Market Expectations of a Warming Climate

Wolfram Schlenker[1,2,*] and Charles A. Taylor[1]

Online Appendix

List of Figures

List of Tables

[1] School of International and Public Affairs, Columbia University, 420 West 118th St., New York, NY 10027.
[2] National Bureau of Economic Research (NBER), 1050 Massachusetts Ave., Cambridge, MA 02138.
[*] Corresponding author. E-mail address: wolfram.schlenker@columbia.edu

A1 Data Appendix

The Chicago Mercantile Exchange (CME) has a process in place to update weather futures 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). Some clearly erroneous prices were excluded, specifically, any price of 0 or prices that were 1 on March 17, 2017 for March 2017 contracts.

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. Specifically, we made the following adjustments:

1. We rescaled by a factor of $\frac{1}{10}$ the following 2011 contracts
    - January at DFW if its price exceeded 6000.
    - July at CVG by if its price exceeded 4000.
    - September at MSP its price exceeded 300.
    - September at ATL, LAS, LGA, and SAC if prices exceeded 1000.
2. We rescaled by a factor of 10 the following 2005 contracts
    - December at CVG and SAC if its price fell below 100.
3. We excluded the following contracts
    - Any SAC contract before April 2003.
    - Any SAC contract in 2020.
    - March SAC contract in 2005.
    - December MSP in 2005.
    - November contracts at SAC, ATL, LAS in 2011.

Figure A6 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.

We can infer 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 flag contracts whose 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. Similarly, we flag contracts that were not traded at least 30 days prior to the start of a contract month. We find similar results when they are included. Volume data from CME on a subset of contracts is summarized in Appendix Table A1.

Leap years are treated differently in various datasets. For example, some of the 21 climate models in the NEX-GDDP data 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 $\frac{28}{29}$.

A2 Oceanic Oscillation Indices

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 be driving the trends in observed warming. One of the strongest and most famous is the El Niño - Southern Oscillation (ENSO), 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 oceanic-atmospheric phenomena, we partial out the effect of these oceanic factors on temperature.

In a first step, we replicate the seasonality regression of equation (3) of the main text. We again regress daily average temperature T_{ad} at airport a on day d on a constant α_a as well as a flexible spline that is a function f of the day of the year.

$$T_{ad} = \alpha_a + \beta_a f(d) + \gamma_a y(d) + \sum_{k=1}^{5} \sum_{s=1}^{3} \sum_{\tau=3}^{0} \delta_{ks\tau} o_{ks[m(d)-\tau]} + \epsilon_{ad} \tag{A1}$$

The additional variables are the $o_{ks[m(d)-\tau]}$ terms: the subscript $k = 1 \ldots 5$ indicates five monthly oceanic oscillation indices: ENSO (El Niño - Southern Oscillation), NAO (North Atlantic Oscillation), PDO (Pacific Decadal Oscillation), PNA (Pacific/ North American Teleconnection Pattern), and AO (Arctic Oscillation).[1] The subscript s allows the coefficient to be (i) the same throughout the year (pooled effects), (ii) be different for the summer months (June - September) or (iii) be different for the winter months (November - March). The subscript $[m(d) - \tau]$ indicates the month m in which the day falls and whether to take

[1] Monthly ENSO values were downloaded for Region 3.4 from https://www.cpc.ncep.noaa.gov/data/indices/sstoi.indices, while values for the remaining indices was obtained from https://www.ncdc.noaa.gov/teleconnections/.

lags τ of the monthly oscillation indices. We purposely fit a very flexible model: each airport is allowed to depend on 60 oceanic variables (5 oceanic indices $\times$ 3 seasons $\times$ 4 lags $(0 \ldots 3)$). We rely on machine learning to pick the optimal model, specifically LASSO regression using the Extended Bayesian Information Criteria (EBIC). Appendix Table A5 summarizes which of the 60 oceanic variables are selected by the LASSO routine. There is a wide range: the procedure selects 6 out of the 60 variables for Las Vegas, but 21 of the 60 for Minneapolis.

In a second step, we then partial out the effect of the observed oceanic oscillation indices

$$\widehat{T_{ad}} = T_{ad} - \sum_{k=1}^{5} \sum_{s=1}^{3} \sum_{\tau=0}^{3} \widehat{\delta_{ks\tau}} o_{ks[m(d)-\tau]} \tag{A2}$$

Since the indices are all normalized to zero, this procedure provides a counterfactual where the index is at its mean rather of the observed anomaly. This allows us to derive heating and cooling degree days after accounting for the influence of the oceanic indices.

Figure A11 shows the nonparametric trend for the raw station-level data in red and in magenta after partialling out the oceanic indices. Note how the magenta line shows a smoother trend, especially for cooling degree days in the left graph. In other words, the standard deviation around the trend is lower for the magenta line than the red line because common oceanic shocks are accounted for. But most importantly, the overall trend is hardly changed which means that the observed warming trend is not due to decadal oscillation patterns.

A3 Prior Trends

We contrast the trend from 2002 to 2020 in our analysis to nonparametric trends in weather station data from 1900 to 2020. Some of the weather stations do not go back to 1900, and we hence replace the weather station data with the fine-scaled (2.5 x 2.5 miles) weather dataset from Schlenker & Roberts (2009) that keeps the weather stations constant. Similar to the previous analysis, we remove airport-by-month fixed effects (e.g., July contracts for Atlanta) from the monthly degree days series. We then average the monthly residuals over the four summer months (June - September) or the five winter months (November - March). The results are shown in Appendix Figure A12. As other authors have emphasized (Burke & Emerick 2016), recent global warming trends become apparent in the data around 1980. We find the same for the eight cities in our sample: from 1940 to 1980 the nonparametric black line is rather flat. Then 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. We also find that there was an early century warming trend from 1900 to 1940, which then plateaued for several decades. The early century warming trend is primarily attributable to natural forcings, while the trend in the latter part of the century is attributable to anthropogenic forcings (Meehl et al. 2004).

References

Burke, Marshall, and Kyle Emerick. 2016. "Adaptation to Climate Change: Evidence from US Agriculture." *American Economic Journal: Economic Policy*, 8(3): 106–140.

Meehl, Gerald A., Warren M. Washington, Caspar M. Ammann, Julie M. Arblaster, T. M. L. Wigley, and Claudia Tebaldi. 2004. "Combinations of Natural and Anthropogenic Forcings in Twentieth-Century Climate." *Journal of Climate*, 17(19): 3721–3727.

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.

Tukey, J. W. 1977. *Exploratory Data Analysis.* Addison–Wesley.

Zebiak, Stephen E., and Mark A. Cane. 1987. "A Model El Niño-Southern Oscillation." *Monthly Weather Review*, 115(10): 2262–2278.

