

MARKET EXPECTATIONS ABOUT CLIMATE CHANGE

Wolfram Schlenker♣ and Charles A. Taylor♠

February 2019

Abstract

An emerging literature examines how agents update their beliefs about climate change. Most studies have relied on indirect belief measures or opinion polls. We analyze a direct measure: prices of financial products whose payouts are tied to future weather outcomes. We compare these market expectations to climate model output for the years 2002 to 2018 as well as observed weather station data across eight cities in the US. All datasets show statistically significant and comparable warming trends. Nonparametric estimates suggest that trends in weather markets follow climate model predictions and are not based on shorter-term variation in observed weather station data. When money is at stake, agents are accurately anticipating warming trends in line with the scientific consensus of climate models.

Keywords: Market expectations, belief formation, weather markets, climate change

♣ Columbia University and NBER. Email: wolfram.schlenker@columbia.edu.

♣ Columbia University. Email: cat2180@columbia.edu.

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 climate change caused by human activities is occurring now.” On the other hand, Oreskes & Conway (2010) argue that a handful of scientists with strong ties to particular industries are “keeping the controversy alive” by spreading doubt. Certain politicians in the US have questioned the evidence on climate change, with some famously calling it an “elaborate hoax.” As a result, personal beliefs about climate change have been shown to vary across geography, political affiliation, educational status, and economic sector (Leiserowitz et al. 2016). Given the divergent beliefs about climate change, the debate surrounding the accuracy of climate change predictions and whether agents incorporate these predictions when deciding on their actions persists.

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, Jones & Olken 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 even infeasible for a one-time unknown weather shock. However, before agents can adapt, they first need to realize that the climate is changing. Again, there is uncertainty regarding how well climate models predict the future and whether agents base their decisions on model forecasts.

We address these questions by examining how market participants update their beliefs about future weather. 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 the 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 future climate.

First, we show that the futures market capitalizes weather shocks, i.e., deviations from historical averages, in the two weeks before an unexpected weather deviation occurs. This is consistent with an earlier finding that for longer horizons of 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 & Diebold 2005). Therefore, futures prices more than 10 days before the start of a month therefore reflect expectations about a month’s weather before a particular year’s outcomes are known.

Second, we find that market expectations as measured by futures prices when weather outcomes are unknown have been trending at the same rate as temperature forecasts in the CMIP5 archive, the latest repository where various climate modeling groups made predictions for 2006 onward, as well as observed temperatures 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. Predictions of climate models have materialized, at least on average, validating model forecasts, and financial speculators with money on the line have fully internalized these forecasts.

Third, while climate models correctly predicted average trends in degree days, the spatial heterogeneity among the eight cities does not match the observed distribution. This is likely due to the limited duration of our time series and the influence of outliers that would otherwise average out over space. It has also been argued that predicting average warming is easier as it relies on a simple balance of energy (radiative forcing), while predicting spatial heterogeneity requires predicting shifts of the atmospheric system like the jet stream (Hsiang & Kopp 2018). There is an active discussion whether a shift in the jet stream will reduce February temperatures as it allows for cold air from the arctic to influence weather on the East Coast. The futures market seems to agree with this theory, as there has been a significant increase in expected heating degree days in February.

Fourth, market expectations have been trending up smoothly in line with climate model predictions and do not seem to respond to year-to-year fluctuation in weather outcomes. In other words, market participants do not myopically update based on weather outcomes in the previous year, but proactively anticipate a warming climate.

Finally, we employ LASSO regressions to examine how oceanic oscillation indices affect temperatures across the eight cities in our sample. Removing these large-scale effects reduces the year-to-year variability, but does not change the observed trend. The observed warming trend is hence not driven by the major oceanic drivers of the natural variability, but rather caused by increased greenhouse gas emissions.

1 Background and Model

Recent research has shown how weather fluctuations affect the corporate sector. The profitability of the food industry responds to fluctuations in the Palmer drought index with direct implications for the sector’s valuation (Hong, Li & Xu 2019). More broadly, corporate earnings of several sectors of the US economy are sensitive to temperature fluctuations (Addoum, Ng & Ortiz-Bobea 2019). Weather markets offer companies a hedge against such fluctuations as well as a direct measure of the market expectation of future climate.

A second strand of literature has emphasized how climate change policy that is designed to limit emissions affects the profitability of various companies. Meng (2017) shows how the stock market incorporates changes in the likelihood of US carbon regulation as measured by betting markets. Furthermore, limiting emissions may render a company’s marginal reserves, i.e., the most costly ones, worthless as they can no longer be extracted (McGlade & Ekins 2015). Thus future climate expectations are key to the sector’s profitability.

A third strand of the literature focuses on how agents adjust their behavior in response to environmental forecasts (Rosenzweig & Udry 2014, Neidell 2009). Shrader (2017) 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 & Emerick (2016) find that yield changes in response to observable long-term temperature trends are not significantly different from yield changes in response to random weather shocks.

What is common to most of the previous studies is that researchers infer indirectly how agents update their beliefs on climate. Some authors have modeled how market participants learn about and adapt to changing weather conditions. For example, Kala (2017) 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, Obradovich & Lehner 2018). On the other hand, public opinion surveys ask respondents to self-report their beliefs. We know there is significant variation in public opinion about climate change across the U.S., which varies by location and demographic characteristics (Howe et al. 2015). Public opinion on climate change also seems to be driven by recent weather events, especially extremes. 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.

We add to this literature by using a different approach to directly measure beliefs about

climate change in financial markets. It is possible to observe the market expectations about climate by looking at the price of futures contracts that are linked to weather outcomes. The predominant contracts are based on heating and cooling degree days which are indexed to 65°F, the temperature considered the most comfortable for humans on average. It is also a common standard for utility companies because heating and cooling systems tend to be turned on that level, respectively.

Cooling degree days (CDD) measure by how much and for how long temperatures exceed 65°F and thus require cooling, hence the name cooling degree days. The exact formula to derive CDD (CDD_{ik}) for day k in location i with average temperatures equal to τ_{ik} is

$$CDD_{ik} = \max\{\tau_{ik} - 65, 0\} \quad (1)$$

These daily measures are then summed over all $k = 1 \dots K$ days of a week w , month m , or season of the year y to derive the total number of cooling degree days per time interval t

$$CDD_{it} = \sum_{k=1}^K CDD_{ik} = \sum_{k=1}^K \max\{\tau_{ik} - 65, 0\} \quad (2)$$

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 (HDD_{ik}) for day k in location i with the daily average temperature equal to τ_{ik} is

$$HDD_{ik} = \max\{65 - \tau_{ik}, 0\} \quad (3)$$

The daily measures are again summed over all K days of a week w , month m , or season of the year y to derive the total number of heating degree days per time interval t

$$HDD_{it} = \sum_{k=1}^K HDD_{ik} = \sum_{k=1}^K \max\{65 - \tau_{ik}, 0\} \quad (4)$$

In the first step, we estimate the timing of when the market updates its beliefs, i.e., capitalizes information into prices, by running a weekly regression of changes in futures prices on the weather of that week and the weeks that follow. Weekly aggregation addresses the discontinuous pricing given weekends, market closures, etc. We approximate the baseline expectations of the seasonality in degree days by regressing the daily observed degree days on

a flexible restricted cubic spline¹ using data from 1950-2018. This gives us the average weekly w degree days for each location i : \overline{DD}_{iw} . We then obtain weather shocks by subtracting this average baseline from the observed weather outcome DD_{it} . We regress weekly price changes $[p_{it} - p_{i[w-1]}]$ on these weather shocks. Since markets are forward looking and should only respond to *news*, we include up to $L = 3$ leads, and vary the set of fixed effects α_{ikw} as described in the regression table.

