Online Appendix for

A Unifying Approach to Measuring Climate Change Impacts and Adaptation

Antonio M. Bento, Noah Miller,
Mehreen Mookerjee, Edson Severnini*

February 2021

(for reference only; not for publication)

This appendix provides details on the construction of the data, descriptive figures, and the tabular results of robustness tests and explorations of heterogeneity using alternate specifications. In Appendix A, we provide further details on the sources of our data and construction of final variables in Subsection A.1, and relevant background on ozone as a local air pollutant in Subsection A.2. We additionally include maps of both weather and ozone monitoring station locations, illustrative figures of our decomposition of temperature and its relationship with ozone concentration, and tables of summary statistics. Appendix B includes additional discussion of alternate specifications, split between those investigating robustness in Subsection B.1, and those examining heterogeneity in Subsection B.2.

*Bento: University of Southern California and NBER, abento@usc.edu. Miller: University of Southern California, nsmiller@usc.edu. Mookerjee: Zayed University, mehreen.mookerjee@zu.ac.ae. Severnini: Carnegie Mellon University and IZA, edsons@andrew.cmu.edu.

Appendix A. Additional Data Discussion

This appendix section provides further details on the data sets discussed in Section III, as well as auxiliary data sets used in alternative specifications. It then includes relevant Figures and Tables as outlined below.

A.1. Further Details on the Construction of the Data

Weather Data — Meteorological data was obtained from the National Oceanic and Atmospheric Administration's Global Historical Climatology Network database (NOAA, 2014). This data set provides detailed weather measurements at over 20,000 weather stations across the country, for which we use the period April-September, 1950-2013, for the contiguous 48 states. In constructing our complete, unbalanced panel of weather stations we make only one restriction: for each weather station in each year, we include only those stations for which valid measurements of maximum and minimum temperature, as well as precipitation, exist for at least 75 percent of the days in the ozone monitoring season (April-September). Figure A1 illustrates the geographical location of the weather stations that we have used from 1950-2013, while Table A1 reports summary statistics for maximum temperature and our decomposed measures of climate norm and temperature shock, averaged across our entire sample for each year 1980-2013. Figure A2 illustrates the variation we have in both components of the maximum temperature, namely, the temperature shocks and the climate norms, using a semi-balanced panel of the comprehensive set of weather stations[1] while Figure A3 depicts similar variation, but using only the temperature assigned to each ozone monitor in our final sample. Notice that there seems to be more variation in the 30-year MA in the latter figure because it includes cross-sectional variation as well. Also, the 30-year MA trends down towards the end of the period of our study due to changes in ozone monitor location over time, as shown in Figure A5.

These weather stations are typically not located adjacent to the ozone monitors. Hence, we develop an algorithm to obtain a weather observation at each ozone monitor in our sample. Using information on the geographical location of pollution monitors and weather stations, we calculate the distance between each pair of pollution monitor and weather station using the Haversine formula. Then, for every pollution monitor we exclude weather stations that lie

[1]To create this semi-balanced panel, we impose an additional restriction on our complete, unbalanced sample: for each weather station, we include only those stations with valid readings in every year 1950-2013.

beyond a 30 km radius of that monitor. Moreover, for every pollution monitor we use weather information from only the closest two weather stations within the 30 km radius. Once we apply this algorithm, we exclude ozone monitors that do not have any weather stations within 30km. We calculate weather at each ozone monitor location as the weighted average of these two weather stations using the inverse of the squared distance between them. Figure A6 illustrates the proximity of our final sample of ozone monitors to these matched weather stations. We additionally assess the robustness of our results to changes in this algorithm by increasing the radius to 80 km and using the 5 closest weather stations, and by varying the weights used – unweighted arithmetic mean and simple inverse distance weighting – in calculating the approximate daily weather at each ozone monitoring location. The results of our model under these alternative specifications is discussed further in Appendix B.1.

Ozone Data — Ambient ozone concentration data was obtained from the Environmental Protection Agency's Air Quality System (AQS) AirData database, which provides daily readings from the nationwide network of the EPA's air quality monitoring stations. The data was made available by a Freedom of Information Act (FOIA) request. In our preferred specification we use an unbalanced panel of ozone monitors. We make only two restrictions to construct our final sample. First, we include only monitors with valid daily information. According to EPA, daily measurements are valid for regulation purposes only if (i) 8-hour averages are available for at least 75 percent of the possible hours of the day, or (ii) daily maximum 8-hour average concentration is higher than the standard. Second, as a minimum data completeness requirement, for each ozone monitor we include only years for which at least 75 percent of the days in the ozone monitoring season (April-September) are valid; years having concentrations above the standard are included even if they have incomplete data.

We have valid ozone measurements for a total of 5,638,273 monitor-days.[2] The number

[2]Note that this value refers to *all* valid ozone measurements, the final samples used in estimation will be smaller due to, e.g., instances where an ozone monitor is not paired with any weather stations under our matching algorithm. For instance, our main estimating sample contains 5,139,523 valid monitor-day

of total valid monitors increased from 1,361 in the 1980s to 1,851 in the 2000s, indicating a growth of 16.6 percent of the ozone monitoring network per decade.[3] The number of monitored counties in our main estimating sample also grew from 585 in the 1980s to 840 in the 2000s. Figure A5 depicts the evolution of our sample monitors over the three decades in our data, and illustrates the expansion of the network over time. Table A2 provides some summary statistics regarding the increase in the number of monitors over time.[4]

Auxiliary Data — In some of our robustness checks and examinations of heterogeneity we incorporate additional datasets. Sources and any necessary data construction steps are outlined below.

In Table B7 we use measures of whether a county is "VOC-limited" or "NOx-limited." These measures were constructed using data collected by the EPA's network of respective monitoring stations. Note, however, that these are often separate pollution monitors from our main sample of ozone monitors. Additionally, data – especially for VOCs – is relatively sparse compared to ozone data. Due to these data constraints, we construct measures of whether a county is VOC-limited or NOx-limited for each 5-year period in our sample, e.g. 1980-1984, which we then match with our sample of ozone monitors at the county level. To construct these measures we first combine the EPA's VOC and NOx data at the county-day level and generate a daily ratio of VOCs to NOx for each county, where possible. Following the scientific literature, observations with a ratio less than or equal to 4 are coded as VOC-limited, while those greater than 15 are coded NOx-limited, and the remainder are coded as non-limited. We then sum these three measures by county across each 5-year interval and denote a county as VOC-limited, NOx-limited, or non-limited for that interval based on whichever measure was the most prevalent. For example, a county with 50 VOC-limited

observations.

[3]For our main estimating sample, these are 1,285 and 1,701, respectively.

[4]Note that not all monitored counties were monitored in every year, and not all monitoring stations were active in every year. Some monitors were phased in to replace others, while others were simply added to the network over time as needed – thus individual years will generally have less unique monitors and monitored counties than existed across an entire decade or the sample period.

day, 20 NOx-limited days, and 30 non-limited days would be marked as VOC-limited for this 5-year window. Admittedly, this creates a somewhat coarse measure of whether a county is VOC- or NOx-limited. Given the available data, however, this appears to be the furthest this question can be pursued at this time, and, if anything, should be expected to bias the observed effect from this heterogeneity towards zero.

In Table B3 we include average daily windspeed and total daily sunlight as additional regressors within our main specification. These data, although less frequently available, are collected at the same weather monitoring stations as our main temperature and precipitation variables. Due to the sparseness of these data we do not decompose them into a long-run climate component and transitory weather shock as we do with temperature and precipitation.

