits the time window to after 2010. Second, the pace of global warming slowed between 1998-2012, and then picked up again around 2012.

Both Panel B and C show that the spatial heterogeneity in trends in the futures data is better aligned with the heterogeneity in the climate model output rather than the trend at the underlying weather station. For this subinterval of accelerated warming, the heterogeneity found in RCP8.5 is a better predictor than RCP4.5. On the one hand, this is not surprising as the early 2000s mostly relied on climate projections from the IPCC Fourth Assessment Report that did not include the more rapid RCP8.5 scenario. On the other hand, as we have argued above, the futures market was quick to pick up on scientific advances related to the polar vortex. Since the IPCC reports are based on published studies, much of the underlying theory might have also been available to interested parties in the early 2000s. We lack a credible proxy for when information is received by the market, so we cannot directly test when market participants update their view on which climate model to follow.

It is noteworthy that across all the time periods considered in Panels A-C, the heterogeneity in the futures price trends more closely mirror climate models than the eventual weather realizations. Combined with the uptick in February futures prices that is not supported by station-level observations, we conclude that market participants are utilizing climate models, or some related source of information, to update their beliefs on future weather rather than just projecting forward historical trends. Moreover, as Appendix Section A3 shows, previous warming trends in the early part of the 20th century have plateaued, and simply forecasting that past trends will continue rather than utilizing climate model forecasts would be a risky endeavor for investors.

Warming trends are predicted to diverge further out in the future as shown in Figure 5, which displays climate model output through 2100. We again remove airport-by-month fixed effects and then average the residuals over the four summer months (June - September) or the five winter months (November - March). The top row again shows cooling degree days, while the bottom row shows heating degree days. The left column shows nonparametric warming paths under the RCP4.5 scenario, while the right column uses RCP8.5. For example, the reduction in heating degree days in Minneapolis under the RCP8.5 scenario (bottom right graph) is almost twice as large as for Atlanta.

3. Conclusion

To the best of our knowledge, this paper is the first to utilize a direct measure of climate change expectations as derived from weather-based futures contracts. The evidence shows that financial markets incorporate warming trends that are consistent with climate model projections. We find that the market has been accurately pricing in a warming climate, and that this began occurring at least since the early 2000s when the weather futures markets were formed. The market also seems to price in recent scientific findings like the polar vortex that can lead to February cooling over the eastern US, an effect neither present in the CMIP5 archive nor detectable in recent weather station observations.

Our findings have direct implications for firms and financial markets. Recent studies have highlighted how the valuations of companies and entire industries are sensitive to weather fluctuations and climate change risk. Since efficient and profit-maximizing behavior requires an accurate assessment of predicted warming, weather markets can provide companies with pertinent information on future weather and climate trends, as well as a hedge against potential lost profit. Relatedly, our findings may have relevance to climate adaptation. Adaptation requires that agents form beliefs about whether the climate is changing and the extent to which it is. This paper suggests that agents, at least those participating in weather markets, have been updating their beliefs that summers are getting hotter and winters colder.

There are also policy implications of our findings, especially since some politicians still question the existence and extent of climate change. The observed annual trend in futures prices shows that the supposedly-efficient financial markets agree that the climate is warming. At least so far, climate models have been very accurate in predicting the average warming trend that's been observed across the US. While our results do not conclude that the market believes the warming is human-induced, per se, anyone doubting the observed warming trend can always attempt to profit from that belief by betting against the observed trend. The price of a summer month cooling degree days contract, for example, has increased by roughly \$50 over 20-year sample period. Since the payout of the financial derivative has a multiplier of 20, this implies an additional \$1,000 in value is on the table per contract. When money is on the line, it is hard to find parties willing to bet against the scientific consensus.

References

- Addoum, J.M., Ng, D.T., Ortiz-Bobea, A., 2020. Temperature shocks and establishment sales. Review of Financial Studies 33, 1331–1366. doi:https://doi.org/10.1093/rfs/hhz126.
- Akerlof, K., Maibach, E.W., Fitzgerald, D., Cedeno, A.Y., Neuman, A., 2013. Do people personally experience global warming, and if so how, and does it matter? Global Environmental Change 23, 81–91. doi:https://doi.org/10.1016/j.gloenvcha.2012.07.006.
- American Association for the Advancement of Science, December 9, 2006. Aaas board statement on climate change.
- Anttila-Hughes, J.K., 2016. Financial market response to extreme events indicating climatic change. European Physical Journal Special Topics 225, 527–538. doi:https://doi.org/10.1140/epjst/e2015-50098-6.
- Auffhammer, M., 2018. Quantifying economic damages from climate change. Journal of Economic Perspectives 32, 33–52. doi:https://doi.org/10.1257/jep.32.4.33.
- Auffhammer, M., Hsiang, S.M., Schlenker, W., Sobel, A., 2013. Using weather data and climate model output in economic analyses of climate change. Review of Environmental Economics and Policy 7, 181–198. doi:https://doi.org/10.1093/reep/ret016.
- Auffhammer, M., Li, B., Wright, B., Yoo, S.J., 2015. Specification and estimation of the transfer function in dendroclimatological reconstructions. Environmental and Ecological Statistics 22, 105–126. doi:https://doi.org/10.1007/s10651-014-0291-6.
- Bauer, P., Thorpe, A., Brunet, G., 2015. The quiet revolution of numerical weather prediction. Nature 525, 47–55. doi:https://doi.org/10.1038/nature14956.