$$[p_{it} - p_{i[w-1]}] = \alpha_{ikw} + \sum_{l=0}^L \beta_l [DD_{i[w+l]} - \overline{DD}_{i[w+l]}] + \epsilon_{iw} \quad (5)$$

In our baseline specification we set leads beyond the end of a month equal to zero. A weather derivative on the number of cooling degree days in June should not respond to weather outcomes in July. For example, price changes in the last week of June should only respond to what happened in that week, while leads on future weeks fall outside the month. In a sensitivity check we limit the analysis to weeks that are at least $L = 2$ weeks from a month's end and find similar results.²

Confirming earlier findings that weather predictions beyond a narrow 10-14 day window are not better than climatic normals, we argue that futures prices at least 10 days before the start of a month represent ex-ante expectations. We estimate the annual time trends in degree days from three different data sources: the market expectation as measured by the price of weather contracts prior to when the random weather is known, predictions from the CMIP5 archive of climate models, and the observed weather station data on which the contracts are based. We aggregate the data by summer and winter season of each year y and regress various contracts that are based on eight city-level weather stations on a simple common time trend

$$DD_{iy} = \alpha_i + \beta y + \epsilon_{iy} \quad (6)$$

In sensitivity checks we allow the time trend to vary by city and by each month within the seasons.

We next relax the linearity assumption and present lowess regression for the average residual after removing location fixed effects, i.e., the mean for each location. Mathematically, we derive the residuals $DD_{it} - \overline{DD}_{it}$ for each year and location, then obtain the annual

¹Our specification uses 5 knots in the day of the year, subject to the constraint that the effect at the first day of the year and the last day of the year are the same.

²The more leads that we include, the more weeks we have to drop.

average residual r_{iy}

$$r_{iy} = \frac{1}{8} \sum_{i=1}^8 [DD_{iy} - \overline{DD}_{iy}] \quad (7)$$

before we run lowess regression on the annual observations for $y = 2002 \dots 2018$.

While we find consistent evidence in warming trends across various data sources, one concern may be that oceanic oscillations, which have been shown to be strong predictors of temperatures (Zebiak & Cane 1987), might have driven the trend in observed warming. To rule this out we link the monthly observed station level data DD_{im} in the 17 years 2002-2008 to the six oceanic indices for the same months as well as a linear time trend.

$$DD_{im} = \alpha_{im} + \beta_{im}y + \sum_{k=1}^K \gamma_{ikm}o_{km} + \epsilon_{im} \quad (8)$$

Given the small degrees of freedom, we rely on machine learning to pick the optimal model, specifically LASSO regression using the Extended Bayesian Information Criteria (EBIC). We then partial out the effect of the observed oceanic oscillation indices, i.e., the γ_{ikm} chosen under the LASSO regression.

2 Data

2.1 Futures 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 in 2018. 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 Figure 1. In the past there were more cities with weather markets, but trading in several cities was halted due to a lack of liquidity. Therefore, we focus on the eight US cities for which contracts are still available in 2018. We have data ranging from the 2001/2002 winter through the 2018 summer. When we present seasonal aggregates, winter months are assigned to the new year, i.e., November 2001 - March 2002 is recorded as winter 2002. Our data ranges from November 2001- September 2018, but is aggregated as winter 2002 -

summer 2018.

Contracts are priced based on cumulative monthly heating degree days (HDD) in winter months and cooling degree days (CDD) in summer months. Degree days are a measure of how much a day's average temperature, calculated as the mean of the day's high and low temperature, deviates from a baseline. The CME temperature product uses the same baseline of 65°F (18°C). Thus, heating degrees for a given day are the number of degrees that the average temperature of that day is below 65°F, and cooling degrees are the number of degrees that the average temperature is above 65°F. The accumulation period of each contract begins on the first calendar day of the contract month and ends on the last calendar day of the contract month. Trading volume of contracts generally increases in the last two months before the start of a month, and is more infrequent before then.

The final settlement price 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 is \$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 CDD for the month of July settled at 330 cooling 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).

The main participants in the weather market are firms seeking to offset 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 a 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 (i.e., private beliefs on climate).

Daily futures prices (end of day) were obtained from Bloomberg terminals. In the absence of market activity, prices are simply carried forward. 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 privately over the counter, and it is missing for most days. 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.

One concern that we will return to in our empirical analysis is the liquidity of the market. CME has a process in place to update prices on days in which no transacted volume occurs. Prices are updated based on the mid-point of outstanding but non-converging bids and offers

that are registered on the exchange. In the absence of live bids and offers, price changes can be derived from option prices linked to the relevant month’s contract and/or the seasonal strip contract (the aggregated contract for the entire winter or summer months). We can infer that there was market activity, i.e., either a market clearing trade or a new bid/offer, when prices change over time. Such price changes are a sufficient but not necessary condition for market activity, as trades might happen at the previous day’s price. Price fluctuations tend to increase the closer one gets to the contract live month. In our baseline we exclude city-year observations if city-month contract prices did not vary over the roughly 45-day period ranging from 14 days prior to the start of the contract month and the end of the contract month for *any* of the months of the season. We find similar results when they are included.

Some data cleaning was necessary because of “sticky fingers,” e.g., sudden price jumps by a factor of 10. For example, a price series was 91, 91, 910. We contacted Bloomberg about whether these were data entry errors, but they assured us that the data had been cleaned. The last trade might reflect an erroneous entry by a trader.³

2.2 Weather Station Data

Daily temperature data come from airport station monitors that are linked to each city’s futures contract. We obtained the station ID of the weather station underlying the contract, and downloaded the data on minimum and maximum temperature from the National Oceanic and Atmospheric Administration’s FTP server. A very small number of days have missing values, in which case we replace the missing value with the previous’ day value. 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 3 above.

2.3 Oceanic Oscillation Indices

Oceans exert strong influences on weather. One of the strongest and most famous is the El Niño - Southern Oscillation, warming in the eastern Pacific Ocean that has been linked to periodic climate shifts across the globe (Zebiak & Cane 1987). To rule out that recent trends in observed weather are driven by trends in oceanic oscillations, we estimate monthly models linking temperature data in a city to six monthly oceanic oscillation indices: ENSO (El Niño - Southern Oscillation), NAO (North Atlantic Oscillation), PNA (Pacific/ North American

³The exact adjustments are listed in the online appendix A1.

Teleconnection Pattern), AO (Arctic Oscillation), and AAO (Antarctic Oscillation). These oscillations were downloaded from a FTP server of NOAA.⁴

2.4 Climate Model Data

The Coupled Model Comparison Project (CMIP) 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 trends in climate from 2006 onwards. These runs were done only using observed climate data prior to 2006 in the calibration. We obtain daily values from NASA NEX-GDDP, a dataset of 21 models that were spatially downscaled to a common grid of 0.5° latitude and longitude. We pick the grid cell in which the weather station is located.

NASA NEX-GDDP is bias corrected. There is a well-documented literature describing how climate models can get the baseline averages wrong for various grids because the models are primarily designed to simulate how *shocks*, i.e., deviations from the average, promulgate through the system. NASA NEX-GDDP therefore adjusts for possible biases by ensuring the baseline average match those observed for the grid cell. Regardless, subtracting the same constant every year will not affect predicted changes in temperature, the relevant factor in this and other trend analyses.

NASA NEX-GDDP has data for two scenarios: Representative Concentration Pathway (RCP) 4.5, assuming 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°C (3.2°F) by 2100. There is large spatial heterogeneity, and warming in the US is usually projected to be higher than the global average by a factor or two, i.e., 6.4°F . 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. This results in a global mean temperature increase of 3.7°C by 2100. However, in the short term (2005-2018) that includes our study period, both models give very similar projections. They are only predicted to diverge towards the end of the century as carbon emissions accumulate over time.