In Tables 5 and B8 we examine heterogeneity in our results when separating counties into low- median- and high-levels of belief regarding the existence of climate change and the use of regulation to reduce carbon emissions. These measures were constructed using county level survey data collected by Howe et al. (2015) in 2013 which estimate the percentage of each county's respective population that hold such beliefs. Notably, we do not rely on the explicitly stated aggregate level of belief, but rather the relative level of belief compared to the rest of our sample. Specifically, we separate counties into low- median- or high-belief terciles based on their stated level of belief in the existence of climate changes – and separately by their belief in the use of regulations on carbon emissions. In this way we arrive at three equally sized groups for which we are able to examine heterogeneity in climate impacts and adaptive response. For reference, Table A4 provides summary statistics of basic demographic characteristics across these three county groupings using data from the 2006-2010 5-year American Community Survey.

In Table B9 we approach the question of hetergeneous beliefs from a different angle, using county-level voting results from the 2008 general presidential election obtained from MIT's Election Data and Science Lab (2018). We construct a simple indicator variable for whether

Barack Obama or John McCain won the popular vote in that county and denote a county as "Democrat" if the former is true.

A.2. Background Details on Ozone

Background on Ozone — The ozone the U.S. EPA regulates as an air pollutant is mainly produced close to the ground (tropospheric ozone).[5] It results from complex chemical reactions between pollutants directly emitted from vehicles, factories and other industrial sources, fossil fuel combustion, consumer products, evaporation of paints, and many other sources. These highly nonlinear Leontief-like reactions involve volatile organic compounds (VOCs) and oxides of nitrogen (NOx) in the presence of sunlight. In "VOC-limited" locations, the VOC/NOx ratio in the ambient air is low (NOx is plentiful relative to VOC), and NOx tends to inhibit ozone accumulation. In "NOx-limited" locations, the VOC/NOx ratio is high (VOC is plentiful relative to NOx), and NOx tends to generate ozone.

As a photochemical pollutant, ozone is formed only during daylight hours, but is destroyed throughout the day and night. It is formed in greater quantities on hot, sunny, calm days. Indeed, major episodes of high ozone concentrations are associated with slow moving, high pressure systems, which are associated with the sinking of air, and result in warm, generally cloudless skies, with light winds. Light winds minimize the dispersal of pollutants emitted in urban areas, allowing their concentrations to build up. Photochemical activity involving these precursors is enhanced because of higher temperatures and the availability of sunlight. Modeling studies point to temperature as the most important weather variable affecting ozone concentrations.[6]

[5]It is not the stratospheric ozone of the ozone layer, which is high up in the atmosphere, and reduces the amount of ultraviolet light entering the earths atmosphere.

[6]Dawson, Adams and Pandisa (2007), for instance, examine how concentrations of ozone respond to changes in climate over the eastern U.S. The sensitivities of average ozone concentrations to temperature, wind speed, absolute humidity, mixing height, cloud liquid water content and optical depth, cloudy area, precipitation rate, and precipitating area extent were investigated individually. The meteorological factor that had the largest impact on ozone metrics was temperature. Absolute humidity had a smaller but appreciable effect. Responses to changes in wind speed, mixing height, cloud liquid water content, and optical depth were rather small.

Ambient ozone concentrations increase during the day when formation rates exceed destruction rates, and decline at night when formation processes are inactive.[7] Ozone concentrations also vary seasonally. They tend to be highest during the late spring, summer and early fall months.[8] The EPA has established "ozone seasons" for the required monitoring of ambient ozone concentrations for different locations within the U.S.[9] Recently, there is growing concern that the ozone season may prolong with climate change (e.g., Zhang and Wang, 2016).

[7]In urban areas, peak ozone concentrations typically occur in the early afternoon, shortly after solar noon when the suns rays are most intense, but persist into the later afternoon.

[8]In areas where the coastal marine layer (cool, moist air) is prevalent during summer, the peak ozone season tends to be in the early fall.

[9]Appendix Table A3 shows the ozone seasons during which continuous, hourly averaged ozone concentrations must be monitored.

Figure A1: Comprehensive Location of all Weather Monitors

Notes: This figure maps the location of all weather stations across the continental U.S. contained in our complete dataset.

Figure A2: Climate Norms and Shocks (semi-balanced sample)

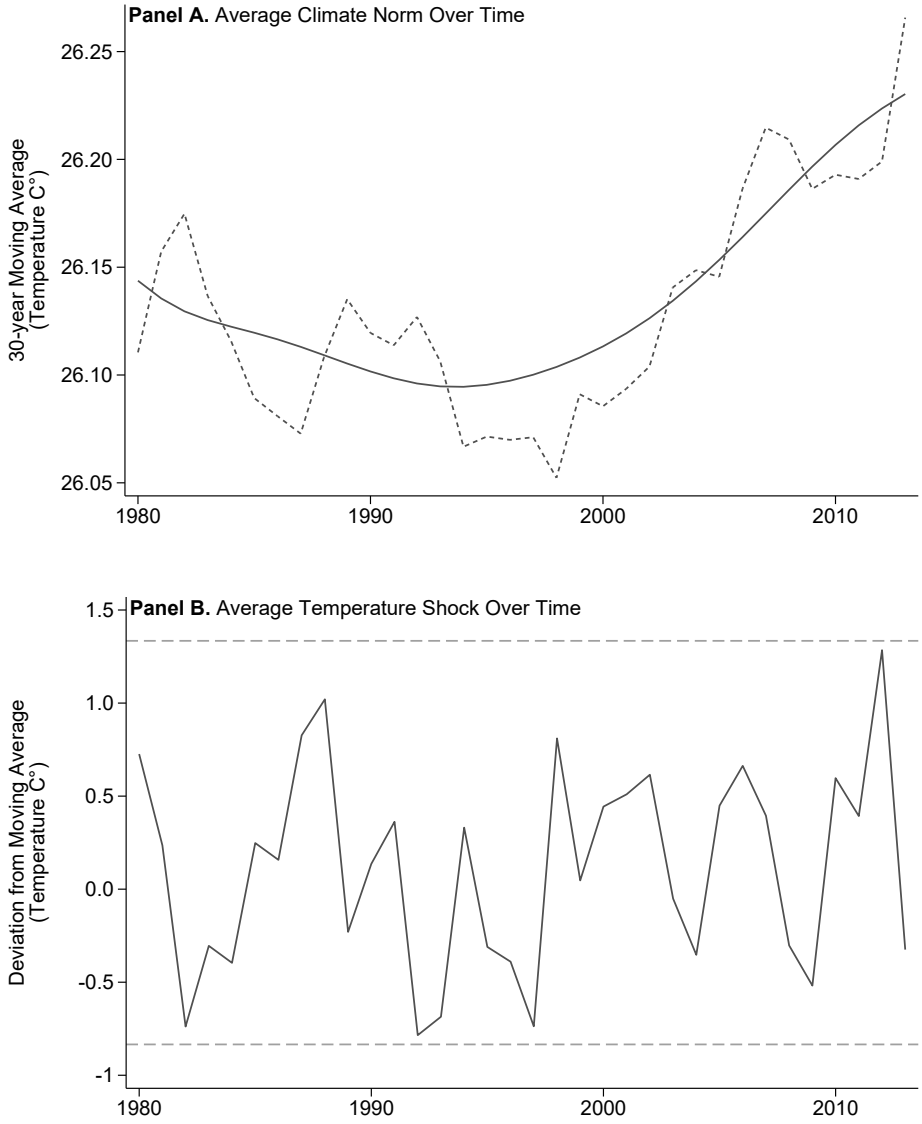

Notes: This figure depicts US temperature over the years in our sample (1980-2013), decomposed into their climate norm and temperature shock components. The climate norm (Panel A) and temperature shocks (Panel B) are constructed from a semi-balanced panel of weather stations across the US from 1950 to 2013, restricting the months over which measurements were gathered to specifically match the ozone season of April–September, the typical ozone season in the US (see Appendix Table A3 for a complete list of ozone seasons by state). Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature. The solid line in Panel A smooths out the annual averages of the 30-year moving averages, and the horizontal dashed lines in Panel B highlights that temperature shocks are bounded in our period of analysis. As described in our data construction, our full sample of weather stations includes only weather stations with valid measurements for at least 75% of the days in the ozone season. Here we further restrict this to a semi-balanced sample, including only those stations with valid readings in every year of our sample.