Figure A1: Location of Eight Airports in Sample

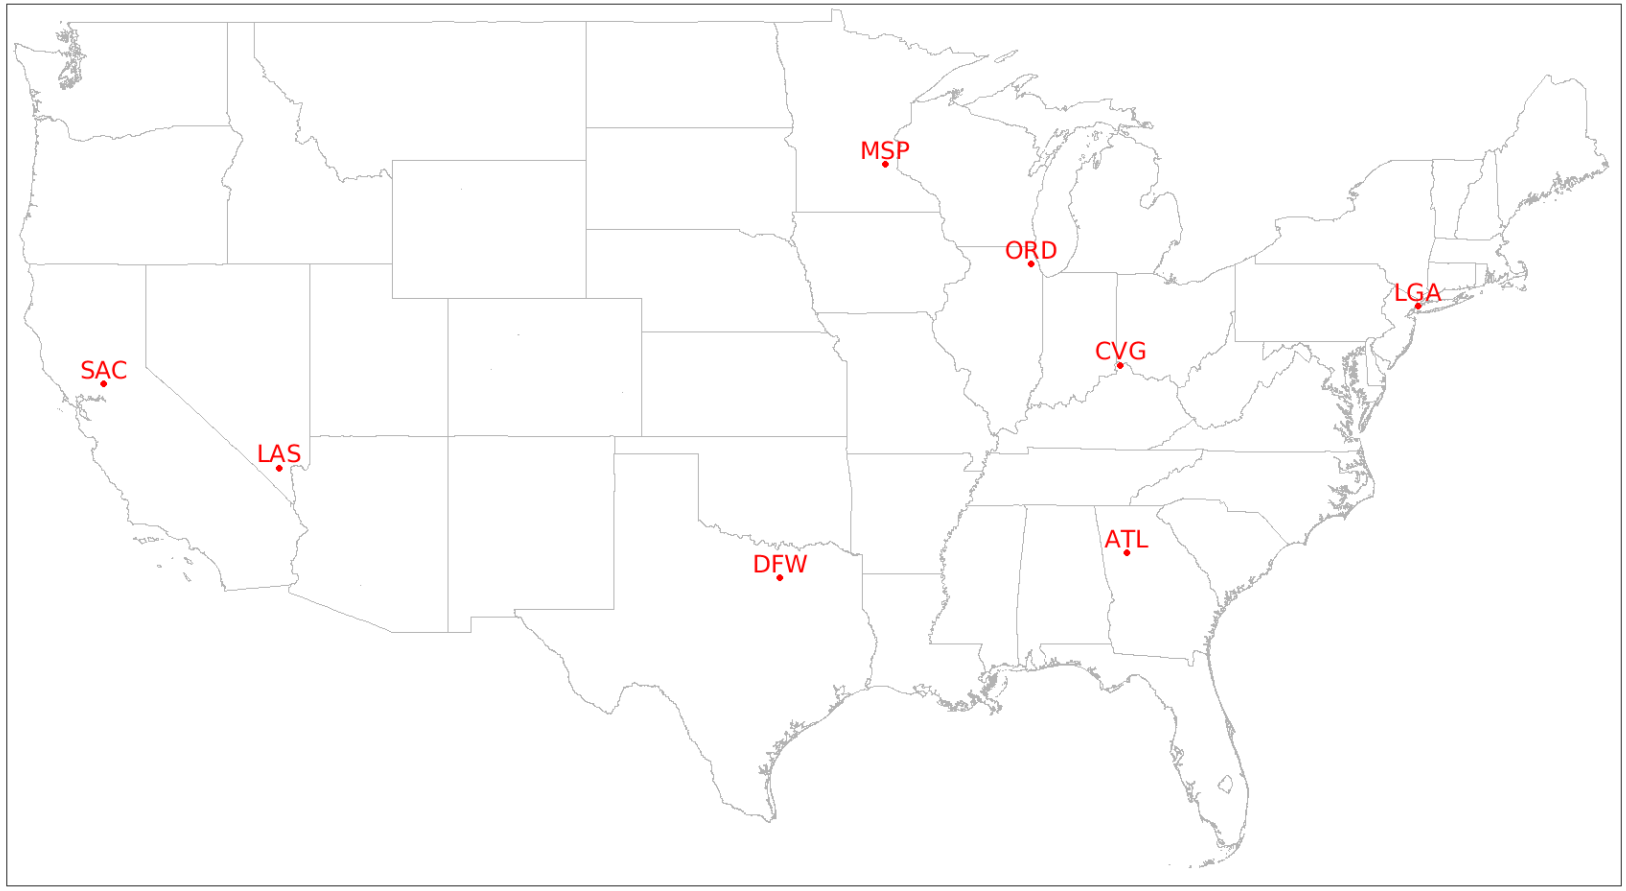

This figure displays the location of the eight weather stations at US airports in our sample that offered weather derivatives. They are from north to south: Minneapolis - Saint Paul (MSP), Chicago O'Hare (ORD), New York LaGuardia (LGA), Cincinnati - Northern Kentucky (CVG), Sacramento (SAC), Las Vegas (LAS), Atlanta (ATL), and Dallas - Fort Worth (DFW). The four airports in the northeastern quadrant are MSP, ORD, LGA, and CVG.

Figure A2: Future Prices Around Maturity

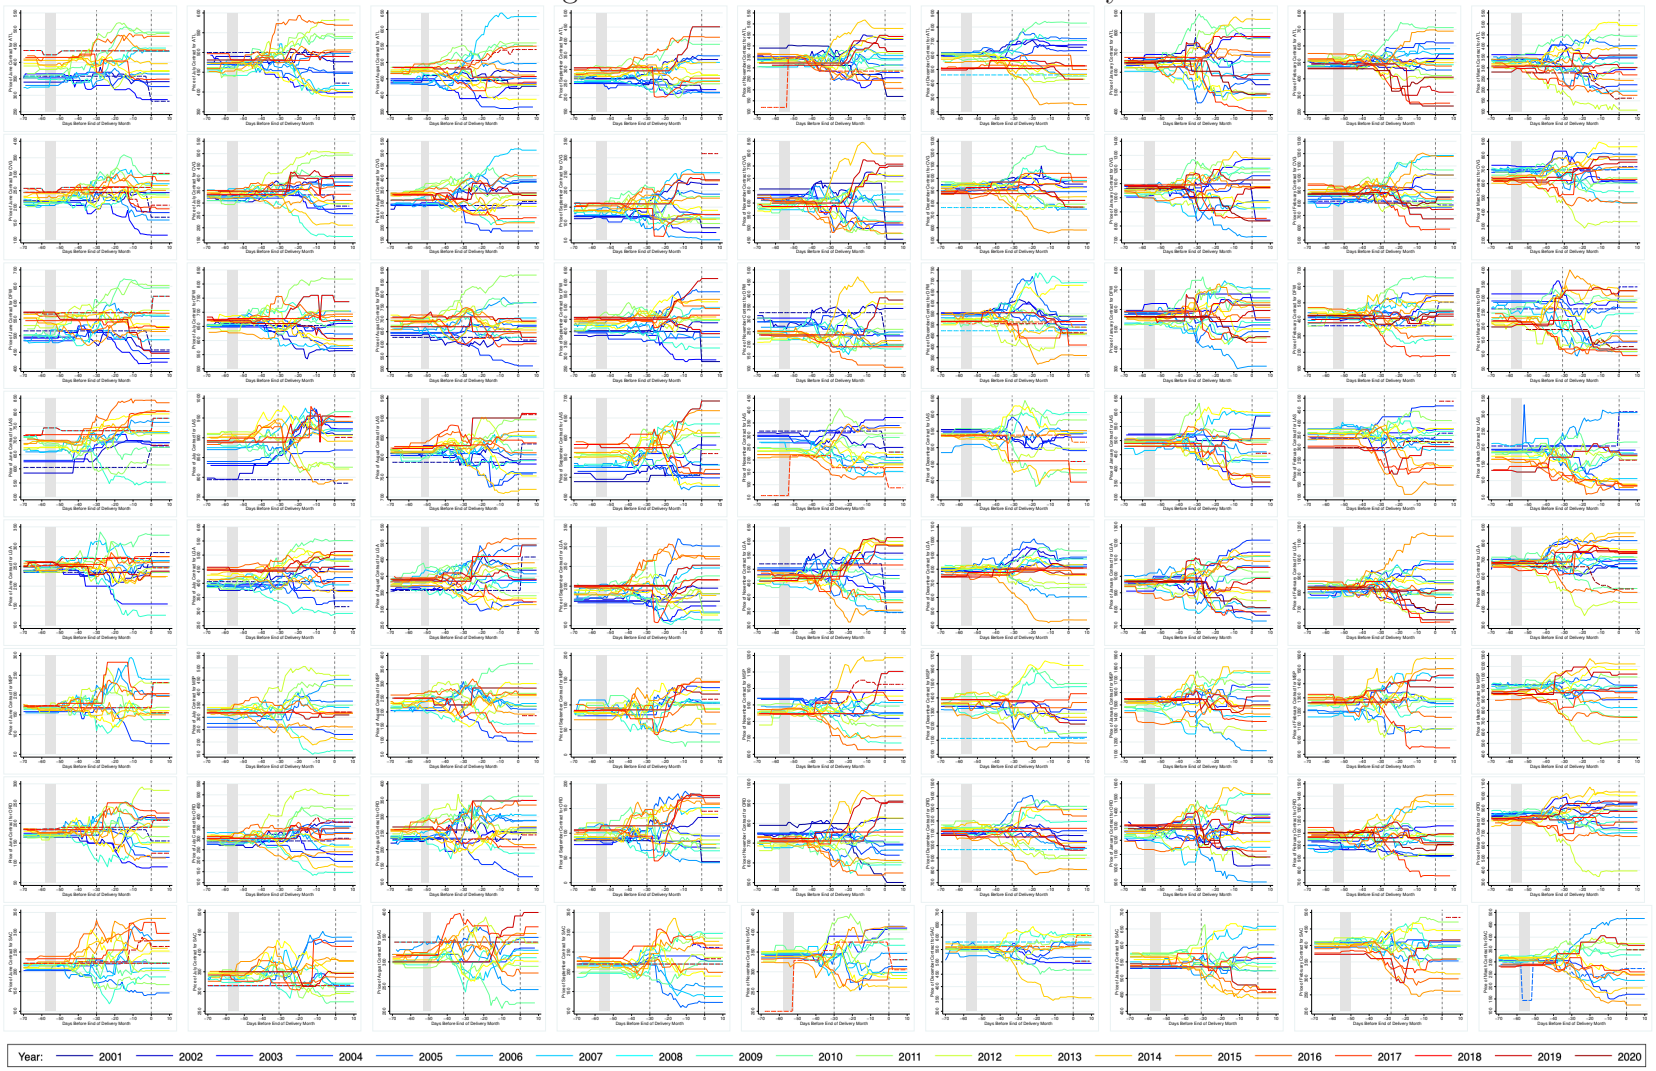

This figure displays the time series of futures prices around maturity. Day 0 is the end of the month on which the weather derivative is based, e.g., day 0 for a June contract is June 30. Rows are the eight airports for which weather contracts are offered throughout our sample period. The first four columns show cooling degree days for the months June, July, August, and September, while the last five columns show heating degree days for the months November, December, January, February, and March. 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.