Figure 2 shows box plots for the number of cooling degree days by month for the eight cities with weather futures contracts. The red line displays the weather station data and

⁴Monthly ENSO values were downloaded from <ftp://ftp.cpc.ncep.noaa.gov/cwlinks/>, while daily values for the remaining indices was obtained from <ftp://ftp.cpc.ncep.noaa.gov/cwlinks/> and averaged over all days of a month.

the blue line shows the climate model data. Both use data from 1950-2005, which was the historical baseline period in the CMIP5 archive.⁵ 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 it might differ from the observed temperature at any given point (weather station) if there is spatial heterogeneity. For example, a city close to a mountain might have a lower temperature than the temperature of the surrounding area when averaged over the entire grid.

Figure 3 shows the analogous plot for heating degree days. Both plots show strong seasonality: more cooling degree days in the summer, and more heating degree days in the winter. As expected, more northerly cities (Chicago, Minneapolis, New York) have relatively more HDD and less CDD, while more 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.

While Figures 2 and 3 compare the historic baseline period in the climate models, Figure 4 compares station level data on cumulative degree days for a given city-month with the settled prices of the city-month's futures contract (averaged over the seven days after the month's close). These data have been cleaned using the process described earlier. As expected, there is close alignment between futures prices and observed weather station data at the month's end, at which point all uncertainty has been resolved, i.e., the weather has been realized. The correlation of the two series is above 0.999. The scatter plot reveals a small number of deviations from the 45-degree line, which may be explained if the futures did not trade at the end of the month and hence the price may not reflect the final tally of observed degree days.

3 Empirical Results

3.1 Market Capitalization of Weather Shocks

We start by analyzing the timing of when weather shocks are capitalized into futures prices, which is akin to the market updating a particular year's weather shock. Forecasting and prediction skill of weather (short-term) and climate (medium to long-term) are closely connected

⁵Some of the 21 climate models do not account for leap years, i.e., February 29. For consistency, we rescale February in both station and climate model data to 28 days, i.e., if a weather station or climate model reports 29 days for February, we multiple the cumulative number of degree days by $\frac{28}{29}$.

(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.

Weather forecasts are widespread and freely available sources of information, and under efficient markets, we would expect that prices of weather futures to adjust based on these forward expectations. Forecast skill is a function of forecast range. There has been a sustained improvement in weather forecasting across all prediction ranges over the past decades. Forecasting skill seems to have plateaued in the early 2000s when the weather futures market was launched. At present, 5-day forecasts are very accurate with 90% skill, 7-day forecasts are good with 75% skill, and 10-day forecasts are poor with less than 50% skill (Bauer, Thorpe & Brunet 2015). Given this, we would expect an inverted U-shape in terms of the impact of weather shocks on current 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. As such, anticipated changes in weather around one week out should have the largest impact on current prices in an efficient weather market.

To test this, we regress changes in weekly futures prices (Friday to Friday) on actual weekly deviations in degree days from historical averages over the contemporaneous week (Saturday to Friday), as well as the three leading weeks as given in equation 5. For CDD, we include the four summer contract months (June to September), the months that showed the highest average number of cooling degree days in Figure 2. The results are shown in Panel A of Table 1. We find that the majority of price updating is driven by weather shocks that occur one week out, when 41-47% of a weather shock (deviation from seasonal average) is priced into futures. The concurrent week follows with an estimate of 17-22%, while shocks two weeks ahead only get reflected 6-8% in the futures price. Leads three weeks into the future have no significant coefficients. The sum of the coefficients for these three weeks $l = 0 \dots 2$ implies that 78% of a weather shock gets capitalized into futures prices during those three weeks, as shown in the most flexible model in column (7). Results are insensitive to the fixed effects included in the regression, which range from no fixed effects in column (1) to city by year by contract month fixed effects as well as week till settlement fixed effects, e.g., a fixed effect for the June 2017 contract in Atlanta as well as fixed effects for the week's position relative to the contract's maturity in column (7).

Panel B of Table 1 presents the regression results for heating degree days during the winter season from November to March. The coefficient for the period one week into the future again has the largest magnitude, being reflected 35% in the current week’s price change. HDD contract prices are also driven by weather shocks two weeks into the future, which are capitalized 23-25%. Together they account for 60% of price changes, as shown in column (7). Neither the contemporaneous week nor the one three weeks out are significant. The insignificant coefficient on the contemporaneous week suggests that short-term weather forecasting skill is better in winter months than summer months in the US.

Earlier we discussed how weather futures are sometimes not traded and hence might not reflect the latest information and the market expectations. Our baseline regression excluded week-on-week price changes if the prices remained constant throughout the entire week. Table A1 includes all the available data, even when no price change occurred during any of the trading days in a week (not just Friday to Friday). This increases the total number of observations by roughly 10%. The resulting coefficients are slightly smaller in magnitude, but not significantly different. The inclusion of these observations may induce attenuation bias as the prices do not reflect true market expectations (the market price has not adjusted). Excluding these observations has its own drawbacks, as the occurrence of a trade is endogenous and might signal extreme weather events that might be easier (or harder) to predict and hence have different forecast skill. However, the results are consistent either way.

In another sensitivity check in Table A2 we exclude weeks where the leads fall beyond the end of the contract month. In the baseline model, these weather shocks were set to zero as they fall outside the time range for which the weather contract is based. We skip the third lead as it was not significant, which allows for an extra week of data. The results are again consistent.

3.2 Linear Trends in Expectations

We now turn to our main analysis of market expectations and climate change. With weather futures contracts, we must be careful to separate price changes driven by short-term weather forecasts and those that reflect longer-term market beliefs on climate change. 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 like to use futures prices well before the contract’s delivery month to ensure that we are capturing market expectations of climate and not short-term weather forecasts. However,

precisely because weather is hard to forecast far in advance, trading does not pick up until one is getting closer to the delivery month. 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 between 30 to 10 days prior to the start of each contract month, e.g., the average price between June 1 and June 20 for a July CDD contract, which ensures that prices reflect future expectations and not contemporaneous weather as confirmed in the previous section. For the CDD contracts, this average price for each contract-month is again summed over the summer months from June to September, and for the HDD contracts, it is again summed over the winter months from November to March. To capture overall trends, Table 2 regresses the total degree days for each city and season, e.g., annual summer cooling degree days from June to September, on an annual time trend as shown in equation 6.

Column (1a)-(1d) all use the same set of observations where futures data are non-missing.⁶ Column (1a) uses the baseline average of futures prices traded 30 to 10 days prior to the start of the contract month. Panel A shows cooling degree days in summer and Panel B shows heating degree days in winter. Both show statistically significant warming in the US with anticipated cooling degree days increasing by about 10 per year during the summer and anticipated heating degree days declining by about 8 per year during the winter. Column (1b) uses weather station data. The observed trends are larger in magnitude with an increase of 12 cooling degree days per year during the summer and a decrease of 17 heating degree days during the winter. The standard errors are much larger given the greater year-to-year swings due to random weather fluctuations, e.g., cold spells during the so-called Polar Vortex. Given these larger standard errors, they are not significantly different from the trends anticipated by the futures market as shown in (1a). The smaller standard errors in the futures data relative to the station level data suggest that we are correctly measuring market expectations over the longer term and not annual weather realizations, which are much noisier. Finally, columns (1c)-(1d) give the predicted trends that climate models had forecast for the same time period in the CMIP5 archive under the RCP4.5 and RCP8.5 model, respectively. Recall that these forecasts were made using climate data from 2005 and before, i.e., the modeling groups had no chance to match their forecast to observed trends. These annual warming trends closely mirror the trend in the futures data, i.e., the markets seem to have fully internalized the projections of climate scientists.