Figure A3: Climate Norms and Shocks (main model sample)

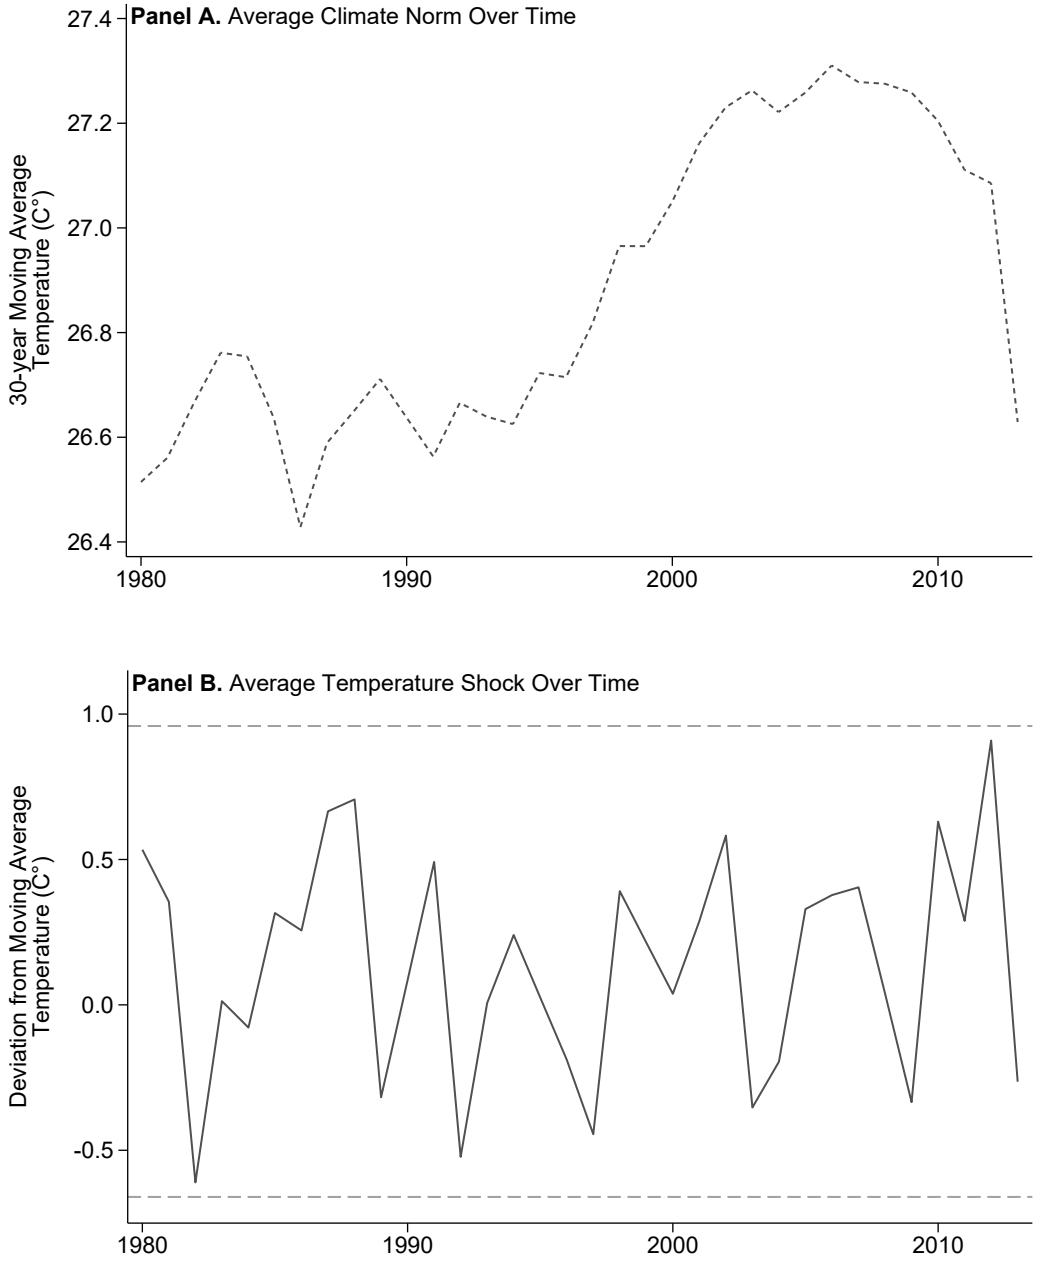

Notes: This figure depicts US temperature over the years in our sample (1980-2013), decomposed into their climate norm and temperature shock components. The climate norm (Panel A) and temperature shocks (Panel B) are constructed from the panel of weather stations included in our main model sample across the US from 1950 to 2013, restricting the months over which measurements were gathered to specifically match the ozone season of April–September, the typical ozone season in the US (see Appendix Table A3 for a complete list of ozone seasons by state). The unbalanced feature of our main sample, with ambient ozone monitors moving north over time (see Figure A4), is the likely driving force behind the downward pattern of the average climate norm at the end of our sample period in Panel A. Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature. The horizontal dashed lines in Panel B highlights that temperature shocks are bounded in our period of analysis.