- Bernstein, A., Gustafson, M.T., Lewis, R., 2019. Disaster on the horizon: The price effect of sea level rise. Journal of Financial Economics 134, 253–272. doi:https://doi.org/10.1016/j.jfineco.2019.03.013.
- Burke, M., Emerick, K., 2016. Adaptation to climate change: Evidence from us agriculture. American Economic Journal: Economic Policy 8, 106–140. doi:https://doi.org/10.1257/pol.20130025.
- Campbell, S.D., Diebold, F.X., 2005. Weather forecasting for weather derivatives. Journal of the American Statistical Association 100, 6–16. doi:https://doi.org/10.1198/016214504000001051.
- Carney, M., 2015. Breaking the tragedy of the horizon climate change and financial stability. Speech at Lloyd's of London by the Governor of the Bank of England and Chairman of the Financial Stability Board .
- Charlton, A.J., Polvani, L.M., 2007. A new look at stratospheric sudden warmings. part i: Climatology and modeling benchmarks. Journal of Climate 20, 449–469. doi:https://doi.org/10.1175/JCLI3996.1.
- Dell, M., Jones, B.F., Olken, B.A., 2014. What do we learn from the weather? the new climate-economy literature. Journal of Economic Literature 53, 740–798. doi:https://doi.org/10.1257/jel.52.3.740.
- Deryugina, T., 2013. How do people update? the effects of local weather fluctuations on beliefs about global warming. Climatic Change 118, 397– 416. doi:https://doi.org/10.1007/s10584-012-0615-1.
- Dorfleitner, G., Wimmer, M., 2010. The pricing of temperature futures at the chicago mercantile exchange. Journal of Banking and Finance 34, 1360–1370. doi:https://doi.org/10.1016/j.jbankfin.2009.12.004.
- Francis, J.A., Vavrus, S.J., 2015. Evidence for a wavier jet stream in response to rapid arctic warming. Environmental Research Letters 10, 1–12. doi:http://dx.doi.org/10.1088/1748-9326/10/1/014005.
- F.W., Hong, H., Li, Xu, J., 2019. Climate risks and mar-Journal of 265 - 281.ket efficiency. Econometrics 208.doi:https://doi.org/10.1016/j.jeconom.2018.09.015.

- Howe, P.D., Mildenberger, M., Marlon, J.R., Leiserowitz, A., 2015. Geographic variation in opinions on climate change at state and local scales in the usa. Nature Climate Change 5, 596–603. doi:https://doi.org/10.1038/nclimate2583.
- Hsiang, S., Kopp, R.E., 2018. An economist's guide to climate change science. Journal of Economic Perspectives 32, 3–32. doi:https://doi.org/10.1257/jep.32.4.3.
- Kala, N., 2019. Learning, adaptation, and climate uncertainty: Evidence from indian agriulcture. Working Paper.
- Kaufmann, R.K., Mann, M.L., Gopal, S., Liederman, J.A., Howe, P.D., Pretis, F., Tang, X., Gilmore, M., 2017. Spatial heterogeneity of climate change as an experiential basis for skepticism. Proceedings of the National Academy of Sciences 114, 67–71. doi:https://doi.org/10.1073/pnas.1607032113.
- Kim, B.M., Son, S.W., Min, S.K., Jeong, J.H., Kim, S.J., Zhang, X., Shim, T., Yoon, J.H., 2014. Weakening of the stratospheric polar vortex by arctic sea-ice loss. Nature Communications 5, 1–8. doi:https://doi.org/10.1038/ncomms5646.
- Leiserowitz, A., Maibach, E., Roser-Renouf, C., Rosenthal, S., Cutler, M., 2017. Climate Change in the American Mind: November 2016. Yale University and George Mason University.
- Li, Y., Johnson, E.J., Zaval, L., 2011. Local warming: Daily temperature change influences belief in global warming. Psychological Science 22, 454– 459. doi:https://doi.org/10.1177/0956797611400913.
- McGlade, C., Ekins, P., 2015. The geographical distribution of fossil fuels unused when limiting global warming to 2c. Nature 517, 187–190. doi:https://doi.org/10.1038/nature14016.
- Meng, K.C., 2017. Using a free permit rule to forecast the marginal abatement cost of proposed climate policy. American Economic Review 107, 748–784. doi:https://doi.org/10.1257/aer.20150781.
- Moore, F., Obradovich, N., Lehner, F., Baylis, P., 2019. Rapidly declining remarkability of temperature anomalies may obscure public perception of