⁶For some cities, the futures contracts were established in later years and hence have a shorter time span. In addition, we excluded contracts that had a constant price over the entire contract month and leading up to it as further outlined in Section A1.

While columns (1a)-(1d) purposefully keeps the set of city-year observations constant, columns (2a)-(2d) replicate the analysis with different subsets of the data. Column (2a) excludes seasons where any of the monthly contracts did not experience a price change between 30 to 10 days before the start of a delivery month, an indication that no trading activity occurred. This exclusion has a very limited effect on the estimated coefficients. Next we address concerns about the endogeneity of this market, e.g., contracts are especially traded in 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 futures data existed) and again find very similar trends.

3.3 Nonlinear Trends in Expectations

Figure 5 relaxes the linearity assumption of the time trend and instead plots lowess regressions of the average annual residuals, i.e., after city fixed effects have been removed to account for different average climates (Atlanta is hotter than Minneapolis) and residuals have been averaged over the eight cities. The exact formula is given in equation (7) above. The lines in green, red, cyan, and blue are the same as in columns (1a)-(1d) of Table 2, respectively. All of them show a steady upward trend for cooling degree days and a downward trend for heating degree days. The year 2017/2018 was especially warm, leading to a sharp drop for that winter.⁷ The observation for summer 2018 is missing because data for June 2018 contracts were not available. Figure A2 replicates the analysis for the summer months from July to September and finds comparable results. Mirroring the heating degree day trend, there is an accelerated increase in warming in 2018.

Next, we look at variability in the non-linear warming trends. The red line showing actual weather outcomes from station data is the most variable as it is influenced by annual and city-level weather anomalies. The magenta line, an additional trend from those included in Table 2, partials out the effects of the ocean oscillation indices that we obtained from the LASSO regression in equation (8) that are summarized in Table 3. Several monthly degree day totals are influenced by the phase of various oceanic oscillation indices. The trend is less variable than the one based on the raw weather station data, i.e., the standard deviation around the trend is lower for the magenta line than the red line, both for cooling and heating degree days. But the overall trend is rather unchanged. In other words, the warming trend observed from station data is not due to the natural variability of oceanic-atmospheric phenomena.

⁷Recall that the 2017-2018 winter is coded as 2018.

Lastly, the smooth green line, which is the trend in futures market expectations, closely follows the cyan and blue line of climate projections, and is not influenced by the red or magenta line, suggesting that the beliefs are not myopically updated based on recently experienced weather but rather tied to the smooth warming trend projected by climate models. Financial markets seem to fully incorporate this scientific consensus.

3.4 Sub-seasonal Trends in Expectations

Tables 4 and 5 separate the aggregate seasonal analysis into months for cooling and heating degree days, respectively. Not surprisingly, there is more variability in these coefficient estimates when the data is summed over a smaller time frame given that anomalous weather shocks at the individual city-level are not being averaged over the course of a season. One unique feature is that futures prices for the month of February show a significant positive trend, i.e., an increase in HDD, which would be consistent with cooling and not warming. All other winter months either show a significant negative trend or insignificant trend with a negative point estimate. This may be explained by recent literature suggesting that melting ice sheets destabilize the jet stream, leading to an increased frequency of stable weather patterns that bring cold arctic air to Europe and North America (Francis & Vavrus 2015). Another paper concludes that “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.” (Zhang et al. 2016). Kim et al. (2014) note that “the mechanism that links sea-ice loss to cold winters remains a subject of debate.” So while there remains an active scientific debate, the futures market seems to anticipate cooling in February consistent with the recent literature. This is striking because observed station level data has exhibited a February warming trend across cities, as predicted by the climate models in the CMIP5 archive that preceded this recent debate of the shift in the polar vortex.

Since the polar vortex is expected to mainly influence the East Coast, we provide another nonlinear regression based on the February futures prices for the six cities in the eastern half of the US (excluding Las Vegas and Sacramento) in Figure 6. The figure shows a clear uptick in the expected number of heating degree days around 2006, which started to plateau around 2013. The futures market seems to have shifted to an expectation of colder Februaries.

Such a shift in the polar vortex is likely to aggravate the estimated economic impacts of climate change. A myriad of studies find that moderate temperatures are optimal for

energy use (Auffhammer & Mansur 2014), labor productivity (Graff Zivin & Neidell 2014), mortality (Barreca et al. 2016), migration (Feng, Krueger & Oppenheimer 2010, Missirian & Schlenker 2017), GDP growth (Burke, Hsiang & Miguel 2015), and agriculture (Schlenker & Roberts 2009, Auffhammer & Schlenker 2014). Warming usually results in damages from higher summer temperatures while yielding benefits from milder winter temperatures. A shift in the polar vortex with colder February temperatures would hence accentuate damages, as a result of the hotter summers and colder winters. Overall, winters on average are still getting warmer and hence yielding benefits, but a shift in the polar vortex with colder February temperatures will offset some of those benefits.

3.5 Spatial Heterogeneity in Warming

We examine spatial heterogeneity in the warming trends in Figure 7. While climate models correctly predicted average annual trends in degree days, the predicted spatial heterogeneity among the eight cities in our sample does not match the observed distribution from station data. A regression of trends in city-level warming expectations derived from futures data on city-level warming trends from the climate models does not provide a significant estimate for spatial patterns in cooling degree days in summer. For heating degree days, the coefficient is negative under the RCP4.5 scenario, suggesting that cities that were predicted to see higher-than-average winter warming in the climate models actually had lower-than-average warming in the futures data. There are two possible explanations.

First, as discussed in the previous section, 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 & 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.

Second, the futures market only came into existence in the fall of 2001 which might be too short a time frame to pick up city-level warming trends that are heavily influenced by city-level outliers that would otherwise average out over space.

3.6 Trends Prior to 2002 and Post 2018

Given the drawbacks of a limited time series of 17 years, we contrast the trend from 2002 to 2018 in our analysis to nonparametric trends in weather station data from 1950 to 2018 in Figure 8, which are available before the futures data started in 2002. Residuals, i.e., deviations in city-level seasonal totals compared to the historic average, are again color-coded by each city’s airport and the nonparametric lowess regression is added as a black line. As other authors have emphasized (Burke & Emerick 2016, e.g.), observable warming trends become apparent in the data around 1980. We find the same for the eight cities in our sample: from 1950 to 1980 the nonparametric black line is rather flat. Starting around 1980 there is a clear uptick in warming over the last four decades as manifested by a higher number of cooling degree days and a lower number of heating degree days.

This warming trend is predicted to intensify in the future as shown in Figure 9, which displays the output from the climate models to the end of the century. 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 shows it for RCP8.5. Our sample period (2002-2018) is indicated by dashed grey lines. While the models accurately predicted the initial trend from 2002 to 2018, which are comparable in both RCP scenarios, there is greater uncertainty about what will happen towards the end of the century as greenhouse gases accumulate in the atmosphere. Note the policy-driven divergence in warming trends projected under the two different climate model scenarios, i.e., much stronger warming under RCP8.5 (right column) than under RCP4.5 (left column). There is also spatial heterogeneity in predicted warming, albeit subject to the caveat discussed in the previous section about the uncertainty around the impact of shifting circulation patterns.