Figure A3: Contracts in Dataset and Fraction of Days With Price Changes

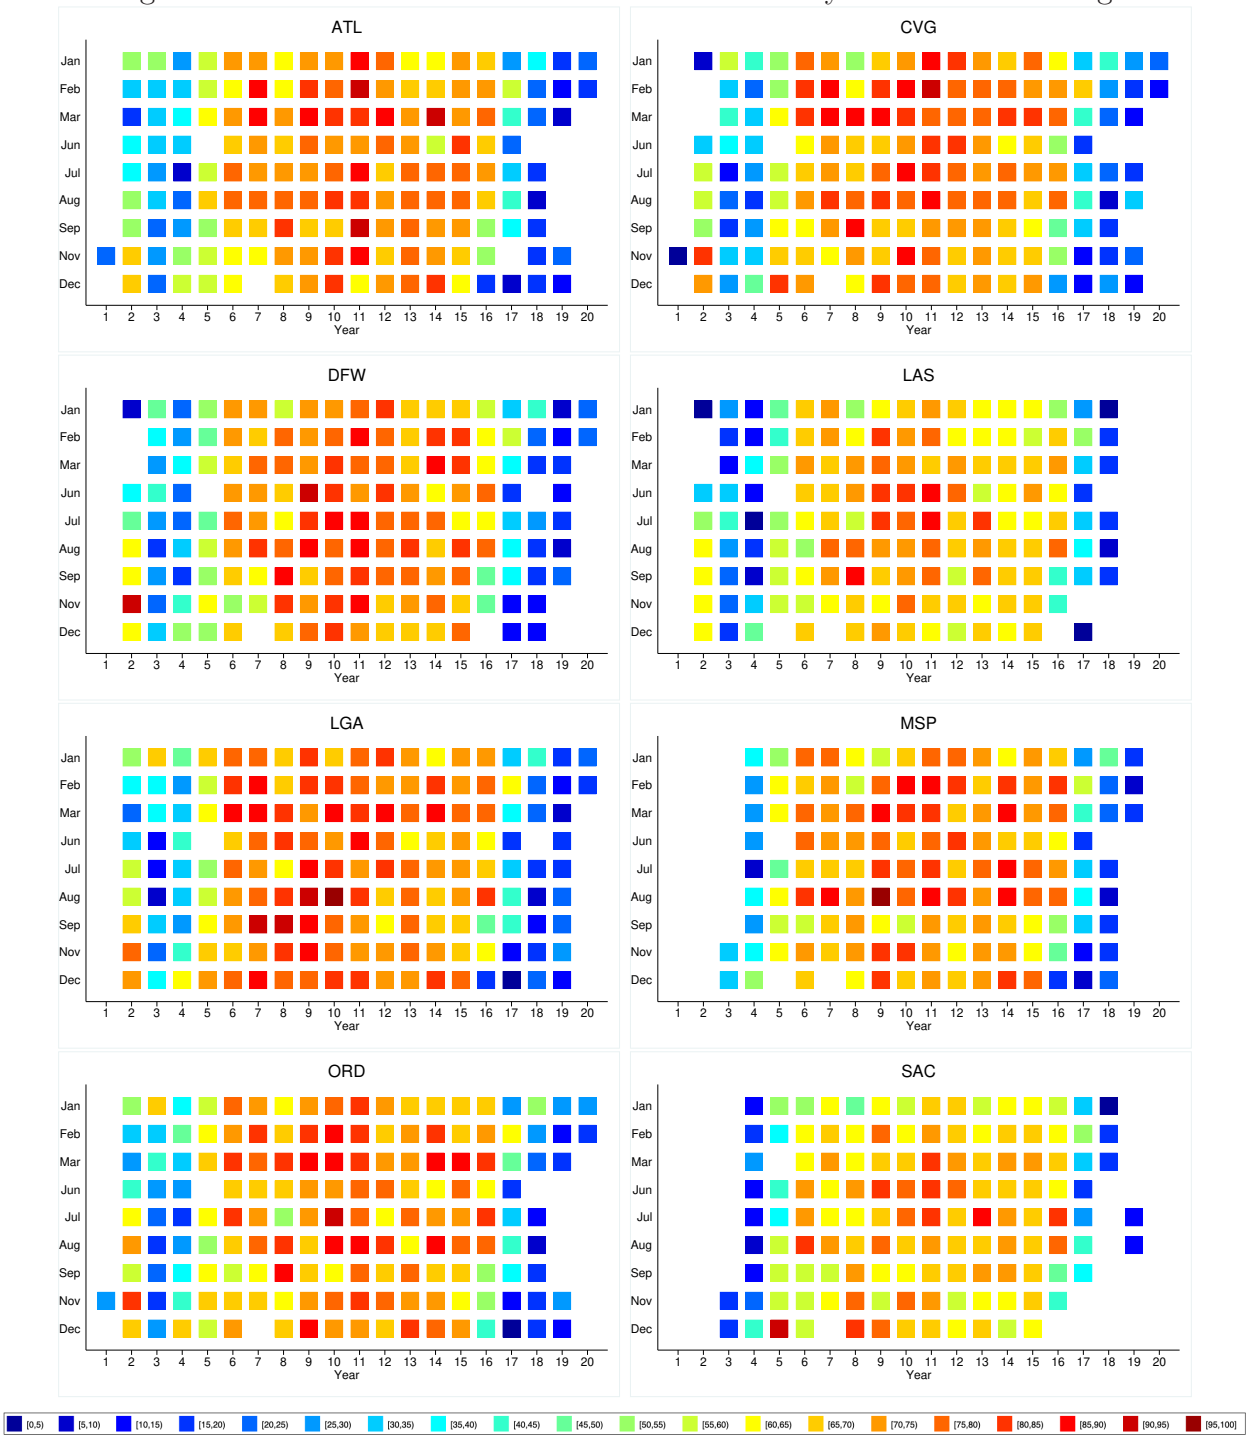

This figure shows what contracts are used in our regression and how often they report price changes. Colors indicate the fraction of trading days in which a future contract had a price change in the Bloomberg data for the two-months period that includes the contract month as well as the month prior. Only contracts that are included in our analysis after passing quality checks are included as color-coded squares, i.e., a missing square indicates the contract is not part of our regressions.

Figure A4: Future Prices Around Maturity - Data from CME

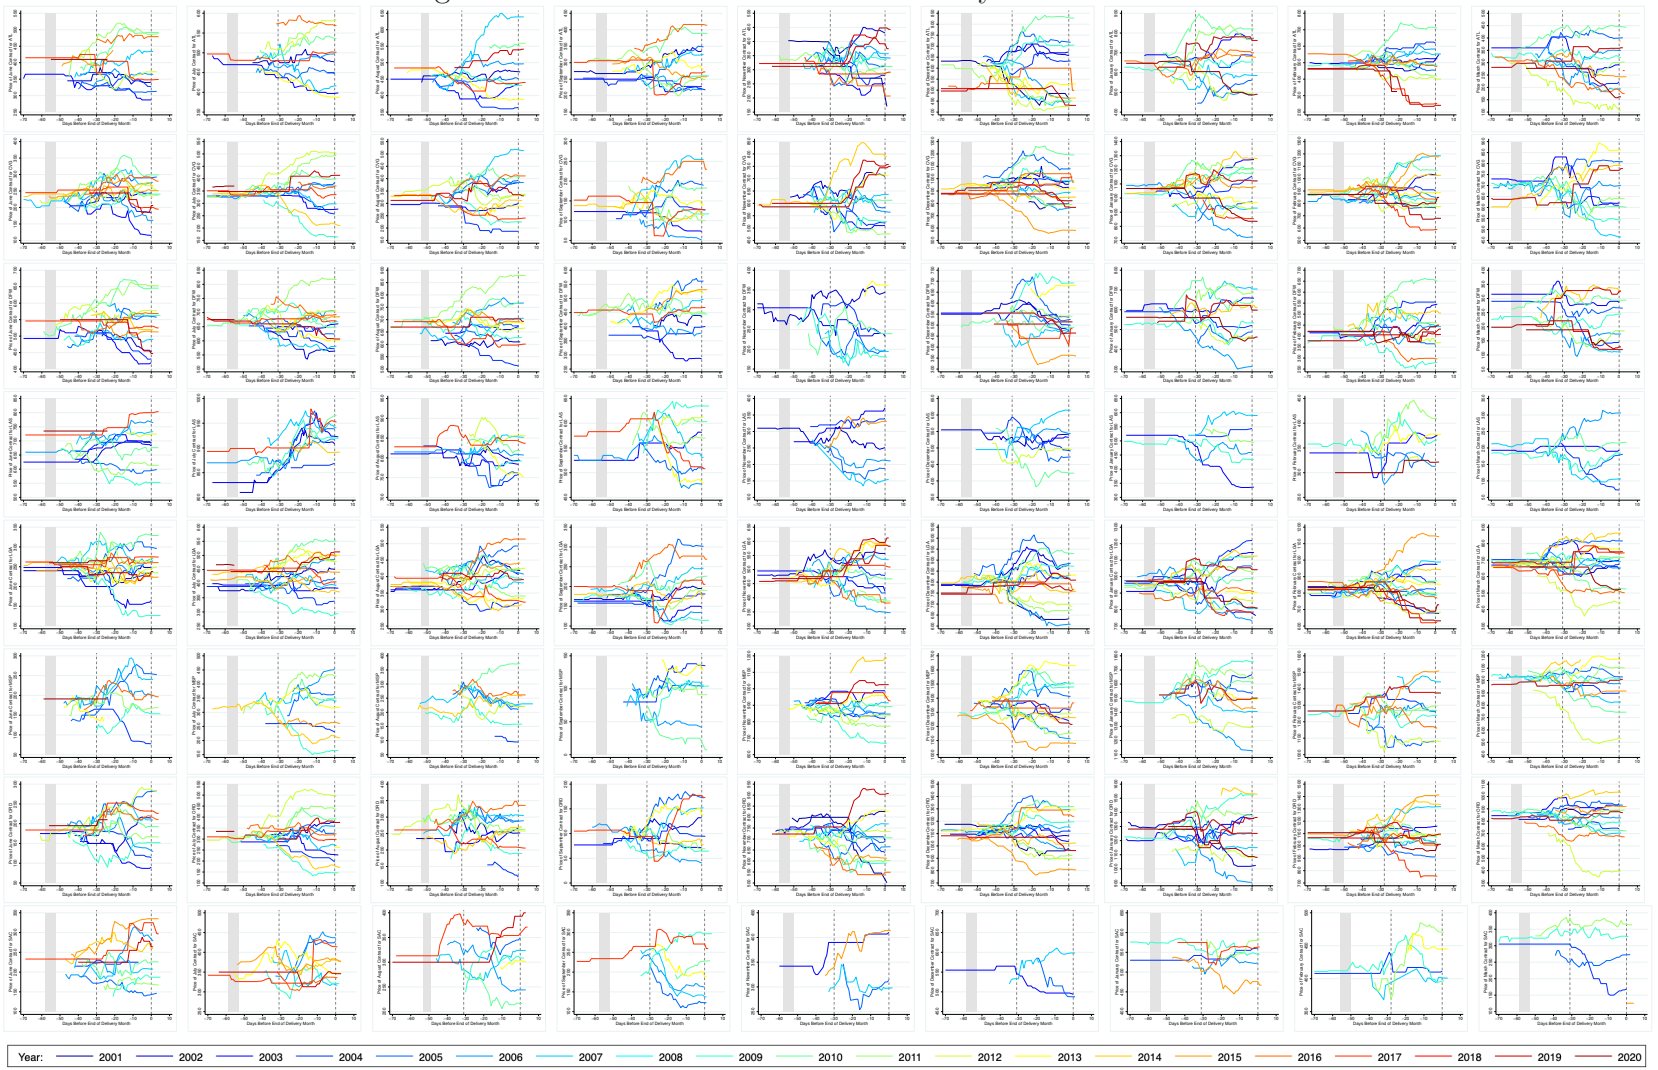

This figure replicates Appendix Figure A2, but uses data from CME instead of Bloomberg terminals.