climate change. Proceedings of the National Academy of Sciences 116, 4905–4910. doi:https://doi.org/10.1073/pnas.1816541116.

- Myers, T.A., Maibach, E.W., Roser-Renouf, C., Akerlof, K., Leiserowitz, A.A., 2013. The relationship between personal experience and belief in the reality of global warming. Nature Climate Change 3, 343–347. doi:https://doi.org/10.1038/nclimate1754.
- Neidell, M., 2009. Information, avoidance behavior, and health: The effect of ozone on asthma hospitalizations. Journal of Human Resources 44, 450–478. doi:https://doi.org/10.3368/jhr.44.2.450.
- Rosenzweig, M.R., Udry, C., 2014. Rainfall forecasts, weather, and wages over the agricultural production cycle. American Economic Review 104, 278–283. doi:https://doi.org/10.1257/aer.104.5.278.
- Severen, C., Costello, C., Deschênes, O., 2018. A forward-looking ricardian approach: Do land markets capitalize climate change forecasts? Journal of Environmental Economics and Management 89, 235–254. doi:ttps://doi.org/10.1016/j.jeem.2018.03.009.
- Shrader, J., 2020. Expectations and adaptation to environmental risks. SSRN Working Paper 3212073 doi:http://dx.doi.org/10.2139/ssrn.3212073.
- Zaval, L., Keenan, E.A., Johnson, E.J., Weber, E.U., 2014. How warm days increase belief in global warming. Nature Climate Change 4, 143–147. doi:https://doi.org/10.1038/nclimate2093.
- Zhang, J., Tian, W., Chipperfield, M.P., Xie, F., Huang, J., 2016. Persistent shift of the arctic polar vortex towards the eurasian continent in recent decades. Nature Climate Change 6, 1094–1099. doi:https://doi.org/10.1038/nclimate3136.



Graphs display the time series of futures prices around maturity. Day 0 is the end of the month on which the weather derivative is based, e.g., day 0 for a June contract is June 30. Top row is for New York LaGuarida airport (LGA), bottom row for Dallas-Fort Worth (DFW). The left column shows cooling degree days for July while the right columns show heating degree days for December. Years are color coded as shown in the bottom legend. Price series that are flagged for quality issues are shown as dashed lines instead of solid lines. The grey shaded area shows the period over which we average futures prices in our baseline specification to derive market expectations, which is four weeks before the start of the month. Contracts for the remaining airports and months are shown in Appendix Figure A2.



This figure displays the results from a distributed lag model. Daily futures price changes Δp_{cd} for contract c on day d are regressed on 21 leads and 7 lags of weather shocks $\Delta \widehat{D_{c[d+\tau]}}$, i.e., the difference compared to the average climate on day $d+\tau$. The regression equation is $\Delta p_{cd} = \alpha_c + \sum_{\tau=-7}^{21} \beta_{\tau} \Delta \widehat{D_{c[d+\tau]}} + \epsilon_{cd}$. The left graph shows the estimated coefficient $\widehat{\beta_{\tau}}$ for the weather shock on a particular lead/lag τ . Negative values of τ on the horizontal axis indicate weather occurring on an earlier day, i.e., in the past, while positive values depict weather at a future date. The right graph shows $\sum_{k=1}^{\tau} \widehat{\beta_k}$, i.e., the cumulative sum of coefficients form day 0 onwards for positive values of the horizontal axis and $\sum_{k=\tau}^{-1} \widehat{\beta_k}$, i.e., the cumulative sum of coefficients before day 0 for negative values of the horizontal axis. The regression pools cooling degree day contracts in June - September and heating degree day contracts for November - March. The estimated coefficients for

and heating degree day contracts for November - March. The estimated coefficients for leads $\tau > 1$ and lags $\tau \leq -1$ are on the weather shock for one day, but the coefficient shown for $\tau = 1$ is for the sum of shocks from today to the previous close given that futures are not traded every day.