4 Conclusion

This paper contributes to the literature on belief formation and expectations about climate change. To the best of our knowledge, we are the first to utilize a direct measure of climate change expectations as derived from weather-based futures contracts. The evidence shows that financial markets fully incorporate climate model projections. We find that the market has been accurately pricing in climate change, largely in line with global climate models, and that this began occurring at least since the early 2000s when the weather futures markets were formed.

This has relevance for the corporate sector. Recent studies have highlighted how the val-

uations of companies and entire industries are sensitive to weather fluctuations. 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.

There are also policy implications of our findings, especially since some politicians still question the existence and extent of climate change. Anyone doubting the observed warming trend can make a significant profit by betting against it in weather markets. However, 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. When money is on the line, it is hard to find parties willing to bet against the scientific consensus.


References

- Addoum, Jawad M., David T. Ng, and Ariel Ortiz-Bobea.** 2019. “Temperature Shocks and Earning News.” *Review of Financial Studies*.
- American Association for the Advancement of Science.** December 9, 2006. “AAAS Board Statement on Climate Change.”
- Auffhammer, Maximilian.** 2018. “Quantifying Economic Damages from Climate Change.” *Journal of Economic Perspectives*, 32(4): 33–52.
- Auffhammer, Maximilian, and Erin T Mansur.** 2014. “Measuring climatic impacts on energy consumption: A review of the empirical literature.” *Energy Economics*, 46: 522–530.
- Auffhammer, Maximilian, and Wolfram Schlenker.** 2014. “Empirical studies on agricultural impacts and adaptation.” *Energy Economics*, 46: 555–561.
- Auffhammer, Maximilian, Solomon M. Hsiang, Wolfram Schlenker, and Adam Sobel.** 2013. “Using Weather Data and Climate Model Output in Economic Analyses of Climate Change.” *Review of Environmental Economics and Policy*, 7(2): 181–198.
- Barreca, Alan, Karen Clay, Olivier Deschênes, Michael Greenstone, and Joseph S. Shapiro.** 2016. “Adapting to Climate Change: The Remarkable Decline in the U.S. Temperature-Mortality Relationship over the 20th Century.” *Journal of Political Economy*, 124: 105–159.
- Bauer, Peter, Alan Thorpe, and Gilbert Brunet.** 2015. “The quiet revolution of numerical weather prediction.” *Nature*, 525: 47–55.
- Burke, Marshall, and Kyle Emerick.** 2016. “Adaptation to Climate Change: Evidence from US Agriculture.” *American Economic Journal: Economic Policy*, 8(3): 106–140.
- Burke, Marshall, Solomon M. Hsiang, and Edward Miguel.** 2015. “Global Non-linear Effect of Temperature on Economic Production.” *Nature*.
- Campbell, Sean D., and Francis X. Diebold.** 2005. “Weather Forecasting for Weather Derivatives.” *Journal of the American Statistical Association*, 100(469): 6–16.

- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken.** 2014. “What Do We Learn from the Weather? The New Climate-Economy Literature.” *Journal of Economic Literature*, 53(3): 740–798.
- Feng, Shuaizhang, Alan B. Krueger, and Michael Oppenheimer.** 2010. “Linkages among climate change, crop yields and MexicoUS cross-border migration.” *Proceedings of the National Academy of Sciences of the United States*, 107(32): 14257–14262.
- Francis, Jennifer A., and Stephen J Vavrus.** 2015. “Evidence for a wavier jet stream in response to rapid Arctic warming.” *Environmental Research Letters*, 10: 014005.
- Graff Zivin, Joshua, and Matthew Neidell.** 2014. “Temperature and the Allocation of Time: Implications for Climate Change.” *Journal of Labor Economics*, 32(1): 1–26.
- Hong, Harrison, Frank Weikai Li, and Jiangmin Xu.** 2019. “Climate risks and market efficiency.” *Journal of Econometrics*, 208(1): 265–281.
- Howe, Peter D., Matto Mildener, Jennifer R. Marlon, and Anthony Leiserowitz.** 2015. “Geographic variation in opinions on climate change at state and local scales in the USA.” *Nature Climate Change*, 5: 596603.
- Hsiang, Solomon, and Robert E. Kopp.** 2018. “An Economist’s Guide to Climate Change Science.” *Journal of Economic Perspectives*, 32(4): 3–32.
- Kala, Namrata.** 2017. “Learning, Adaptation, and Climate Uncertainty.” *Working Paper*.
- Kaufmann, Robert K., Michael L. Mann, Sucharita Gopal, Jackie A. Lieder-
man, Peter D. Howe, Felix Pretis, Xiaojing Tang, and Michelle Gilmore.** 2017. “Spatial heterogeneity of climate change as an experiential basis for skepticism.” *Proceedings of the National Academy of Sciences*, 114(1): 67–71.
- Kim, Baek-Min, Seok-Woo Son, Seung-Ki Min, Jee-Hoon Jeong, Seong-Joong
Kim, Xiangdong Zhang, Taehyoun Shim, and Jin-Ho Yoon.** 2014. “Weakening of the stratospheric polar vortex by Arctic sea-ice loss.” *Nature Communications*, 5: 4646 EP.
- Leiserowitz, Anthony, Edward Maibach, Connie Roser-Renouf, Seth Rosenthal,
and Matthew Cutler.** 2016. *Climate Change in the American Mind: November 2016*. Yale University and George Mason University.


- McGlade, Christophe, and Paul Ekins.** 2015. “The geographical distribution of fossil fuels unused when limiting global warming to 2C.” *Nature*, 517: 187–190.
- Meng, Kyle C.** 2017. “Using a Free Permit Rule to Forecast the Marginal Abatement Cost of Proposed Climate Policy.” *American Economic Review*, 107(3): 748–784.
- Missirian, Anouch, and Wolfram Schlenker.** 2017. “Asylum Applications Respond to Temperature Fluctuations.” *Science*, 358: 1610–1614.
- Moore, Frances, Nicholas Obradovich, and Flavio Lehner.** 2018. “Rapidly-Adjusting Perceptions of Temperature in a Changing Climate.” *Working Paper*.
- Neidell, Matthew.** 2009. “Information, Avoidance Behavior, and Health: The Effect of Ozone on Asthma Hospitalizations.” *Journal of Human Resources*, 44(2): 450–478.
- Oreskes, Naomi, and Erik M. Conway.** 2010. *Merchants of Doubt*. Bloomsbury Press.
- Rosenzweig, Mark R., and Christopher Udry.** 2014. “Rainfall Forecasts, Weather, and Wages over the Agricultural Production Cycle.” *American Economic Review*, 104(5): 278–283.
- Schlenker, Wolfram, and Michael J. Roberts.** 2009. “Nonlinear Temperature Effects Indicate Severe Damages to U.S. Crop Yields under Climate Change.” *Proceedings of the National Academy of Sciences*, 106(37): 15594–15598.
- Shrader, Jeffrey.** 2017. “Expectations and adaptation to environmental risks.” *Working Paper*.
- Turkey, J. W.** 1977. *Exploratory Data Analysis*. AddisonWesley.
- Zebiak, Stephen E., and Mark A. Cane.** 1987. “A Model El Niño-Southern Oscillation.” *Monthly Weather Review*, 115: 2262–2278.
- Zhang, Jiankai, Wenshou Tian, Martyn P. Chipperfield, Fei Xie, and Jinlong Huang.** 2016. “Persistent shift of the Arctic polar vortex towards the Eurasian continent in recent decades.” *Nature Climate Change*, 6: 1094–1099.

Figure 1: Location of Eight Airports in Sample




Notes: Figure displays the location of the eight airports for which weather derivatives were traded in 2018. They are from north to south: Minneapolis - Saint Paul (MSP), Chicago O'Hare (ORD), New York LaGuardia (LGA), Cincinnati - Northern Kentucky (CVG), Sacramento (SAC), Las Vegas (LAS), Atlanta (ATL), and Dallas Fort Worth (DFW).