Figure A4: Evolution of Maximum Ambient Ozone Concentration

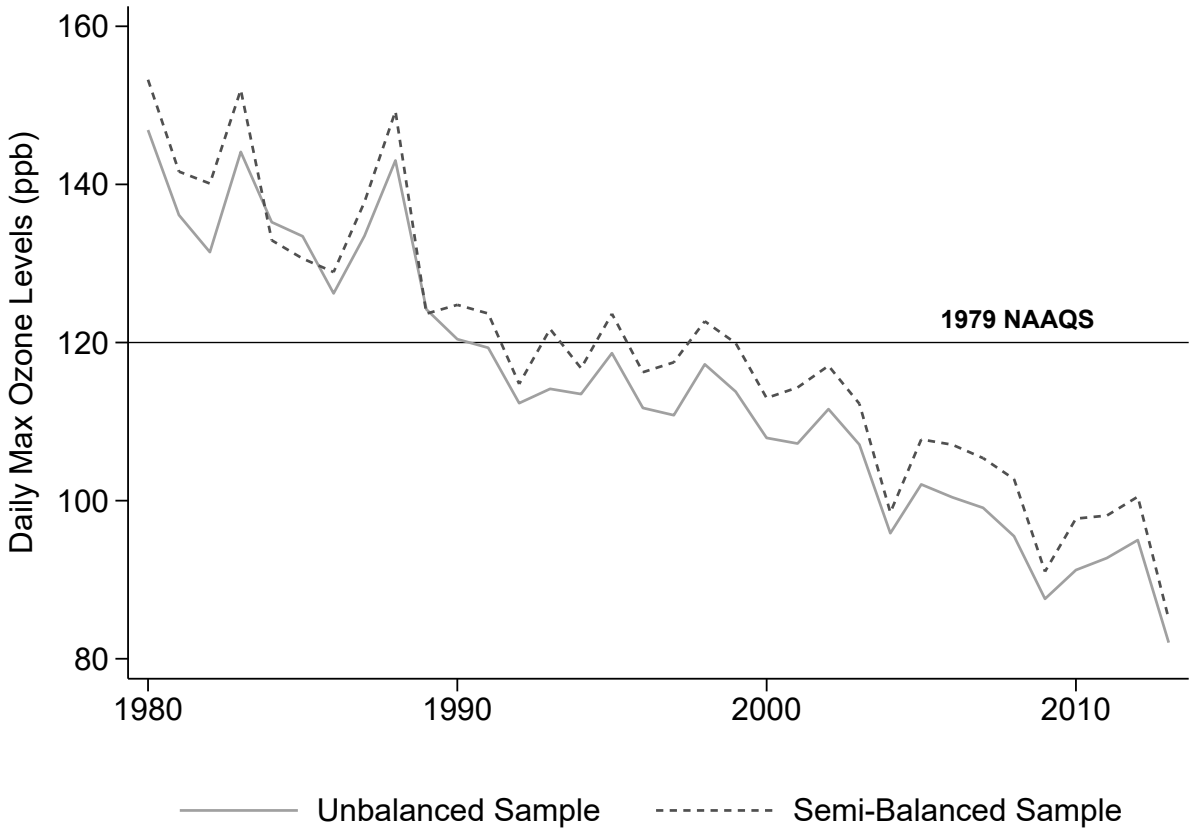

Notes: This figure depicts the evolution of the daily maximum 1-hour ambient ozone concentrations over time in the US for both our complete (unbalanced) sample and our restricted (semi-balanced) sample. For reference the horizontal line depicts the 1979 National Ambient Air Quality Standard for Ozone, which was based on an observed 1-hour maximum ambient ozone concentration of 120 ppb or higher.

Figure A5: Ozone Monitor Location by Decade of First Appearance

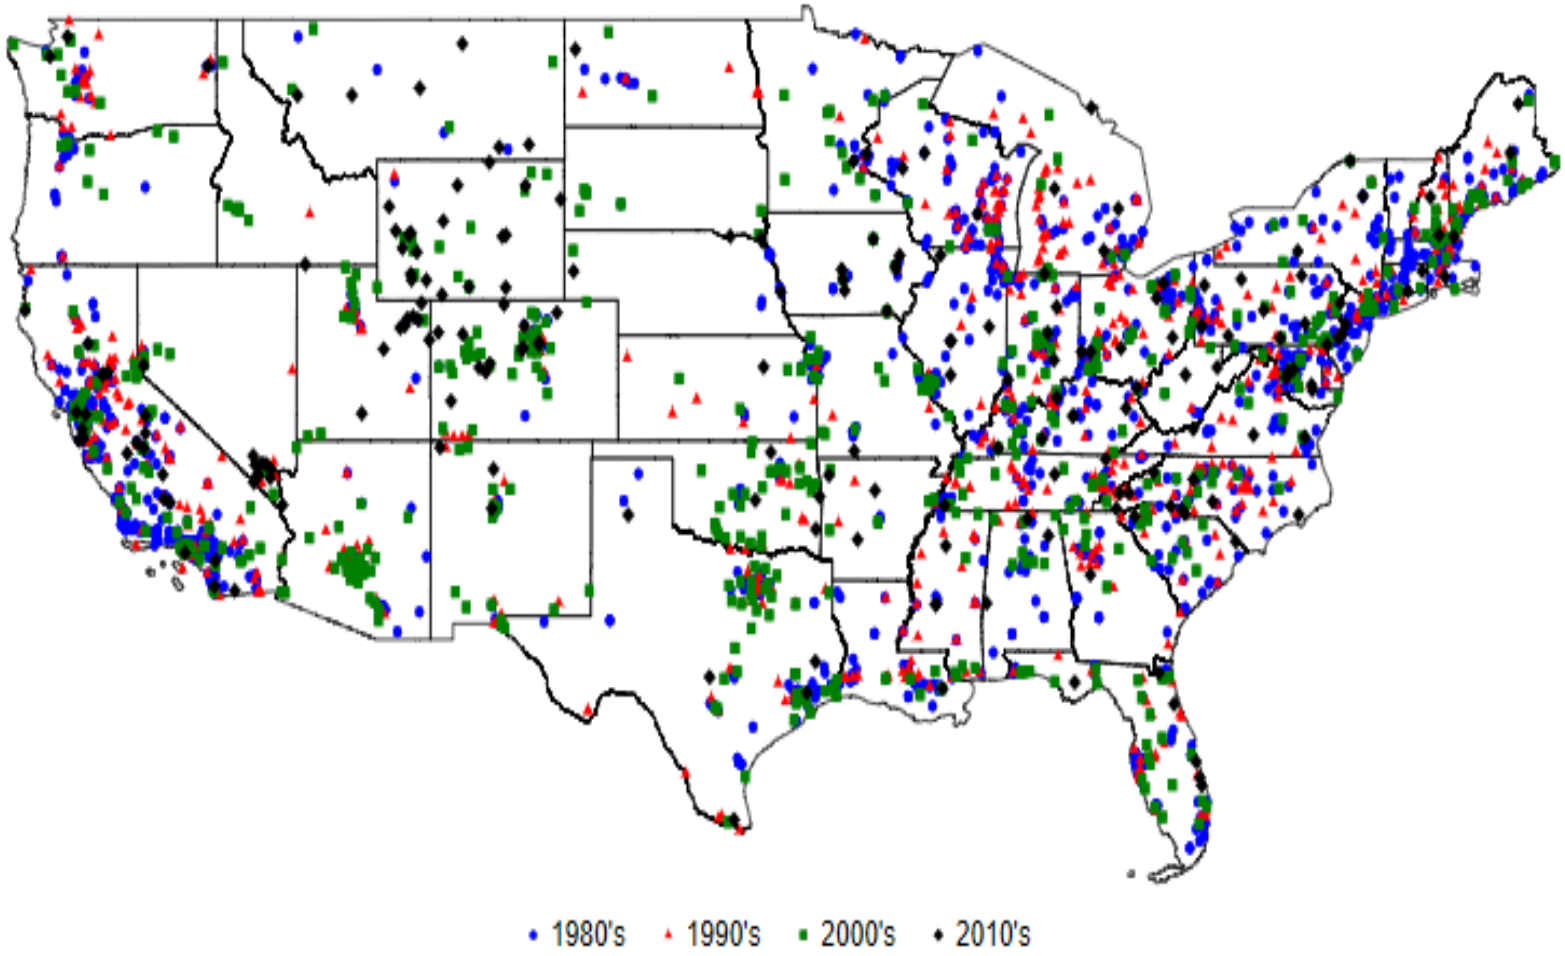

Notes: This figure maps the location of each ozone monitor in our final sample, by decade of first appearance.