Figure 3: Nonlinear Time Trends in Futures Prices and Weather

This figure estimates nonlinear time trends using restricted cubic splines in time (knots at 2003, 2008, 2013, and 2018) on the residuals, which are obtained by subtracting airport-by-month fixed effects $\widehat{\beta_{am}}$ among the eight airports and four summer months (June - September) in the left graph or eight airports and five winter months (November-March) in the right graph, excluding February for the four northeastern airports. The green line uses futures prices four weeks before the start of the contract month. The red line shows the results for the observed weather station data. The blue lines use model output from NASA NEX-GDDP. In each case we subtract the average for the airport and month (i.e., airport-by-month fixed effect). The horizontal axis reports the year a season ends, e.g., winter 2001/2002 is recorded as 2002. The 95% confidence bands are added as shaded areas.



Figure 4: Nonlinear Time Trend in February Futures at Northeastern Airports

This figure estimates nonlinear time trends using restricted cubic splines in time (knots at 2003, 2008, 2013, and 2018) on the residuals of February contracts among the four airports in the northeastern quadrant in Appendix Figure A1. Residuals are obtained and after removing airport fixed effects and are displayed for the four airports. The green line uses futures prices four weeks before the start of the contract month. The 95% confidence band is added as shaded area.



Figure 5: Predicted Change in Degree Days in Climate Models

This figure shows nonparametric time trends by airport averaged over the 21 climate models in the NASA NEX-GDDP database. The y-axis gives the predicted average change in monthly cooling or heating degree days. Top row shows the results for the change in monthly cooling degree days in the summer (June - September) and the bottom row for the change in monthly heating degree days in the winter (November - March). Left column uses the predictions under the RCP4.5 scenario, while the right column uses RCP8.5. Specifically, a nonparametric lowess regression is fit to the annual average of the monthly residuals after removing airport-by-month fixed effects.

	Table 1: Linear Time Trends in Degree Days								
	(1a)	(1b)	(1c)	(1d)	(2a)	(2b)	(2c)	(2d)	
	Panel A: CDD June - September								
Trend	2.432^{***}	2.998^{***}	2.286^{***}	2.774^{***}	2.148^{***}	2.676^{***}	2.167^{***}	2.432^{***}	
	(0.160)	(0.887)	(0.169)	(0.174)	(0.330)	(0.772)	(0.173)	(0.160)	
Observations	522	522	522	522	222	576	576	522	
	Panel B: HDD November - March								
Trend	-1.000^{**}	-2.081	-1.662^{***}	-1.854^{***}	-1.175^{**}	-1.677	-1.734^{***}	-1.000**	
	(0.415)	(1.723)	(0.354)	(0.370)	(0.573)	(1.524)	(0.314)	(0.415)	
Observations	676	676	676	676	322	760	760	676	
Panel C: HDD November - March (excluding February in Northeast)									
Trend	-1.719^{***}	-1.856	-1.527^{***}	-1.710^{***}	-2.224^{***}	-1.610	-1.643^{***}	-1.719^{***}	
	(0.384)	(1.731)	(0.362)	(0.336)	(0.478)	(1.529)	(0.329)	(0.384)	
Observations	604	604	604	604	281	684	684	604	
Data	Futures	Station	RCP4.5	RCP 8.5	Futures	Station	RCP4.5	RCP 8.5	
Years	Common	Common	Common	Common	Traded	All	All	All	

This table reports the estimated annual increase/decrease in degree days $\hat{\beta}$. Each entry is from a separate regression where degree days D_{amy} at airport *a* for month *m* in year *y* are regressed on airport-by-month fixed effects as well as a linear time trend: $D_{amy} = \alpha_{am} + \beta y + \epsilon_{amy}$. Panel A regresses cooling degree days (CDD) for the summer months June - September, while panel B and C use heating degree days (HDD) for November - March. Panel C excludes February for the four northeastern airports in Appendix Figure A1. The data set ranges from winter 2001/2002 through winter 2019/2020. Columns (a) uses the average futures price $\overline{p_{amy}}$ four weeks before the start of each contract month, e.g., the average price between May 4 and May 10 for a June contract. Columns (b) uses observed station-level data for the month, while columns (c) and (d) use climate model forecasts in the NASA NEX-GDDP database under the RCP4.5 and RCP8.5 scenarios for the month. Columns (1a)-(1d) estimate the trends for a consistent set of observations where futures data are available. Columns (2a)-(2d) conduct sensitivity checks to the included years. Columns (2a) exclude contracts where the price did not change during the fourth week preceding the contract month. Columns (2b)-(2d) include all years even if futures data is not available. Stars indicate significance levels: * 10%, ** 5%, *** 1%.