Figure 2: Cooling Degree Days by Month and Weather Dataset for Eight Airports




Notes: Figure displays boxplots of cooling degree days, i.e., the amount average daily temperatures exceed 65F, summed over all days of a month. The box plots show the variation for the calendar month across years in 1950-2005. Blue bar graphs use station level data from the airport, while the red bar graph uses the daily data from NASA NEX-GDDP of the grid cell in which the airport is located. Boxes indicate the 25-75% range, with the median shown by a horizontal bar. Whiskers extend to the lower and upper adjacent value using STATA's default parameter from Turkey (1977).

Figure 3: Heating Degree Days by Month and Weather Dataset for Eight Airports




; *Notes:* Figure displays boxplots of cooling degree days, i.e., the amount average daily temperatures exceed 65F, summed over all days of a month. The box plots show the variation for the calendar month across years in 1950-2005. Blue bar graphs use station level data from the airport, while the red bar graph uses the daily data from NASA NEX-GDDP of the grid cell in which the airport is located. Boxes indicate the 25-75% range, with the median shown by a horizontal bar. Whiskers extend to the lower and upper adjacent value using STATA's default parameter from Turkey (1977).

Figure 4: Observed Cooling Degree Days versus Futures Prices



Notes: Left graph displays a scatter plot of the 513 observed monthly cooling degree day (CDD) totals against the average futures price the week following the end of the summer months June-September. Right graph displays a scatter plot of the 630 observed monthly heating degree day (HDD) totals against the average futures price the week following the end of the winter months November-March. Both graphs use data from winter 2001/2002 through summer 2018 across the eight airports. Graphs exclude 8 CDD (12 HDD) futures contracts where prices did not move in the 14-day window spanning between 7 days before the end of the month or seven days after the end of the month, as it is not clear that they reflect final totals.

Figure 5: Nonparametric Futures Prices and Weather




Notes: Figure estimates nonparametric trends using lowess regression on the average annual residual among the eight airports, i.e., the annual residuals after taking out airport fixed effects (means) are averaged by year before running lowess on the annual observations. Red line shows the results for the observed station data, the magenta line partials out ocean oscillation indices that affect weather. The green line uses futures prices before the weather is realized, and the blue and cyan lines use model output from NASA NEX-GDDP. The left column shows the results for summer (June-September) CDD and the right column for winter (November-March) HDD.

Figure 6: Nonparametric Futures Prices for February Futures in Eastern US




Notes: Figure estimates nonparametric trends using lowess regression on the average annual residual among the six eastern airports, i.e., the annual residuals after taking out airport fixed effects (means) are averaged by year before running lowess on the annual observations. Residuals are color coded by airport and the lowess regression line is added in black.

Figure 7: Spatial Heterogeneity of Trends in Station, Futures, and Climate Model Data




Notes: Graphs show scatter plots of trends in station, futures, and climate model data by airport as shown in Figure 1. Airports are sorted from north to south. Left column shows the results for summer (June-September) CDD and the right column for winter (November-March) HDD.

Figure 8: Nonparametric Trends in Station Weather



Notes: Figure plots residuals from a regression on airport fixed effects are color-coded by airport. The lowest nonparametric regression line is added in black. Top row shows the results for summer (June-September) CDD and the bottom row for winter (November-March) HDD.

Figure 9: Nonparametric Predicted Trends in Climate Data Weather



Notes: Graphs show lowest nonparametric trend by airport for the 21 climate models in the NASA NEX-GDDP data base. Top row shows the results for summer (June-September) CDD and the bottom row for winter (November-March) HDD. Left column uses the predictions under the RCP4.5 scenario, while the right column uses RCP8.5.

Table 1: Market Updates: Futures Returns on Weather Shocks

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Cooling Degree Days June - September							
Shock in week t	0.170*** (0.025)	0.185*** (0.025)	0.170*** (0.024)	0.183*** (0.025)	0.194*** (0.026)	0.212*** (0.031)	0.221*** (0.032)
Shock in week t+1	0.413*** (0.024)	0.424*** (0.024)	0.416*** (0.024)	0.427*** (0.024)	0.433*** (0.024)	0.446*** (0.027)	0.470*** (0.028)
Shock in week t+2	0.058** (0.029)	0.067** (0.028)	0.061** (0.028)	0.065** (0.028)	0.076** (0.029)	0.072** (0.032)	0.087*** (0.032)
Shock in week t+3	-0.039 (0.028)	-0.028 (0.029)	-0.041 (0.028)	-0.037 (0.029)	-0.027 (0.030)	-0.047 (0.034)	-0.044 (0.037)
Observations	3344	3344	3344	3344	3344	3344	3344
Clusters	67	67	67	67	67	67	67
Fixed Effects		4	17	67	133	513	520
Panel B: Heating Degree Days November - March							
Shock in week t	0.050 (0.045)	0.057 (0.045)	0.048 (0.045)	0.051 (0.047)	0.045 (0.048)	0.049 (0.060)	0.045 (0.061)
Shock in week t+1	0.347*** (0.038)	0.354*** (0.038)	0.347*** (0.039)	0.348*** (0.039)	0.349*** (0.040)	0.348*** (0.044)	0.352*** (0.043)
Shock in week t+2	0.244*** (0.048)	0.252*** (0.048)	0.243*** (0.048)	0.242*** (0.048)	0.242*** (0.048)	0.231*** (0.051)	0.230*** (0.050)
Shock in week t+3	0.058 (0.042)	0.068 (0.042)	0.057 (0.042)	0.053 (0.043)	0.058 (0.044)	0.042 (0.050)	0.047 (0.052)
Observations	4199	4199	4199	4199	4199	4199	4199
Clusters	85	85	85	85	85	85	85
Fixed Effects		5	17	85	134	642	649
Fixed Effects	-	M	Y	YM	AY	AYM	AYM,W

Notes: Table regresses weekly cooling (CDD) and heating (HDD) degree day future returns (Friday-Friday) on deviations in degree days from seasonal averages. Both the contemporaneous week as well as three leads are included. Errors are clustered at the delivery month (e.g., November 2010). Columns control for different fixed effects across the 17 years in our sample (2002-2018), months (June-September for CDD and November-March for HDD) and 8 airports for which contracts are available. Fixed effects are: monthly (M), yearly (Y), year-month (YM), airport-year (AY), airport-year-month (AYM), as well as AYM plus week relative to contract maturity fixed effects (AYM,W). Tables excludes observations that did not see a price change during the entire week.

Table 2: Trends in CDD and HDD

	(1a)	(1b)	(1c)	(1d)	(2a)	(2b)	(2c)	(2d)
Panel A: CDD June-September								
Time trend (years)	10.282*** (0.804)	12.085*** (3.357)	9.134*** (0.593)	9.957*** (0.569)	8.995*** (1.206)	10.626*** (2.853)	9.374*** (0.513)	10.668*** (0.471)
Observations	117	117	117	117	90	136	136	136
Panel B: HDD November-March								
Time trend (years)	-8.472*** (1.995)	-16.743** (8.398)	-8.133*** (1.340)	-7.875*** (1.379)	-9.429*** (2.997)	-9.827 (7.189)	-8.971*** (1.173)	-9.223*** (1.173)
Observations	115	115	115	115	87	136	136	136
Airport FE	Yes							
Data	Futures	Station	RCP4.5	RCP 8.5	Futures	Station	RCP4.5	RCP 8.5
Years	Common	Common	Common	Common	Traded	All	All	All

Notes: Panel A regresses cooling degree days (CDD) for the summer months June-September on a linear time trend, while panel B looks at heating degree days for November-March in the years 2002-2018. Column (a) uses the average futures price 30 to 10 days before the start of each contract month, e.g., the average price between June 1 and June 20 for a July contract. Column (b) uses observed station level data for the month, while columns (c) and (d) use climate change forecasts in the NASA NEX-GDDP database under the RCP4.5 and RCP8.5 scenarios. Columns (1a)-(1b) estimate the trends for a consistent set of observations where futures data are available. Columns (2a)-(2b) conduct sensitivity checks to the included years. Columns (2a) exclude years where the price did not change over that time frame as the contract might not have been traded and hence not reflect changes in market expectations. Columns (2b)-(2d) include all years even if futures data is not available.