Figure A5: Degree Days by Month and Weather Dataset for Eight Airports

Panel A: Cooling Degree Days (CDD)

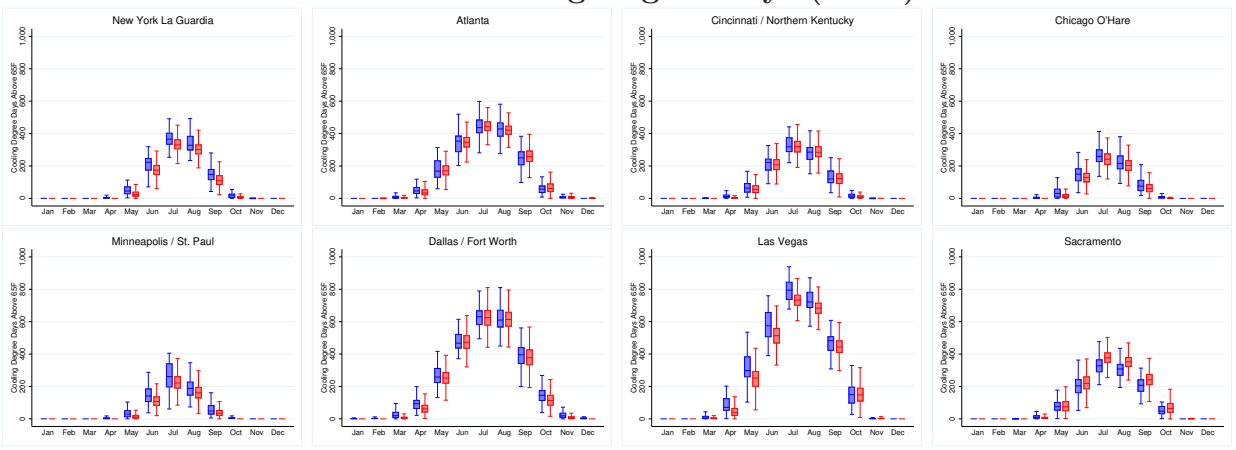

Panel B: Heating Degree Days (HDD)

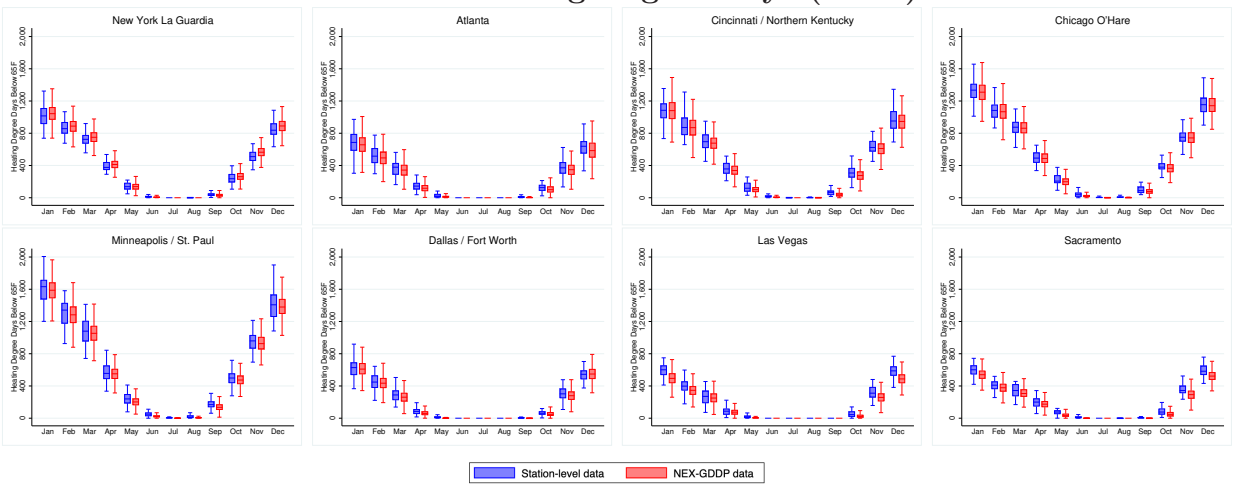

This figure compares degree days at weather stations and in climate model grids. Panel A displays box plots of cooling degree days, i.e., the amount average daily temperatures exceed 65°F, summed over all days of a month. Panel B gives the analogous amount that average temperatures fall below 65°F. The 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 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 Tukey (1977).

Figure A6: Degree Days At Weather Station versus Futures Prices at Settlement

Panel A: Excluding Contracts With No Price Change Around Settlement

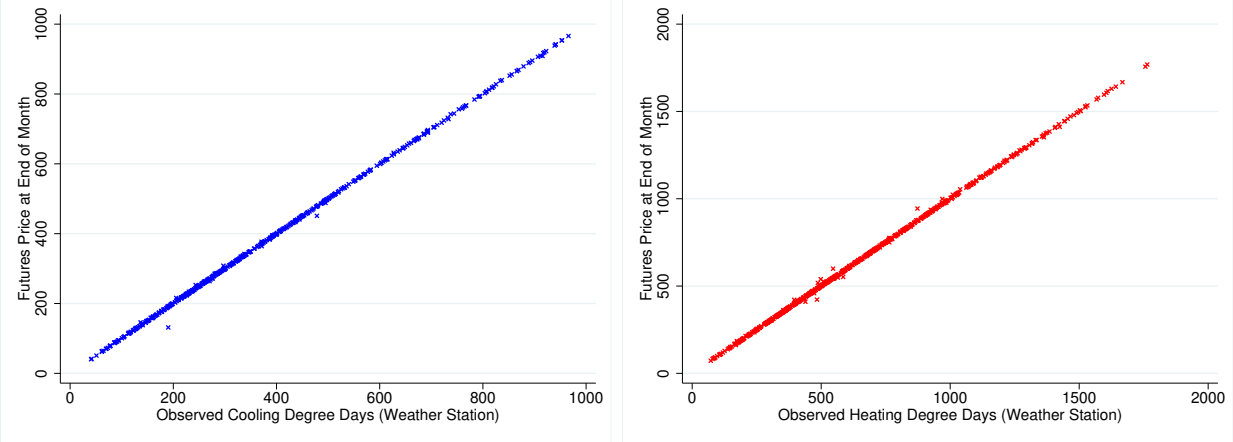

Panel B: Including Contracts With No Price Change Around Settlement

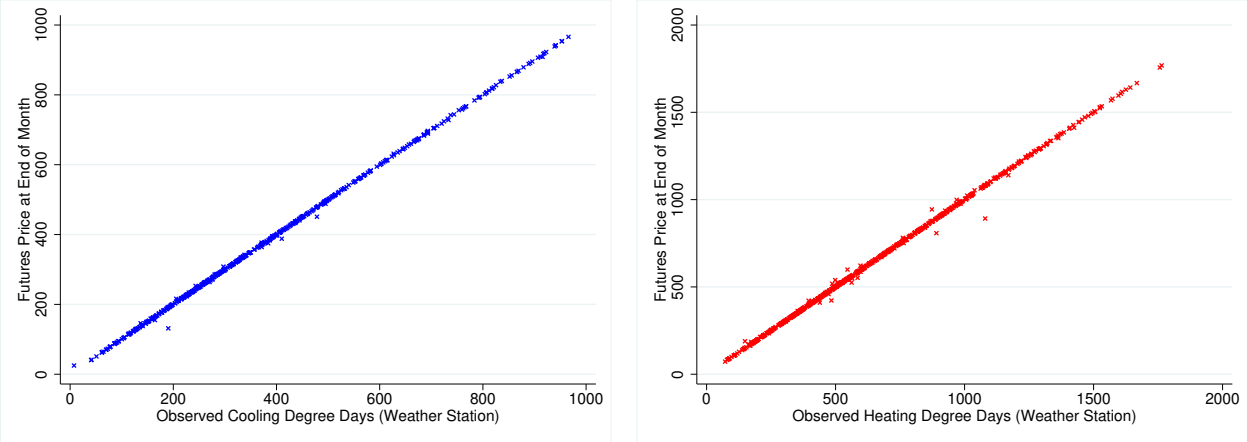

This figure displays scatter plot of the observed monthly cooling degree day (CDD) totals against the average futures price the week following the end of the contract month for June - September in the left column. The right column displays scatter plots of the observed monthly heating degree day (HDD) totals against the average futures price the week following the end of the contract month for November - March. Both graphs use data from winter 2001/2002 through winter 2019/2020 across the eight airports. Top row (panel A) excludes five CDD (11 HDD) futures contracts where prices did not change in the 14-day window spanning between 7 days before the end of the contract month or seven days after the end of the contract month, as it is not clear that they reflect final totals. This resulted in 517 CDD and 665 HDD observations. Bottom row (panel B) includes those observations.