မ္မ

	(1)	(2)	(3)	(4)	(5)	(6)		
	Panel A: CDD June - September							
Trend	2.451^{***}	2.428^{***}	2.432***	2.385^{***}	2.431^{***}	2.448^{***}		
	(0.147)	(0.142)	(0.160)	(0.189)	(0.239)	(0.314)		
Observations	520	522	522	522	522	522		
	Panel B: HDD November - March							
Trend	-0.905**	-0.900**	-1.000**	-1.224^{***}	-1.356^{***}	-1.628^{**}		
	(0.408)	(0.405)	(0.415)	(0.431)	(0.482)	(0.697)		
Observations	672	676	676	676	676	676		
Panel C: HDD November - March (exl. Feb in NE)								
Trend	-1.656^{***}	-1.642^{***}	-1.719^{***}	-1.908^{***}	-2.077^{***}	-2.202***		
	(0.371)	(0.363)	(0.384)	(0.414)	(0.465)	(0.727)		
Observations	600	604	604	604	604	604		
Airport FE	Yes	Yes	Yes	Yes	Yes	Yes		
Weeks Prior	6	5	4	3	2	1		

Table 2: Sensitivity of Linear Trend to When Expectations are Taken

This table shows a sensitivity analysis of column (1a) of Table 1, now column (3), to the time window over which future prices are averaged to evaluate expectations. The last row displays how many weeks prior to the start of the contract month the futures prices are averaged over, ranging from one to six weeks. Stars indicate significance levels: * 10%, ** 5%, *** 1%.

	(1a)	(1b)	(1c)	(1d)	(2a)	(2b)	(2c)	(2d)	
	Panel A: All Years								
Trend at Weather Station	0.251^{**}			0.135	0.302^{***}			0.137	
	(0.101)			(0.113)	(0.100)			(0.105)	
Trend in CMIP5 - RCP4.5		0.628^{***}		0.432		0.840^{***}		0.627	
		(0.155)		(0.474)		(0.137)		(0.426)	
Trend in CMIP5 - RCP8.5			0.501^{***}	0.066			0.666^{***}	0.080	
			(0.126)	(0.392)			(0.120)	(0.376)	
	Papel B: Vears 2006 2020								
Trend at Weather Station	-0.056		1 ai	_0 130**	_0 014	2020		-0 107*	
field at weather Station	(0.054)			(0.058)	(0.014)			(0.063)	
Trend in CMIP5 - BCP4 5	(0.004)	0.175		0.000)	(0.000)	0.422		(0.003) 0.262	
field in Chill 5 fiel 4.5		(0.260)		(0.295)		(0.318)		(0.323)	
Trend in CMIP5 - BCP8 5		(0.200)	0 346**	(0.255) 0 457***		(0.010)	0 400**	(0.020) 0 415***	
field in chill 5 fiel 0.5			(0.167)	(0.166)			(0.172)	(0.153)	
			()	()				()	
	Panel C: Years 2011-2020								
Trend at Weather Station	-0.037			-0.046	-0.062			-0.062	
	(0.049)			(0.046)	(0.056)			(0.052)	
Trend in CMIP5 - RCP4.5		0.158		-0.034		0.159		-0.033	
		(0.192)		(0.186)		(0.197)		(0.188)	
Trend in CMIP5 - RCP8.5			0.716^{***}	0.730^{***}			0.733^{***}	0.737^{***}	
			(0.155)	(0.162)			(0.166)	(0.168)	
Observations	72	72	72	72	68	68	68	68	

Table 3: Comparing Spatial and Temporal Heterogeneity in Trends

This table examines spatial and temporal heterogeneity in various data sources. A separate linear time trend $\widehat{\beta_{am}}$ is fit for each month and airport: $D_{amy} = \alpha_{am} + \beta_{am}y + \epsilon_{amy}$. We then regress the trend in the futures data $\widehat{\beta_{am}^f}$ on the trend in the weather station data $\widehat{\beta_{am}^s}$ as well as the trends in the climate model forecasts $\widehat{\beta_{am}^{4.5}}, \widehat{\beta_{am}^{8.5}}$ under RCP 4.5 and RCP8.5, respectively. Columns (1a)-(1d) include all months (November - March for heating degree days and June - September for cooling degree days). Columns (2a)-(2d) exclude February for the four northeastern airports. Panels vary the years over which the time trends are estimated. Stars indicate significance levels: * 10%, ** 5%, *** 1%.

30