Table 3: LASSO Regression of Monthly Weather on Oceanic Oscillation

		Panel A: Cooling Degree Days																			
		ENSO				NAO				PNA				AO				AAO			
Month		6	7	8	9	6	7	8	9	6	7	8	9	6	7	8	9	6	7	8	9
ATL				X	X			X				X	X			X	X				X
CVG				X				X				X				X					X
DFW		X								X				X							
LAS																					
LGA														X							
MSP				X	X				X												
ORD									X												
SAC							X														

		Panel B: Heating Degree Days																								
		ENSO					NAO					PNA					AO					AAO				
Month		11	12	1	2	3	11	12	1	2	3	11	12	1	2	3	11	12	1	2	3	11	12	1	2	3
ATL							X	X																		
CVG		X					X	X																		
DFW							X																			
LAS					X					X					X					X					X	
LGA								X																		
MSP		X					X					X		X			X					X				
ORD			X					X																		
SAC					X					X																

Notes: Table summarizes LASSO regression (STATA package lassopack) of monthly cooling and heating degree days on six Oceanic Oscillation indices: ENSO (El Niño - Southern Oscillation), NAO (North Atlantic Oscillation), PNA (Pacific/ North American Teleconnection Pattern), AO (Arctic Oscillation), AAO (Antarctic Oscillation). regressions are done separately for each month while linking it to the values of the six oscillations for the same month as well as a linear time trend. Variables that are selected based on the Extended Bayesian Information Criteria (EBIC) are marked with an X.

Table 4: Trends in CDD by Month

	(1a)	(1b)	(1c)	(1d)	(2a)	(2b)	(2c)	(2d)
Panel A: CDD June								
Time trend (years)	3.377*** (0.361)	5.022*** (1.089)	1.587*** (0.225)	1.989*** (0.257)	3.612*** (0.434)	4.425*** (0.965)	1.574*** (0.193)	1.834*** (0.228)
Observations	119	119	119	119	105	136	136	136
Panel B: CDD July								
Time trend (years)	2.351*** (0.311)	1.920 (1.333)	2.635*** (0.208)	3.086*** (0.219)	2.194*** (0.375)	1.681 (1.260)	2.765*** (0.198)	3.199*** (0.207)
Observations	131	131	131	131	110	136	136	136
Panel C: CDD August								
Time trend (years)	2.550*** (0.292)	1.413 (1.239)	2.532*** (0.221)	3.118*** (0.217)	2.227*** (0.340)	1.192 (1.171)	2.667*** (0.211)	3.147*** (0.205)
Observations	131	131	131	131	114	136	136	136
Panel D: CDD September								
Time trend (years)	2.602*** (0.291)	3.912*** (1.017)	2.353*** (0.225)	2.491*** (0.218)	2.579*** (0.339)	3.327*** (0.981)	2.368*** (0.214)	2.487*** (0.208)
Observations	132	132	132	132	114	136	136	136
Airport FE	Yes							
Data	Futures	Station	RCP4.5	RCP 8.5	Futures	Station	RCP4.5	RCP 8.5
Years	Common	Common	Common	Common	Traded	All	All	All

Notes: Table replicates Panel A of Table 2 but estimates a separate trend for each month. It regresses cooling degree days (CDD) for the summer months June-September on a linear time trend in the years 2002-2018. Column (a) uses the average futures price 30 to 10 days before the start of each contract month, e.g., the average price between June 1 and June 20 for a July contract. Column (b) uses observed station level data for the month, while columns (c) and (d) use climate change forecasts in the NASA NEX-GDDP database under the RCP4.5 and RCP8.5 scenarios. Columns (1a)-(1b) estimate the trends for a consistent set of observations where futures data are available. Columns (2a)-(2b) conduct sensitivity checks to the included years. Columns (2a) exclude years where the price did not change over that time frame as the contract might not have been traded and hence not reflect changes in market expectations. Columns (2b)-(2d) include all years even if futures data is not available.

Table 5: Trends in HDD by Month

	(1a)	(1b)	(1c)	(1d)	(2a)	(2b)	(2c)	(2d)
Panel A: HDD November								
Time trend (years)	-4.229*** (0.535)	-0.287 (1.746)	-1.018** (0.452)	-1.053*** (0.362)	-4.663*** (0.644)	0.708 (1.704)	-1.169*** (0.430)	-1.104*** (0.347)
Observations	131	131	131	131	113	136	136	136
Panel B: HDD December								
Time trend (years)	-1.033* (0.613)	-4.985* (2.621)	-1.135** (0.498)	-0.900** (0.453)	-1.085 (0.674)	-1.932 (2.193)	-1.222*** (0.454)	-1.650*** (0.446)
Observations	115	115	115	115	108	136	136	136
Panel C: HDD January								
Time trend (years)	-0.242 (0.661)	-0.583 (2.414)	-1.838*** (0.405)	-2.784*** (0.438)	-0.391 (0.793)	0.008 (2.381)	-1.851*** (0.394)	-2.844*** (0.427)
Observations	134	134	134	134	114	136	136	136
Panel D: HDD February								
Time trend (years)	1.927*** (0.612)	-6.275** (2.437)	-1.945*** (0.421)	-1.575*** (0.456)	1.974*** (0.652)	-5.388** (2.322)	-1.980*** (0.394)	-1.757*** (0.428)
Observations	131	131	131	131	122	136	136	136
Panel E: HDD March								
Time trend (years)	-5.290*** (0.679)	-1.551 (2.247)	-2.483*** (0.427)	-1.615*** (0.420)	-5.531*** (0.737)	-3.222 (2.148)	-2.749*** (0.410)	-1.868*** (0.403)
Observations	131	131	131	131	123	136	136	136
Airport FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Data	Futures	Station	RCP4.5	RCP 8.5	Futures	Station	RCP4.5	RCP 8.5
Years	Common	Common	Common	Common	Traded	All	All	All

Notes: Table replicates Panel B of Table 2 but estimates a separate trend for each month. It regresses heating degree days (HDD) for the winter months November-March in the years 2002-2018. Column (a) uses the average futures price 30 to 10 days before the start of each contract month, e.g., the average price between June 1 and June 20 for a July contract. Column (b) uses observed station level data for the month, while columns (c) and (d) use climate change forecasts in the NASA NEX-GDDP database under the RCP4.5 and RCP8.5 scenarios. Columns (1a)-(1b) estimate the trends for a consistent set of observations where futures data are available. Columns (2a)-(2b) conduct sensitivity checks to the included years. Column (2a) excludes years where the price did not change over that time frame as the contract might not have been traded and hence not reflect changes in market expectations. Columns (2b)-(2d) include all years even if futures data is not available.