Figure A7: Seasonality in Average Temperature

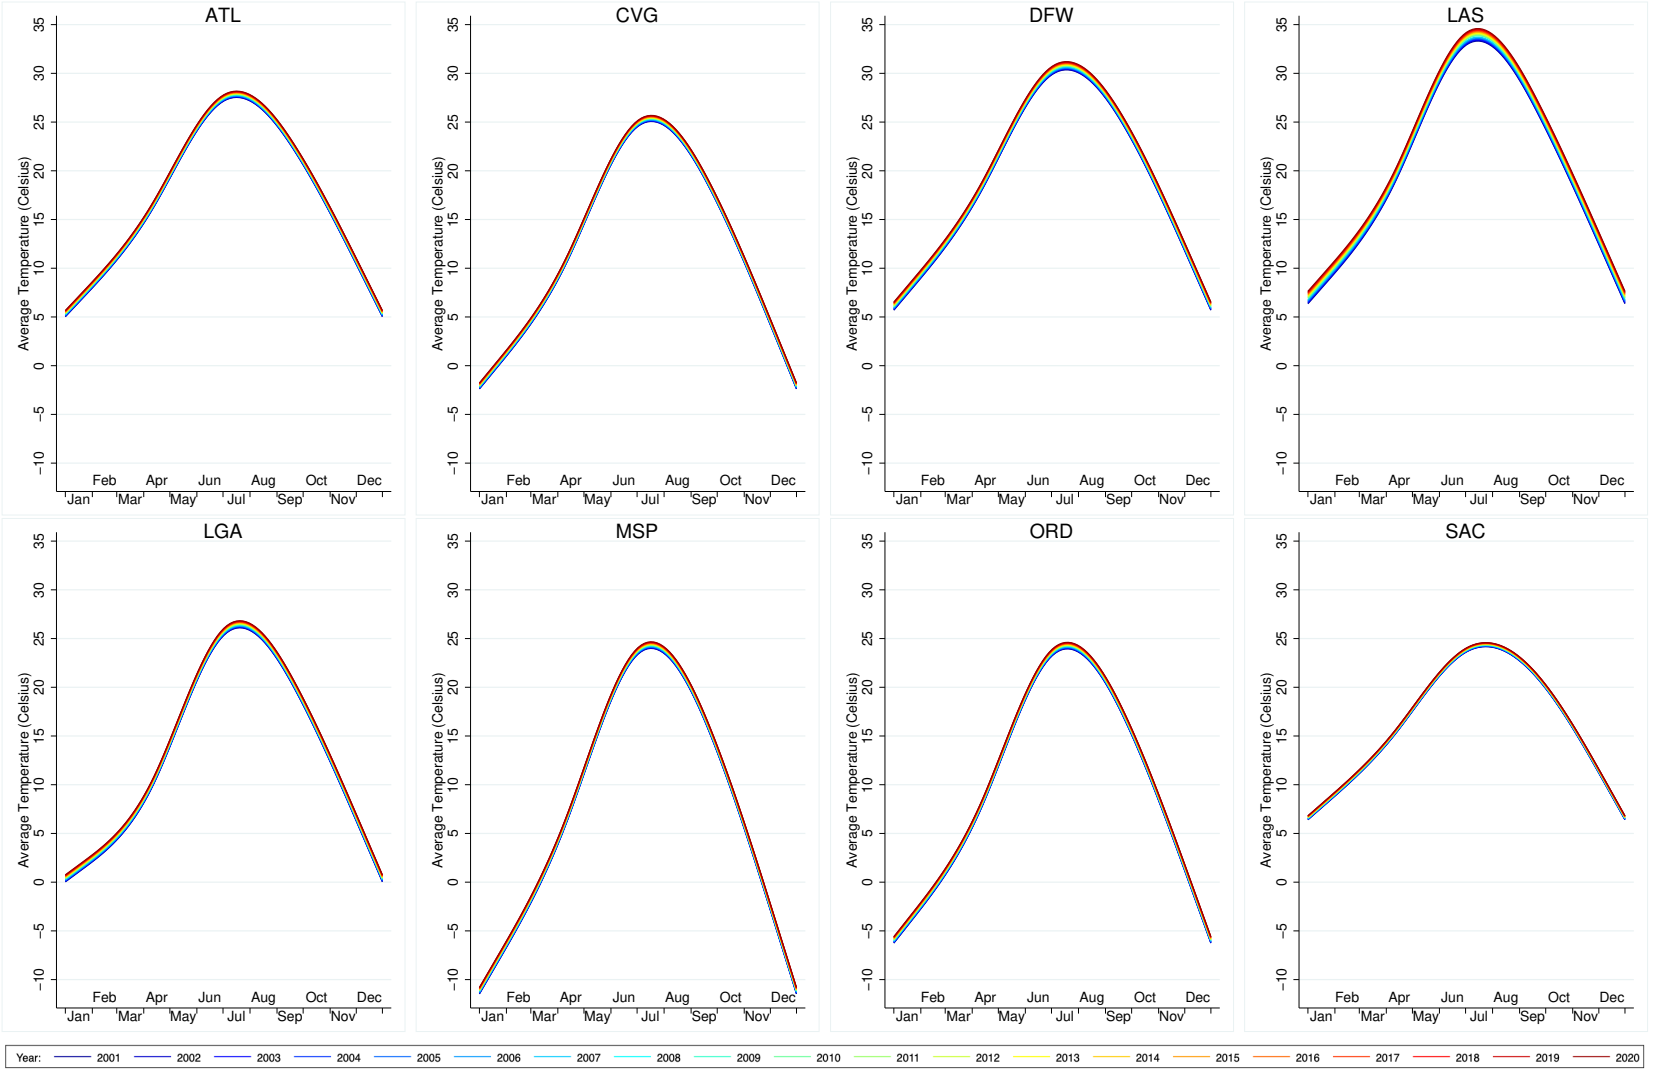

This figure displays the seasonality of average temperature over the year at each of the eight airports. The color-coding displays how average temperature have been trending upward. Both the seasonality and warming trend are estimated using data from 1980-2020. Because of leap years we normalize each year so January 1 is 0 and December 31st is 1. The seasonality is estimated using restricted cubic splines with five knots (0.05, 0.275, 0.5, 0.725 0.95) subject to the constraint that the average temperature at the end of the year equals the average temperature at the beginning of the year. Warming is estimated by including a linear time trend on the year a day falls in.

Figure A8: Capitalization of Weather Shocks - Separating CDD and HDD

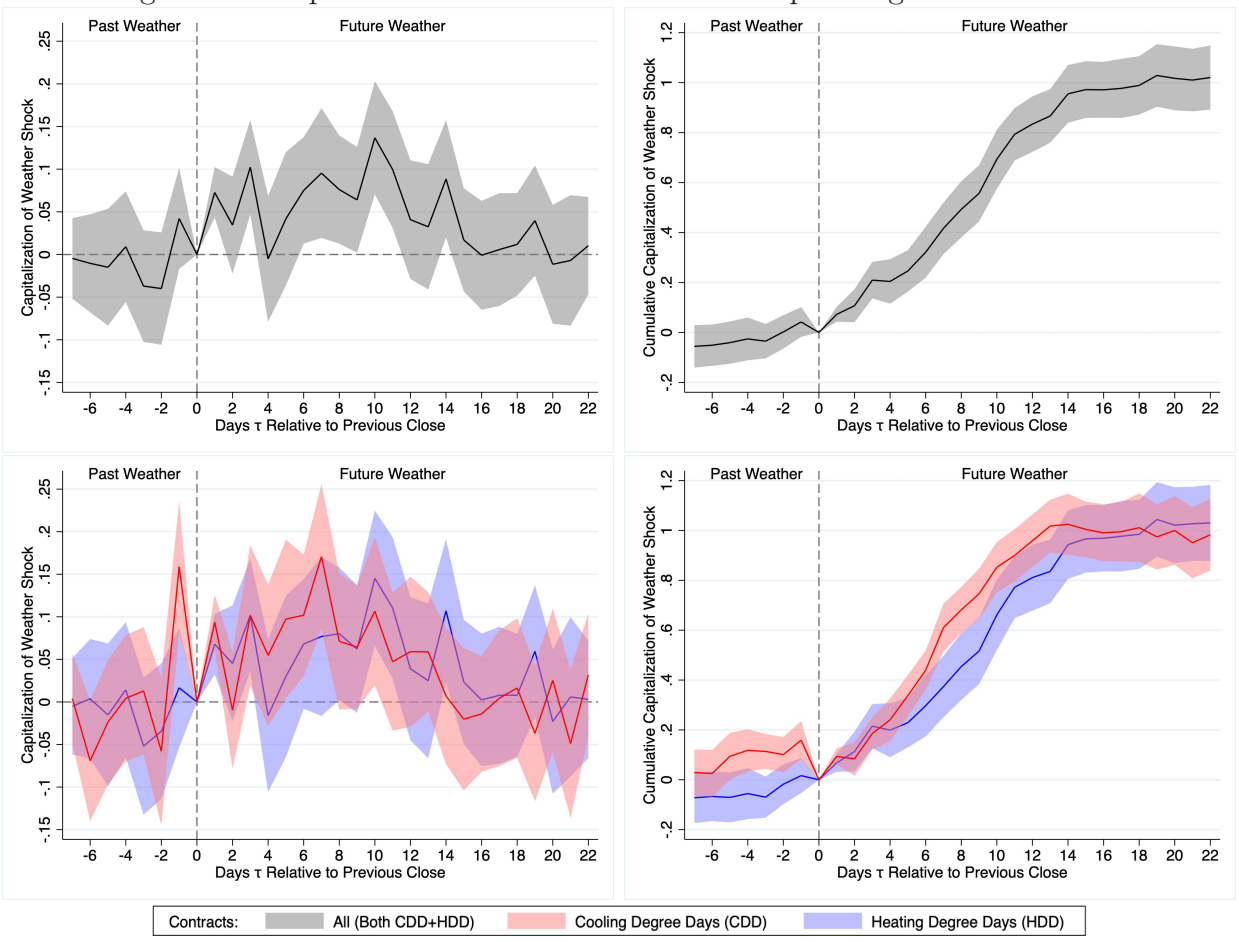

This figure replicates Figure 2 of the main paper in the top row. It adds the bottom row, which splits the regression into cooling degree days for June - September (red) and heating degree days November - March (blue). The interpretation of the coefficients is otherwise the same as in Figure 2.