Figure A6: Ozone Monitors and their Matched Weather Monitors

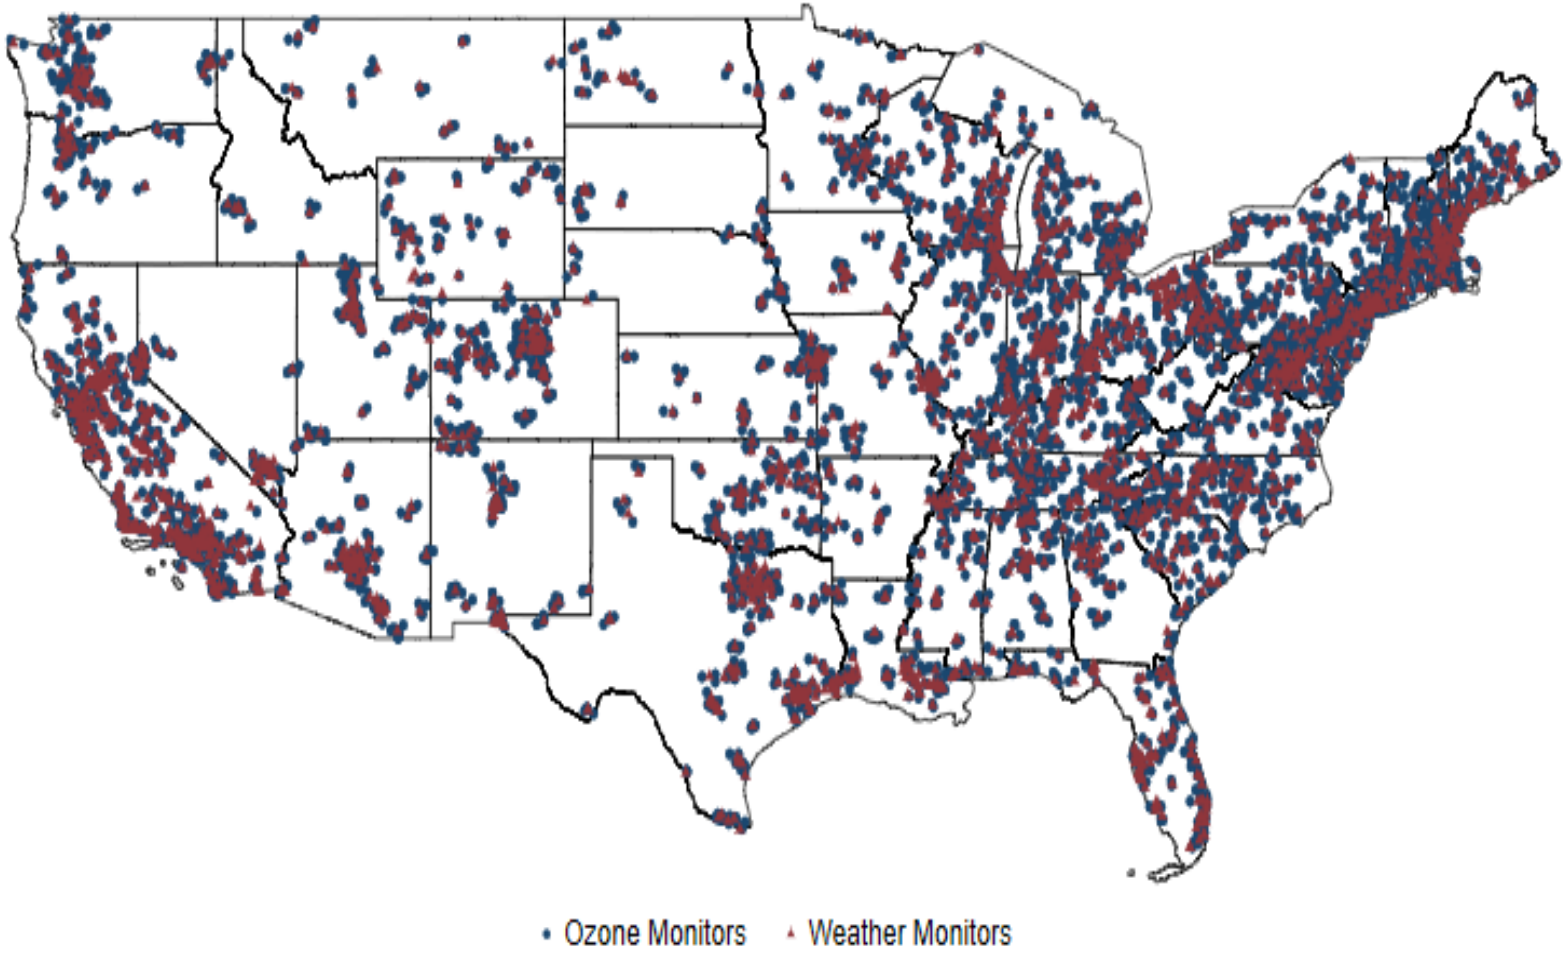

Notes: This figure maps the location of each ozone monitor in our final sample, and their matched weather stations. For each ozone monitor, the closest 2 stations within a 30 km radius have been used in the matching.