A1 Data Appendix


Some data cleaning was necessary for the futures data because of “sticky fingers,” e.g., sudden price jumps by a factor of 10. For example, a price series was 91, 91, 910. We contacted Bloomberg about whether these were data entry errors, but they assured us that the data had been cleaned. Specifically, we rescaled

1. The January 2011 contract in DFW by $\frac{1}{10}$ if its price exceeded 6000.
2. The July 2011 contract for CVG by $\frac{1}{10}$ if its price exceeded 4000.
3. The September 2011 contract for ATL, LAS, LGA and SAC by $\frac{1}{10}$ if its price exceeded 1000.
4. The September 2011 contract in MSP by $\frac{1}{10}$ if its price exceeded 300.
5. The January, February, and March contracts in 2002 and 2003 for SAC by 10 if their price was below 60.
6. The November and December contracts in 2002 for SAC by 10 if their price was below 100.
7. The December 2016 contracts for SAC by 10 if their price was below 100.

Furthermore, we excluded observations


1. Where the price was zero
2. Where prices were missing for 30 days before the start date of a month. This occurred for some December contracts in 2003 and 2005.
3. Where prices never changed between two weeks (14 days) before the start of the month and the end of the month, i.e., there was no price change over the 45-day period.

Figure A1: Observed Cooling and Heating Degree Days versus Futures Prices



Notes: Figure replicates Figure 4, but includes the additional 8 CDD (12 HDD) futures contracts where prices did not move in the 14-day window spanning between 7 days before the end of the month or seven days after the end of the month. Left graph again displays a scatter plot of the 513 observed monthly cooling degree day (CDD) totals against the average futures price the week following the end of the summer months June-September. Right graph displays a scatter plot of the 642 observed monthly heating degree day (HDD) totals against the average futures price the week following the end of the winter months November-March. Both graphs use data for the 18 years (2001-2018) across the eight airports. Some months have missing data.

Figure A2: Nonparametric Futures Prices and Weather for July-September



Notes: Figure replicates the left column of Figure 5 for CDD in the months July-September as data for June was not available in 2018. It estimates nonparametric trends using lowess regression on the average annual residual among the eight airports, i.e., the annual residuals after taking out airport fixed effects (means) are averaged by year before running lowess on the annual observations. Red line shows the results for the observed station data, the magenta line partials out ocean oscillation indices that affect weather. The green line uses futures prices before the weather is realized, and the blue lines use model output from NASA NEX-GDDP.

Table A1: Futures Returns on Weather Shocks - Including Weeks with no Price Change

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Cooling Degree Days June - September							
Shock in week t	0.156*** (0.021)	0.168*** (0.022)	0.158*** (0.021)	0.170*** (0.022)	0.178*** (0.022)	0.193*** (0.026)	0.203*** (0.027)
Shock in week t+1	0.372*** (0.024)	0.381*** (0.024)	0.373*** (0.023)	0.381*** (0.024)	0.386*** (0.024)	0.392*** (0.026)	0.418*** (0.028)
Shock in week t+2	0.048* (0.026)	0.057** (0.026)	0.048* (0.026)	0.051* (0.026)	0.059** (0.027)	0.051* (0.030)	0.065** (0.029)
Shock in week t+3	-0.046* (0.024)	-0.037 (0.024)	-0.045* (0.024)	-0.047* (0.025)	-0.033 (0.026)	-0.062** (0.030)	-0.060* (0.032)
Observations	3778	3778	3778	3778	3778	3778	3778
Clusters	67	67	67	67	67	67	67
Fixed Effects		4	17	67	133	513	520
Panel B: Heating Degree Days November - March							
Shock in week t	0.053 (0.041)	0.060 (0.041)	0.052 (0.041)	0.055 (0.043)	0.051 (0.044)	0.054 (0.053)	0.049 (0.054)
Shock in week t+1	0.325*** (0.035)	0.330*** (0.035)	0.324*** (0.035)	0.323*** (0.036)	0.324*** (0.036)	0.320*** (0.039)	0.328*** (0.038)
Shock in week t+2	0.232*** (0.046)	0.239*** (0.046)	0.231*** (0.046)	0.228*** (0.046)	0.230*** (0.046)	0.219*** (0.048)	0.220*** (0.047)
Shock in week t+3	0.056 (0.038)	0.065* (0.038)	0.055 (0.038)	0.048 (0.039)	0.054 (0.040)	0.033 (0.044)	0.039 (0.046)
Observations	4694	4694	4694	4694	4694	4694	4694
Clusters	85	85	85	85	85	85	85
Fixed Effects		5	17	85	134	642	649
Fixed Effects	-	M	Y	YM	AY	AYM	AYM,W

Notes: Tables replicates Table 1, but also includes weeks that did not see any price change. It regresses weekly cooling (CDD) and heating (HDD) degree day future returns (Friday-Friday) on deviations in degree days from seasonal averages. Both the contemporaneous week as well as three leads are included. Errors are clustered at the delivery month (e.g., November 2010). Columns control for different fixed effects across the 17 years in our sample (2002-2018), months (June-September for CDD and November-March for HDD) and 8 airports for which contracts are available. Fixed effects are: monthly (M), yearly (Y), year-month (YM), airport-year (AY), airport-year-month (AYM), as well as AYM plus week relative to contract maturity fixed effects (AYM,W).

Table A2: Futures Returns on Weather Shocks - Excluding Weeks Where Leads Fall Beyond the End of the Contract Month

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: Cooling Degree Days June - September							
Shock in week t	0.163*** (0.027)	0.178*** (0.028)	0.161*** (0.027)	0.170*** (0.028)	0.186*** (0.030)	0.219*** (0.036)	0.220*** (0.037)
Shock in week t+1	0.418*** (0.024)	0.434*** (0.024)	0.420*** (0.024)	0.438*** (0.025)	0.444*** (0.025)	0.469*** (0.030)	0.472*** (0.030)
Shock in week t+2	0.025 (0.024)	0.043* (0.024)	0.024 (0.024)	0.030 (0.025)	0.053** (0.026)	0.049 (0.031)	0.047 (0.031)
Observations	2743	2743	2743	2743	2743	2743	2743
Clusters	67	67	67	67	67	67	67
Fixed Effects		4	17	67	133	513	518
Panel B: Heating Degree Days November - March							
Shock in week t	0.016 (0.046)	0.031 (0.047)	0.011 (0.047)	0.023 (0.050)	0.002 (0.050)	-0.004 (0.067)	-0.008 (0.067)
Shock in week t+1	0.345*** (0.040)	0.356*** (0.040)	0.341*** (0.040)	0.338*** (0.041)	0.340*** (0.041)	0.314*** (0.046)	0.316*** (0.045)
Shock in week t+2	0.207*** (0.035)	0.215*** (0.035)	0.200*** (0.035)	0.189*** (0.037)	0.192*** (0.035)	0.145*** (0.045)	0.139*** (0.045)
Observations	3493	3493	3493	3493	3493	3493	3493
Clusters	85	85	85	85	85	85	85
Fixed Effects		5	17	85	134	642	647
Fixed Effects	-	M	Y	YM	AY	AYM	AYM,W

Notes: Tables replicates Table 1, but does not include weeks where the leads fall beyond the end of the contract months (they were previously set to zero). It regresses weekly cooling (CDD) and heating (HDD) degree day future returns (Friday-Friday) on deviations in degree days from seasonal averages. Both the contemporaneous week as well as two leads are included. Errors are clustered at the delivery month (e.g., November 2010). Columns control for different fixed effects across the 17 years in our sample (2002-2018), months (June-September for CDD and November-March for HDD) and 8 airports for which contracts are available. Fixed effects are: monthly (M), yearly (Y), year-month (YM), airport-year (AY), airport-year-month (AYM), as well as AYM plus week relative to contract maturity fixed effects (AYM,W). Tables excludes observations that did not see a price change during the entire week.