Figure A9: Capitalization of Weather Shocks - CME Data

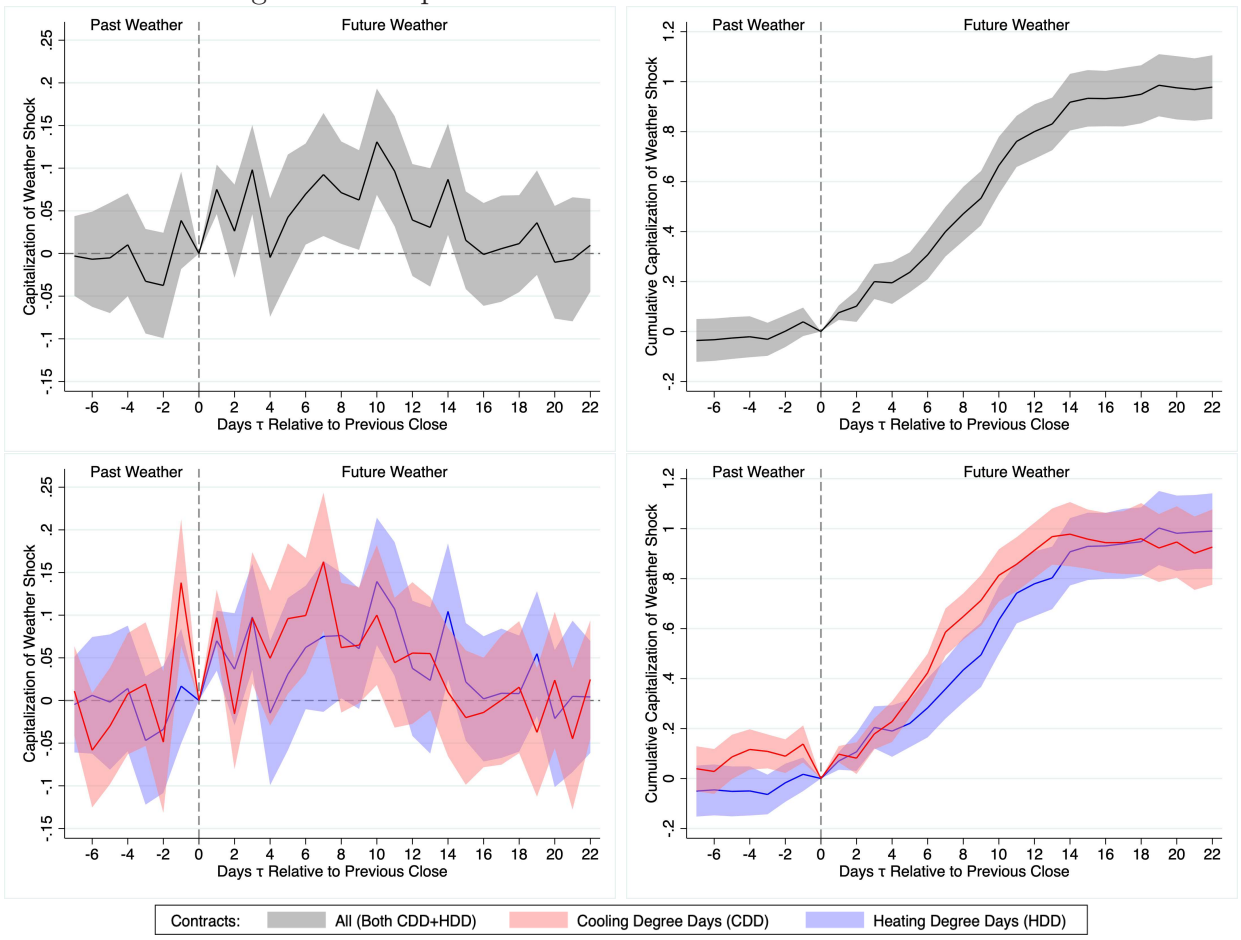

This figure replicates Appendix Figure A8, but uses CME data instead of the data from Bloomberg terminals. The interpretation of the coefficients is otherwise the same as in Figure 2 of the main paper.

Figure A10: Heterogeneity of Monthly Trends by Region

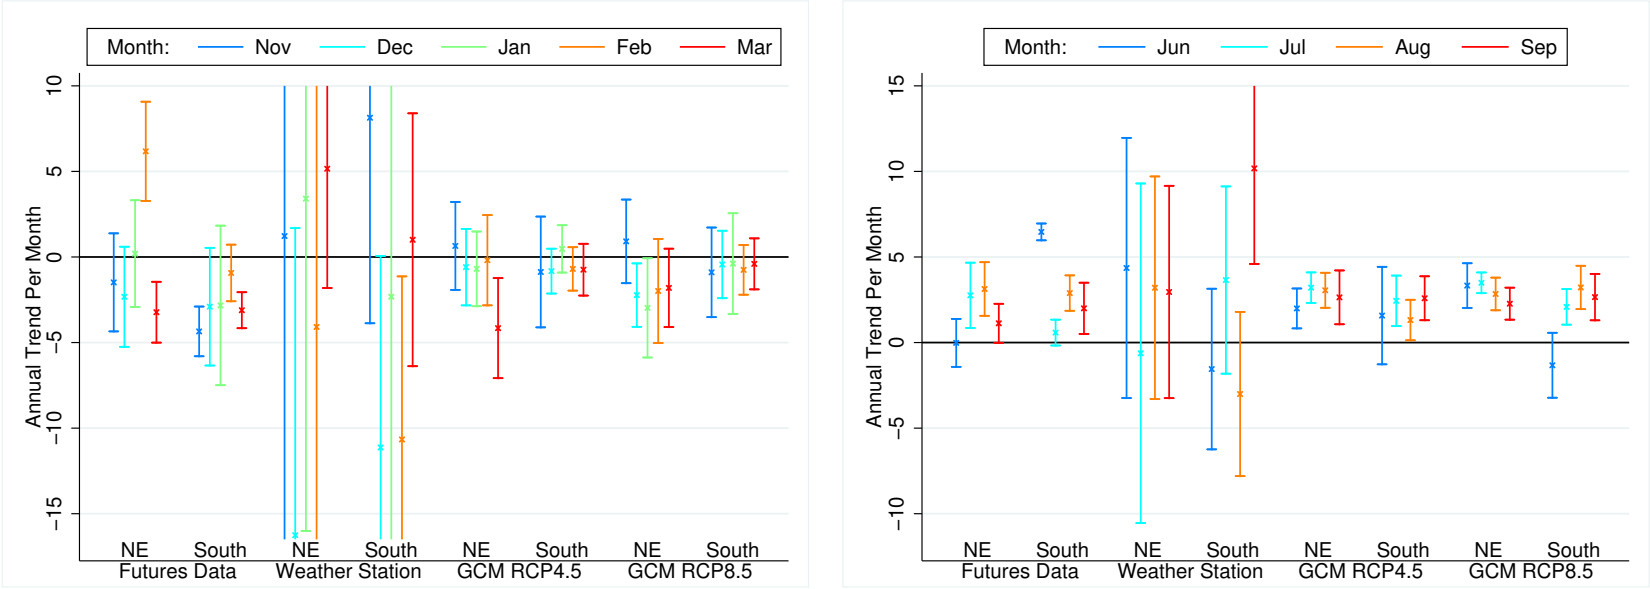

This figure displays the annual change in degree days by month and region, i.e., the coefficient estimate $\widehat{\beta_{r(a)m}}$ of the regression $D_{amy} = \alpha_{am} + \beta_{r(a)m} y + \epsilon_{amy}$, which is allowed to vary for two regions $r(a)$ and by month m. The eight airports are split into a northeastern quadrant consisting of Cincinnati - Northern Kentucky (CVG), New York LaGuardia (LGA), Minneapolis - Saint Paul (MSP), and Chicago O'Hare (ORD) and the rest, i.e., southern airports consisting of Atlanta (ATL), Dallas - Fort Worth (DFW), Las Vegas (LAS), and Sacramento (SAC). See Appendix Figure A1 for a map of the airports. The left graph summarizes annual trends for heating degree days (November - March), while the right graph shows the results for cooling degree days (June - September).

Figure A11: Nonparametric Trends in Futures Prices and Weather

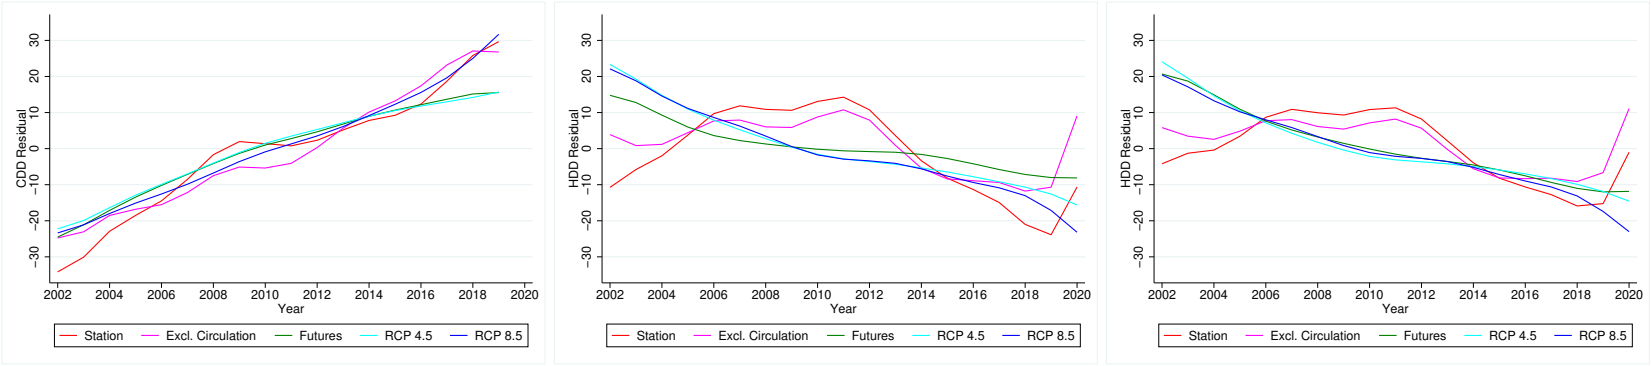

This figure estimating locally-weighted regression (lowess regression in STATA) on the annual average of the monthly residuals after removing airport-by-month fixed effects. The left two graphs use the same residuals as in Figure 3 of the main paper, but fit a lowess regression instead of a time trend using restricted cubic splines. The right graph excludes February for the northeastern airports. The green lines use futures prices four weeks before the start of the contract month. The red lines show the results for the observed weather station data. The blue lines use climate model output from NASA NEX-GDDP under RCP4.5 and RCP8.5 scenarios. The left column shows the results for cooling degree days in the summer months (June - September) and the right two columns for heating degree days in the winter months (November - March).