Figure A7: Relationship between Ozone and Decomposed Temperature

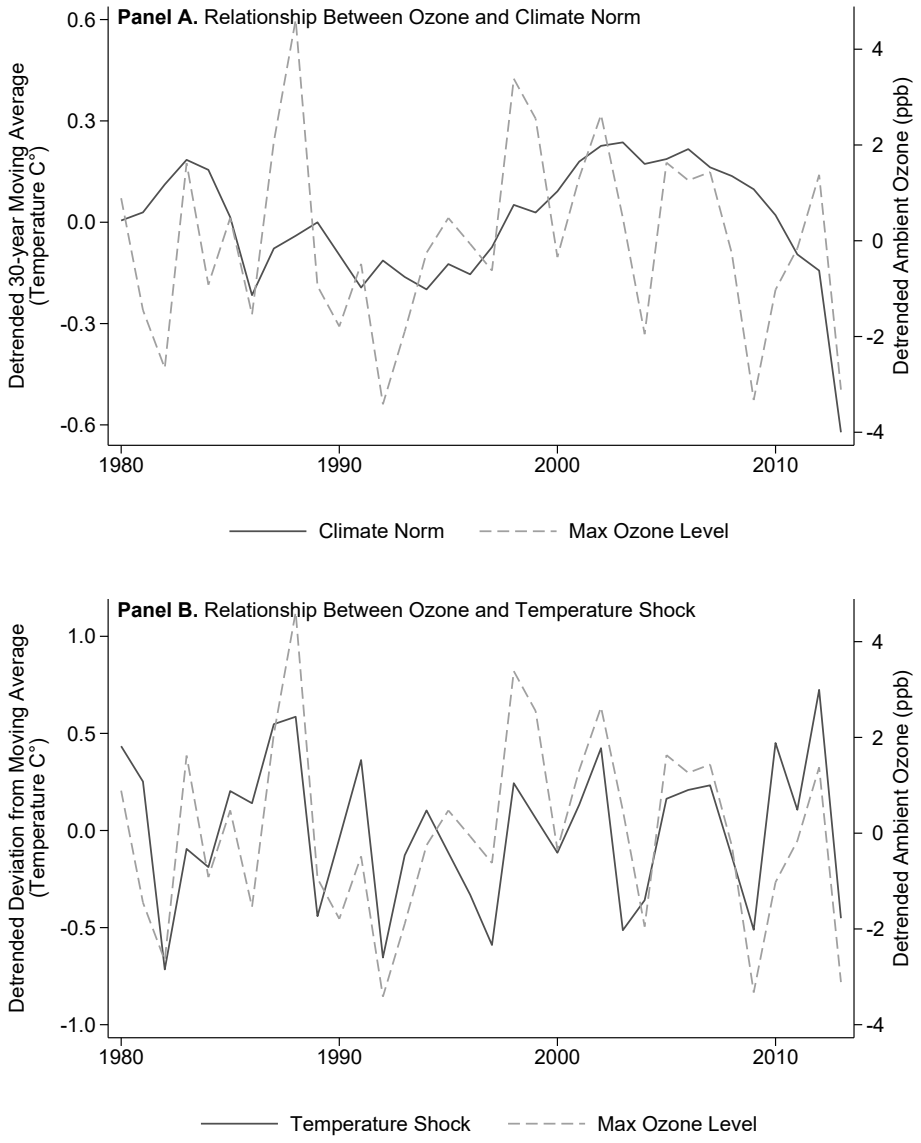

Notes: This figure depicts the general relationship between daily maximum ozone concentrations and temperature over the years in our sample (1980-2013) after decomposing temperature into our measure of climate norm and temperature shock and de-trending the data. Both the climate norm (Panel A) and the temperature shock (Panel B) appear to have a close correlation with ozone concentrations, although the relationship in Panel A appears weaker than that in Panel B, providing suggestive evidence of adaptative behavior. Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature.

Figure A8: Decomposition of Temp. Norms & Shocks – Illustration (Los Angeles, All Years)

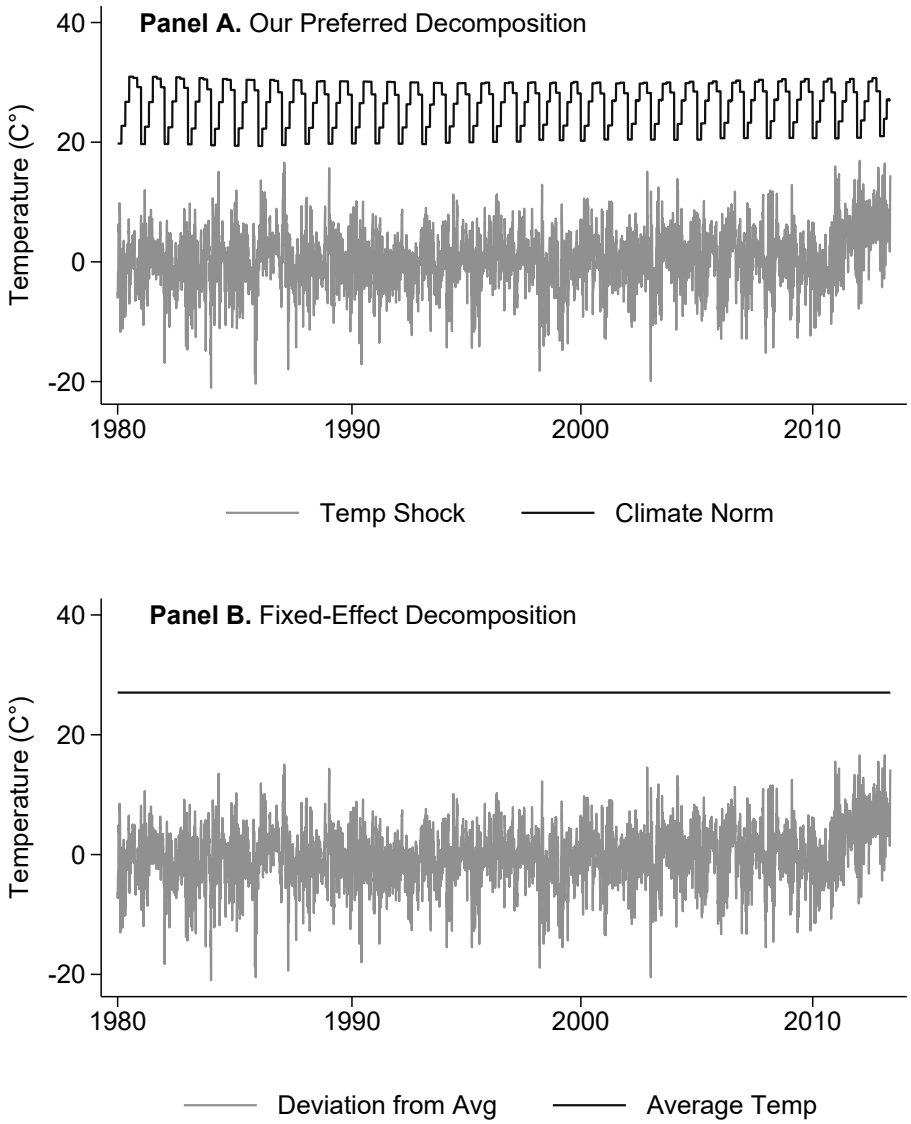

Notes: This figure compares our preferred temperature decomposition method with a standard fixed-effects approach using data for the Los Angeles region duging the ozone season across all years in our sample, illustrating the benefit of our unifying approach as outlined in Equation (6) relative to the standard fixed-effects approach outlined in Equation (2). Specifically, Panel A depicts the daily measure of temperature, as well as its decomposition into climate norm and temperature shock. By contrast, Panel B depicts the same daily measure of temperature, but instead decomposed into a typical fixed-effect average temperature and the deviations from this constant value after additionally controlling for month-year fixed-effects. The black solid line at the top of each panel indicates line represents long-run norms. The gray solid line at the bottom of each panel indicates temperature shocks. Notice that the Temperature Shocks in our preferred decomposition are nearly identical to the deviations in the fixed-effects decomposition, as would be expected from the Frisch-Waugh-Lovell theorem, and illustrate the source of variation used for identifying β_W and β_{FE} respectively. Additionally, Panel A highlights the source of variation in climate used to identify β_C in our proposed approach, while the fixed-effects decomposition lacks any such variation in the measure of climate, as the LA fixed effect is collinear with average temperature. Recall that for our proposed approach the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature.

Table A1: Yearly Summary Statistics for Daily Maximum Temperature

Year	Max Temp	Climate Trend	Temp Shock
(1)	(2)	(3)	(4)
1980	27.0	26.5	0.5
1981	26.9	26.6	0.4
1982	26.1	26.7	-0.6
1983	26.8	26.8	0.0
1984	26.7	26.8	-0.1
1985	27.0	26.6	0.3
1986	26.7	26.4	0.3
1987	27.3	26.6	0.7
1988	27.4	26.6	0.7
1989	26.4	26.7	-0.3
1990	26.7	26.6	0.1
1991	27.1	26.6	0.5
1992	26.1	26.7	-0.5
1993	26.6	26.6	0.0
1994	26.9	26.6	0.2
1995	26.8	26.7	0.0
1996	26.5	26.7	-0.2
1997	26.4	26.8	-0.4
1998	27.3	27.0	0.4
1999	27.2	27.0	0.2
2000	27.1	27.1	0.0
2001	27.4	27.2	0.3
2002	27.8	27.2	0.6
2003	26.9	27.3	-0.4
2004	27.0	27.2	-0.2
2005	27.6	27.3	0.3
2006	27.7	27.3	0.4
2007	27.7	27.3	0.4
2008	27.3	27.3	0.0
2009	26.9	27.3	-0.3
2010	27.8	27.2	0.6
2011	27.4	27.1	0.3
2012	28.0	27.1	0.9
2013	26.4	26.6	-0.3

Notes: This table outlines the evolution of maximum temperature in our sample from the years 19802013 in Column (2). Columns (3) and (4) decompose this into our respective measures of climate norm and temperature shock. Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature.