Figure A12: Nonparametric Trends in Past Weather

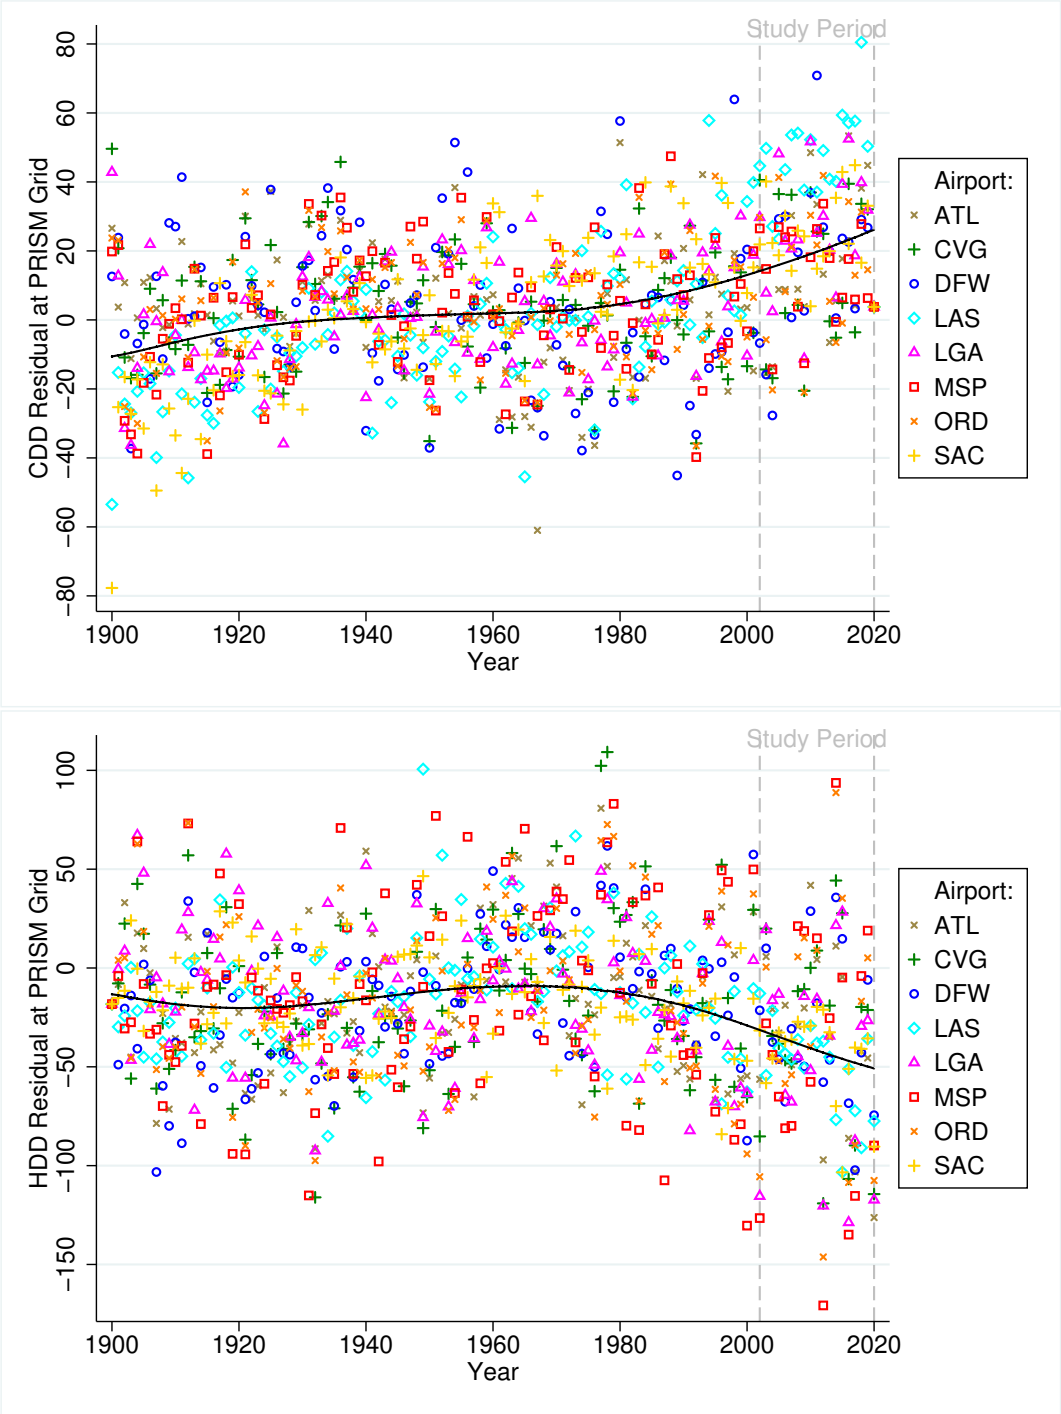

This figure plots average annual residuals in degree days after removing airport-by-month fixed effects, averaged over the four summer months (June - September) in the top graph and five winter months (November - March) in bottom graph. They are color-coded by each airport. A nonparametric lowess regression is added as a black line.

Table A1: Descriptive Statistics: Trade Volume

	5 Winter Months			4 Summer Months		
	Volume (1a)	**Value (1b)**	**Contracts (1c)**	**Volume (2a)**	**Value (2b)**	**Contracts (2c)**
			Panel A: By Year			
2002	127	2	15	993	7	21
2003	3036	47	30	4456	35	24
2004	4023	58	39	13451	69	28
2005	45555	728	35	75725	516	32
2006	73855	1144	38	41511	266	32
2007	68777	1085	34	69116	504	32
2008	95545	1534	40	73139	543	32
2009	50772	803	40	32653	197	29
2010	35747	574	31	31132	274	29
2011	42819	691	36	14886	99	25
2012	8196	126	27	5549	49	18
2013	5961	119	18	8608	75	20
2014	6350	108	25	5750	39	17
2015	3240	52	19	4600	41	13
2016	3603	48	19	3175	22	19
2017	6151	96	16	21250	181	28
2018	8594	111	23	.	.	.
2019	13000	161	12	8875	74	14
2020	19625	263	24	.	.	.
			B: By Airport			
ATL	60732	606	77	31958	253	52
CVG	53223	865	77	39865	229	59
DFW	37907	320	66	56003	653	59
LAS	11920	87	41	25272	385	46
LGA	185621	2884	86	153748	975	64
MSP	49934	1181	66	30007	133	39
ORD	89765	1757	81	50273	228	54
SAC	5874	50	27	27743	139	40

This table summarizes the trade volume in the dataset we obtained from CME, which is a subset of the contracts included in the main analysis. Panel A summarizes the data by season, while panel B summarizes the data by airport. The first three columns (1a)-(1d) present the data for the five months November - March with heating degree days. Seasons are shown in the year in which the season ends, e.g., the winter November 2001-March 2002 is shown in the row for 2002. The last three columns (2a)-(2c) present the data for the four summer months June - September with cooling degree days. Columns (a) give the volume (number) of trades, columns (b) give the value of those trades (volume $\times$ price $\times$ 20) in million dollars, column (c) gives the number of contracts for which positive volume data was provided.

Table A2: Futures Price Changes as Weather Forecast

Airports	All	ATL	CVG	DFW	LAS	LGA	MSP	ORD	SAC
Panel A1: Weather Shocks for Days 1-14									
Price change	0.586***	0.595***	0.584***	0.525***	0.591***	0.618***	0.576***	0.614***	0.567***
	(0.043)	(0.064)	(0.077)	(0.072)	(0.067)	(0.064)	(0.056)	(0.070)	(0.057)
Panel A2: Weather Shocks for Days 1-3									
Price change	0.096***	0.111***	0.088***	0.100***	0.099***	0.097***	0.077***	0.102***	0.122***
	(0.015)	(0.018)	(0.021)	(0.021)	(0.017)	(0.021)	(0.023)	(0.022)	(0.019)
Panel A3: Weather Shocks for Days 4-7									
Price change	0.191***	0.196***	0.189***	0.199***	0.225***	0.172***	0.173***	0.179***	0.259***
	(0.018)	(0.029)	(0.032)	(0.028)	(0.026)	(0.028)	(0.027)	(0.027)	(0.026)
Panel A4: Weather Shocks for Days 8-11									
Price change	0.212***	0.216***	0.223***	0.188***	0.193***	0.241***	0.216***	0.214***	0.156***
	(0.019)	(0.028)	(0.031)	(0.033)	(0.026)	(0.026)	(0.030)	(0.031)	(0.026)
Panel A5: Weather Shocks for Days 12-14									
Price change	0.088***	0.072***	0.085***	0.039*	0.075***	0.108***	0.110***	0.118***	0.030
	(0.015)	(0.021)	(0.021)	(0.021)	(0.017)	(0.021)	(0.022)	(0.023)	(0.018)
Panel A6: Weather Shocks for Days 15-21									
Price change	0.005	-0.002	-0.031	-0.016	-0.029	0.018	0.027	0.051	-0.027
	(0.023)	(0.031)	(0.034)	(0.026)	(0.027)	(0.035)	(0.040)	(0.033)	(0.028)
Observations	49019	6467	6469	6386	5859	6627	5534	6465	5212
Clusters	166	159	159	157	144	163	136	159	128
Panel B: Weather Shocks for Next Two Weeks (Days 1-14) - HDD Contracts									
Price change	0.613***	0.620***	0.609***	0.573***	0.607***	0.659***	0.586***	0.648***	0.573***
	(0.054)	(0.078)	(0.093)	(0.097)	(0.096)	(0.078)	(0.063)	(0.084)	(0.083)
Observations	27448	3694	3655	3493	3128	3734	3167	3734	2843
Clusters	93	91	90	86	77	92	78	92	70
Panel C: Weather Shocks for Next Two Weeks (Days 1-14) - CDD Contracts									
Price change	0.489***	0.494***	0.472***	0.403***	0.566***	0.459***	0.518***	0.467***	0.561***
	(0.043)	(0.076)	(0.071)	(0.059)	(0.082)	(0.067)	(0.099)	(0.074)	(0.077)
Observations	21571	2773	2814	2893	2731	2893	2367	2731	2369
Clusters	73	68	69	71	67	71	58	67	58

This table presents the inverse regression of Figure 2 of the main paper. It regresses the sum of future weather shocks $\sum_{\tau=\tau_0}^{\tau_1} \widehat{\Delta D_{c[d+\tau]}}$ on a given day's price change Δp_{cd} of weather contract c. Panels A1-A6 vary the time period $\tau_0 - \tau_1$ over which weather shocks are aggregated. For example, Panel A1 sums shocks for the next two weeks (days 1-14 into the future), while panels A2-A6 sum shocks over the next 1-3, 4-7, 8-11, 12-14, and 15-21 days, respectively. Panels B and C split the pooled analysis for days 1-14 into heating and cooling degree days, respectively. The first column pools all observations, while the remaining eight columns provide the results by airport. All regressions include contract fixed effects and are clustered by month to allow for common shocks in a given month. Stars indicate significance levels: * 10%, ** 5%, *** 1%.