Table A2: Yearly Summary Statistics for Ozone Monitoring Network

Year	# Observations	# Counties	# Ozone Monitors
(1)	(2)	(3)	(4)
1980	88426	361	609
1981	100459	399	659
1982	102111	402	661
1983	102429	408	653
1984	103828	390	649
1985	105457	388	648
1986	103820	375	634
1987	110366	392	668
1988	113232	409	686
1989	119938	425	725
1990	126535	443	757
1991	132046	466	792
1992	137754	482	821
1993	146023	511	863
1994	149400	520	876
1995	154230	528	902
1996	153019	530	894
1997	160024	550	931
1998	164491	568	960
1999	168901	585	982
2000	172686	592	999
2001	180872	616	1047
2002	186261	630	1071
2003	188462	641	1082
2004	189868	653	1087
2005	187709	649	1082
2006	188298	650	1075
2007	190824	661	1092
2008	190682	660	1099
2009	194184	678	1116
2010	196439	688	1130
2011	199948	716	1159
2012	199723	703	1148
2013	148306	658	1039

Notes: This table outlines the summary statistics of our main data sample. The construction of our main sample follows EPA guidelines by including all monitor-days for which 8-hour averages were recorded for at least 18 hours of the day and monitor-years for which valid monitor-days were recorded for at least 75% of days between April 1st and September 30th.

Table A3: Ozone Monitoring Season by State

State	Start Month - End	State	Start Month - End
Alabama	March - October	Nevada	January - December
Alaska	April - October	New Hampshire	April - September
Arizona	January - December	New Jersey	April - October
Arkansas	March - November	New Mexico	January - December
California	January - December	New York	April - October
Colorado	March - September	North Carolina	April - October
Connecticut	April - September	North Dakota	May - September
Delaware	April - October	Ohio	April - October
D.C.	April - October	Oklahoma	March - November
Florida	March - October	Oregon	May - September
Georgia	March - October	Pennsylvania	April - October
Hawaii	January - December	Puerto Rico	January - December
Idaho	April - October	Rhode Island	April - September
Illinois	April - October	South Carolina	April - October
Indiana	April - September	South Dakota	June - September
Iowa	April - October	Tennessee	March - October
Kansas	April - October	Texas[1]	January - December
Kentucky	March - October	Texas[1]	March - October
Louisiana	January - December	Utah	May - September
Maine	April - September	Vermont	April - September
Maryland	April - October	Virginia	April - October
Massachusetts	April - September	Washington	May - September
Michigan	April - September	West Virginia	April - October
Minnesota	April - October	Wisconsin	April 15 - October 15
Mississippi	March - October	Wyoming	April - October
Missouri	April - October	American Samoa	January - December
Montana	June - September	Guam	January - December
Nebraska	April - October	Virgin Islands	January - December

Notes: This table shows, for each state, the season when ambient ozone concentration is required to be measured and reported to the U.S. EPA.
[1]The ozone season is defined differently in different parts of Texas.
Source: USEPA (2006, p.AX3-3).

Table A4: County Summary Statistics by Belief in Climate Change

	Count	Mean	Std. Dev.	Minimum	Maximum
Panel A. Low Belief Counties					
Population (1000's)	334	80.6	106.9	0.8	837.5
Average Education (Years)	334	12.7	0.6	11.0	14.3
Median Income ($1000/year)	334	48.6	10.4	21.9	83.3
Average Income ($1000/year)	334	61.7	11.3	36.9	111.9
Panel B. Median Belief Counties					
Population (1000's)	335	174.7	297.4	1.9	3,951.0
Average Education (Years)	335	13.2	0.6	11.8	15.1
Median Income ($1000/year)	335	53.8	12.4	26.3	109.8
Average Income ($1000/year)	335	68.2	14.6	39.2	142.2
Panel C. High Belief Counties					
Population (1000's)	336	466.7	780.8	1.3	9,758.3
Average Education (Years)	336	13.6	0.7	11.5	16.1
Median Income ($1000/year)	336	60.5	16.8	30.4	125.7
Average Income ($1000/year)	336	79.6	21.3	41.1	146.0

Notes: This table reports summary statistics of underlying demographics for each of the terciles of counties used in Table 5. Demographic data were obtained from the 2006-2010 5-year American Community Survey, with income reported in 2015 dollars, and average years of education based on a population weighted average of educational attainment status for the county population over 25 years of age.

Appendix B. Further Robustness Checks and Heterogeneity

This appendix provides further elaboration of the alternative specifications used for robustness checks, as discussed in Section IV, and for examining heterogeneity, as discussed in Section V. It then includes relevant Tables as outlined below.

Table B1. Alternative Criteria for Selection of Weather Stations

Table B2. Comparison to Alternative Estimation Methods (Semi-Balanced Panel)

Table B3. Further Robustness Checks

Table B4. Bootstrapped Standard Errors

Table B5. Non-Linear Effects of Temperature

Table B6. Results by Decade

Table B7. Adaptation by VOC- or NOx-limited Atmosphere

Table B8. Adaptation by Belief in Climate Change Regulation

Table B9. Adaptation by Political Leaning

B.1. Further Robustness Checks

Alternative Criteria for Selection of Weather Stations — While our robustness checks presented in Table 2 have addressed potential concerns with the manner in which we construct our regressors by decomposing temperature, a possible additional concern arises from the fact that temperature monitors are not necessarily sited next to ozone monitors. Because of this, we do not have an exact measure of temperature at the same geographic point as our measure of ozone. As discussed in our data section, we define temperature at an ozone monitoring station as the mean of the reported daily maximum temperatures at the two closest weather stations within 30 kilometers, weighted by the inverse squared distance to the ozone monitor. In so doing, we are likely to approximate a good measure of the daily maximum temperature for the local region as a whole, while also maintaining a close geographic boundary around the ozone monitoring station so as not to influence this approximation with temperature readings from a weather station further away that may be subject to a different set of meteorological conditions. It's possible, however, that a less strongly distance weighted mean would provide a more accurate measure of temperature for the overall local region – although likely less accurate at the ozone monitoring station itself – or that the 2-station and 30-kilometer cutoffs are too restrictive. We investigate the effects of lessening the distance weighting in the calculation of expected temperature at the ozone monitoring station, as well as relaxing the constraints on both the number of included weather stations and distance from the ozone monitor in Table B1. Specifically, columns (1) and (2) report results of our main specification when we maintain the 2-station/30-kilometer restriction, but decrease the weighting scheme to either the simple arithmetic mean in column (1), or a non-squared inverse distance weight in column (2). Columns (3) and (4) use the same weighting schemes as in (1) and (2), but now include temperature readings from the 5 closest weather monitoring stations within 80 kilometers. Results in all four columns are relatively stable and consistent with our main specification.

Non-Random Siting of Ozone Monitors — In recent work, Muller and Ruud (2018) argue that the location of pollution monitors is not necessarily random. The U.S. EPA maintains a dense network of pollution monitors in the country for two major reasons: (i) to provide useful data for the analysis of important questions linking pollution to its varied impacts, and (ii) to check and enforce regulations on criteria pollutants. These are conflicting interests: while monitors should be placed in regions having different levels of pollution to provide representative data, they might be placed in areas where pollution levels are the highest to maintain oversight. Not surprisingly, the authors find that most of the monitors tend to be in areas where pollution levels have been high, and compliance with the regulation is a question.