Table A3: Heterogeneity in Linear Time Trends by Airport

	Coefficient Estimates				Difference to Average			
	(1a)	(1b)	(1c)	(1d)	(2a)	(2b)	(2c)	(2d)
Panel A: CDD June - September								
Trend	2.432***	2.998***	2.286***	2.774***				
	(0.160)	(0.887)	(0.169)	(0.174)				
at ATL	2.686***	4.117***	1.974***	2.396***	0.254	1.119	-0.312	-0.378
	(0.392)	(1.532)	(0.229)	(0.259)	(0.423)	(1.771)	(0.285)	(0.312)
at CVG	2.161***	1.406	2.475***	3.374***	-0.271	-1.593	0.189	0.601*
	(0.277)	(1.547)	(0.222)	(0.287)	(0.320)	(1.784)	(0.279)	(0.336)
at DFW	3.069***	3.637**	2.957***	2.893***	0.637	0.638	0.671	0.119
	(0.404)	(1.659)	(0.386)	(0.396)	(0.435)	(1.881)	(0.422)	(0.432)
at LAS	4.382***	3.823***	2.274***	2.606***	1.950***	0.825	-0.012	-0.168
	(0.539)	(1.300)	(0.350)	(0.337)	(0.563)	(1.573)	(0.388)	(0.379)
at LGA	2.018***	2.324**	1.832***	2.304***	-0.414	-0.675	-0.454*	-0.470*
	(0.284)	(1.138)	(0.200)	(0.213)	(0.326)	(1.443)	(0.262)	(0.275)
at MSP	1.792***	1.695	2.543***	2.947***	-0.640	-1.304	0.257	0.174
	(0.483)	(1.825)	(0.390)	(0.361)	(0.509)	(2.029)	(0.425)	(0.400)
at ORD	1.552***	2.898**	2.771***	3.397***	-0.880***	-0.101	0.485	0.624*
	(0.231)	(1.428)	(0.252)	(0.267)	(0.281)	(1.681)	(0.304)	(0.318)
at SAC	1.156***	4.218***	1.103***	2.132***	-1.276***	1.220	-1.183***	-0.641
	(0.341)	(1.490)	(0.401)	(0.409)	(0.376)	(1.734)	(0.436)	(0.444)
Panel B: HDD November - March								
Trend	-1.000**	-2.081	-1.662***	-1.854***				
	(0.415)	(1.723)	(0.354)	(0.370)				
at ATL	-1.673***	-4.946**	-1.520***	-0.865*	-0.673	-2.865	0.142	0.988*
	(0.502)	(2.036)	(0.413)	(0.465)	(0.651)	(2.667)	(0.544)	(0.595)
at CVG	-0.605	-2.021	-2.018***	-1.744***	0.395	0.060	-0.356	0.110
	(0.689)	(2.676)	(0.498)	(0.543)	(0.804)	(3.183)	(0.611)	(0.657)
at DFW	-2.160***	-1.570	-1.527***	-1.181***	-1.160*	0.511	0.135	0.673
	(0.480)	(1.830)	(0.397)	(0.409)	(0.634)	(2.514)	(0.532)	(0.551)
at LAS	-2.958***	-6.034***	-0.888*	-1.398***	-1.958***	-3.953	0.774	0.456
	(0.444)	(1.765)	(0.501)	(0.492)	(0.607)	(2.467)	(0.613)	(0.616)
at LGA	-0.573	-1.362	-1.789***	-2.283***	0.428	0.719	-0.126	-0.429
	(0.405)	(2.321)	(0.319)	(0.372)	(0.580)	(2.891)	(0.476)	(0.525)
at MSP	-0.804	1.688	-2.011**	-2.939***	0.196	3.769	-0.349	-1.085
	(1.031)	(3.558)	(0.903)	(0.763)	(1.111)	(3.954)	(0.970)	(0.848)
at ORD	0.670	0.832	-2.148***	-2.565***	1.670*	2.913	-0.486	-0.711
	(0.749)	(2.644)	(0.518)	(0.525)	(0.856)	(3.156)	(0.627)	(0.642)
at SAC	-0.433	-4.717**	-0.707	-2.107***	0.567	-2.635	0.955	-0.253
	(0.396)	(2.099)	(0.556)	(0.543)	(0.574)	(2.715)	(0.659)	(0.657)
Data	Futures	Station	RCP4.5	RCP 8.5	Futures	Station	RCP4.5	RCP 8.5

This table extends Table 1 of the main paper by allowing the estimated annual linear change in degree days $\widehat{\beta_a}$ to vary by airport: $D_{amy} = \alpha_{am} + \beta_a y + \epsilon_{amy}$. The first row in each panel replicates the pooled results from Table 1, while the next eight rows estimate a separate trend by airport. Columns (2a)-(2d) present the difference between the trend at an airport and the pooled trend as well as whether the difference is statistically significant. Stars indicate significance levels: * 10%, ** 5%, *** 1%.

Table A4: Heterogeneity in Linear Time Trends by Month

	Coefficient Estimates				Difference to Average			
	(1a)	(1b)	(1c)	(1d)	(2a)	(2b)	(2c)	(2d)
Panel A: CDD June - September								
Trend	2.432***	2.998***	2.286***	2.774***				
	(0.160)	(0.885)	(0.169)	(0.173)				
in Jun	2.956***	4.495***	1.705***	2.026***	0.524*	1.497	-0.581*	-0.747**
	(0.232)	(1.547)	(0.271)	(0.257)	(0.282)	(1.782)	(0.319)	(0.310)
in Jul	2.281***	2.002	2.633***	3.131***	-0.151	-0.996	0.347	0.358
	(0.315)	(1.704)	(0.311)	(0.307)	(0.353)	(1.920)	(0.354)	(0.352)
in Aug	2.301***	1.542	2.418***	3.222***	-0.131	-1.456	0.133	0.448
	(0.323)	(1.980)	(0.288)	(0.274)	(0.360)	(2.168)	(0.334)	(0.325)
in Sep	2.308***	4.365***	2.250***	2.529***	-0.124	1.366	-0.036	-0.244
	(0.323)	(1.567)	(0.382)	(0.361)	(0.361)	(1.800)	(0.418)	(0.400)
Panel B: HDD November - March								
Trend	-1.000**	-2.081	-1.662***	-1.854***				
	(0.414)	(1.721)	(0.353)	(0.370)				
in Nov	-2.596***	3.749	-1.085	-1.313*	-1.596**	5.830	0.577	0.541
	(0.570)	(3.738)	(0.877)	(0.736)	(0.705)	(4.115)	(0.945)	(0.824)
in Dec	-1.859***	-6.441*	-0.759	-1.212**	-0.858	-4.360	0.904	0.642
	(0.711)	(3.310)	(0.650)	(0.577)	(0.823)	(3.731)	(0.740)	(0.685)
in Jan	0.330	-1.982	-2.107***	-2.964***	1.331	0.099	-0.445	-1.110
	(0.948)	(4.071)	(0.725)	(0.704)	(1.035)	(4.420)	(0.806)	(0.796)
in Feb	2.148***	-6.031	-2.204***	-2.267**	3.148***	-3.950	-0.542	-0.413
	(0.780)	(3.692)	(0.731)	(1.081)	(0.883)	(4.074)	(0.812)	(1.143)
in Mar	-3.701***	0.195	-1.968**	-1.213*	-2.700***	2.276	-0.306	0.640
	(0.466)	(3.349)	(0.827)	(0.708)	(0.624)	(3.765)	(0.899)	(0.799)
Data	Futures	Station	RCP4.5	RCP 8.5	Futures	Station	RCP4.5	RCP 8.5

This table extends Table 1 of the main paper by allowing the estimated annual linear change in degree days $\widehat{\beta_m}$ to vary by month: $D_{amy} = \alpha_{am} + \beta_m y + \epsilon_{amy}$. The first row in each panel replicates the pooled results from Table 1, while the next rows estimate a separate trend by month. Columns (2a)-(2d) present the difference between the trend at an airport and the pooled trend as well as whether the difference is statistically significant. Stars indicate significance levels: * 10%, ** 5%, *** 1%.

Table A5: LASSO Regression of Average Temperature on Oceanic Oscillation Indices

	Season	**All Months**				**Summer**				**Winter**					**Season**	**All Months**				**Summer**				**Winter**			
	Lags	**0**	**1**	**2**	**3**	**0**	**1**	**2**	**3**	**0**	**1**	**2**	**3**		**Lags**	**0**	**1**	**2**	**3**	**0**	**1**	**2**	**3**	**0**	**1**	**2**	**3**
ATL	ENSO	X									X			CVG	ENSO			X							X		
	NAO	X				X	X				X	X			NAO	X							X			X	
	PDO	X										X			PDO	X											
	PNA	X	X			X		X							PNA		X					X					
	AO									X					AO				X					X			
DFW	ENSO	X			X							X		LGA	ENSO										X		
	NAO			X		X	X		X		X				NAO	X					X					X	X
	PDO	X				X						X			PDO	X					X						X
	PNA		X												PNA		X										
	AO									X					AO				X					X			
LAS	ENSO			X				X						MSP	ENSO			X		X					X		
	NAO	X													NAO		X						X			X	
	PDO	X	X			X									PDO	X								X			X
	PNA														PNA	X	X					X		X	X		X
	AO														AO	X		X	X					X	X		X
SAC	ENSO				X		X			X	X		X	ORD	ENSO										X		
	NAO	X			X			X					X		NAO	X					X					X	
	PDO	X	X			X		X							PDO	X								X			X
	PNA		X								X				PNA		X					X		X			X
	AO	X								X		X			AO				X					X	X		

This table summarizes LASSO regression (STATA package lassopack) of daily average temperatures T_{ad} on monthly measures $m(d)$ of five $k = 1 \dots 5$ oceanic oscillation indices $o_{ks[m(d)-\tau]}$: ENSO (El Niño - Southern Oscillation), NAO (North Atlantic Oscillation), PDO (Pacific Decadal Oscillation), PNA (Pacific/ North American Teleconnection Pattern), AO (Arctic Oscillation). The effect is allowed to vary for three subsets s: (i) common among all months, (ii) different for summer months (June - September), or (iii) different for winter months (November - March). Each regression also includes up to three lags $\tau = 0 \dots 3$ of the oceanic indices. All regressions include the same restricted cubic spline $f(d)$ to capture seasonality as in Appendix Figure A7 and a linear time trend in the year $y(d)$. The regression equation is $T_{ad} = \alpha_a + \beta_a f(d) + \gamma_a y(d) + \sum_{k=1}^{5} \sum_{s=1}^{3} \sum_{\tau=0}^{3} \delta_{ks\tau} o_{ks[m(d)-\tau]} + \epsilon_{ad}$. Variables that are selected for each airport a based on the Extended Bayesian Information Criteria (EBIC) are marked with an X.