Following those authors' results, we can expect that ozone monitors that have consistently been in our sample across all years may be located in areas having very high pollution levels, thus commanding constant monitoring and regulation by the EPA. To check if this claim is accurate, we re-run our main analysis using a *balanced* sample of ozone monitors. Starting from our original sample, and using only monitors that have been in the data for every year from 1980-2013, we are left with 89 pollution monitors. The results are reported in Table B2. We find that a 1°C temperature shock leads to a rise in ozone concentrations by 2.03ppb, while a 1°C increase in the climate norm leads to a rise of 1.49ppb, implying an adaptation level of 0.54ppb. As expected, the temperature effects obtained from the balanced panel are *larger* than those in our main results, although the level of adaptation remains largely unchanged. The balanced panel leads to an overestimation of the climate penalty. Therefore, our preferred, unbalanced sample of monitors includes areas with different levels of air pollution, and our estimates should be more representative of the entire country.

Further Robustness Checks — In addition to all prior robustness checks, we conduct three final checks in Table B3. First, it may be a concern that our climate norm variable structures the long-run climate normal temperature as the 30-year *monthly* moving average, despite the fact that seasonal – or within-season – shifts in temperature are unlikely to exactly follow

the calendar at a monthly level. We examine the sensitivity of our results to this decision by alternatively constructing this variable as a 30-year *daily* moving average, allowing it to vary arbitrarily within each month. Results of our main specification, substituting daily moving averages for the standard monthly ones, are presented in column (1). Both coefficients of interest are nearly identical to our original findings. Ultimately, we prefer the monthly moving average because it is likely that individuals recall climate patterns by month, not by day of the year, making the interpretation of adaptation more intuitive. Indeed, broadcast meteorologists often talk about how a month has been the coldest or warmest in the past 10, 20, or 30 years, but not how a particular day of the year has deviated from a daily norm.

Second, it may be a concern that our proposed methodology is heavily reliant on high-frequency data in order to successfully decompose temperature into its climatological and meteorological components. While this concern does not pose a threat to identification in our context per se, if valid it would reduce the generalizability of our method to other contexts with less temporally rich data. We examine this concept by aggregating our data to the monitor-month level, taking the arithmetic mean of all variables for each ozone monitor, by month, for each year of our sample and running our preferred specification on this aggregate sample. As the climate norm variable is already identified from variation in monthly moving averages, we would not expect this coefficient to change other than due to the aggregation of our dependent variable and the temperature deviations, which both would otherwise vary daily. It is less clear, however, how this "smoothing" of daily ozone and temperature deviations might affect the coefficient on temperature shocks. Although our sample size is greatly reduced, now consisting of 178,175 observations compared to the previous 5,139,523 we find qualitatively similar results, reported in column (2). As expected, the coefficient on climate norm is nearly identical, while the coefficient on temperature shock is slightly larger in magnitude than in our full sample model, though statistically very similar.

Lastly, although temperature is the primary meteorological factor affecting tropospheric (ambient) ozone concentrations, other factors such as wind speed and sunlight have also

been noted as potential contributors. High wind speed may prevent the build-up of ozone precursors locally, and dilute ozone concentrations. Ultraviolet solar radiation should trigger chemical reactions leading to the formation of more ground-level ozone. To test whether our main estimates are capturing part of the effects of wind speed and sunlight, we control for these variables in an alternative specification using a smaller sample containing those variables. Column (3) presents our main results from estimating Equation (13) plus controls for average daily wind speed (meters/second) and total daily sunlight (minutes). As expected, higher wind speeds lead to lower ozone concentrations, and more sunlight leads to higher concentrations. From column (3), we find that a 1 meter/second increase in average daily wind speed would decrease ozone concentrations by 2.3ppb, whereas a 1 minute increase in daily sunlight leads to 0.02ppb increase in ozone concentrations. Controlling for these two effects, we find that a shock in daily maximum temperature of 1°C leads to a 1.75ppb increase in daily maximum ozone whereas a 1°C increase in the climate norm leads to a 1.13ppb increase in ozone, implying a measure of overall adaptation of 0.62ppb – all statistically similar to our main results. Our primary estimates of the impact of temperature on ozone concentrations, and hence our measures of adaptation, do not seem to rely crucially on other potentially important meteorological factors.

B.2. Heterogeneity

Nonlinear Effects of Temperature — Because ozone formation may be intensified with higher temperatures, we also examine the heterogeneous nonlinear effects of daily maximum temperature on ambient ozone concentrations. Similar to our previous investigations we start by creating indicator variables denoting whether the contemporaneous daily maximum temperature at a given ozone monitor falls within a certain 5°C temperature bin. The lowest bin is below 20°C (just over the 10th percentile of our temperature distribution), and the highest bin is above 35°C (90th percentile of our temperature distribution). Table B5 presents the re-

sults of our preferred specification when interacting each of these temperature bin indicators with our other covariates in column (1). The implied measure of adaptation is presented in column (2). In this way, the *marginal* effect of a 1°C change in either component of temperature is constrained to be constant within its respective 5°C temperature bin, but is allowed to vary across each bin. As expected, we find that higher temperatures increasingly lead to higher ozone concentrations.

Below 20°C, a 1°C temperature shock would raise ozone levels by an additional 0.69ppb, while a similar increase in the climate norm would raise ozone concentrations by 0.14ppb. Above 20°C, however, both effects drastically increase, with temperature shocks increasing ozone by 1.69ppb between 20-25°C, and by over 2ppb above 25°C. Interestingly, while the effect of a 1°C increase in the climate norm is increasing with temperature up to 30°C – at 1.28ppb and 1.83ppb for 20-25 and 25-30°C, respectively – it is actually decreasing with temperature above 30°C – at 1.5ppb and 0.90ppb for 30-35 and above 35°C, respectively. This would imply a more than doubling of our full-sample measure of adaptation above 35°C, at 1.15ppb, suggesting that agents may make extra effort to reduce ozone precursor emissions on months where they've experienced average temperatures above 35°C in the past, and thus might reasonably expect ozone to be a particular issue during those months.

This relatively high level of adaptation above 35°C can be plausibly explained by at least two reasons. First, regions having temperatures above 35°C might have higher incidence of sunlight which might lead to more extensive use of solar panels to generate electricity. Since the U.S. as a whole is predominantly NOx limited, we would expect that changes in electricity usage drastically affect ozone concentrations.[10] Higher temperatures might be creating an environment that is more suited to shifts away from conventional and dirtier sources of power generation, thus leading to higher levels of adaptation. Second, absent any adaptation, days that are exceptionally hot are more likely to cause exceptionally high levels

[10]Electricity generation is a major source of NOx, and, since ozone formation has a Leontief-like production function in terms of NOx and VOCs, changes in electricity use in a NOx limited region would imply large changes in ozone formation.

of ozone, which could trigger additional regulatory oversight. In order to avoid this, firms would be most likely to concentrate adaptation efforts on days where the "climate normal" temperature is itself the hottest.

Results by Decade — To examine temporal heterogeneity, Table B6 mirrors Figure 5 and reports our results by decade. We split our sample into three "decades" 1980-90, 1991-2001, and 2002-2013 so that we have roughly the same number of years in each. We find that the effects of contemporaneous daily maximum temperature, and its two components of our decomposition, are decreasing over time, as shown in column (1). Nevertheless, looking at column (2), we find evidence that adaptation by economic agents reduced slightly from the 1980's to the 1990's, but stabilized back at its original levels in the 2000's. The average adaptation across all counties in our sample drops from 0.54ppb in the 1980's to 0.43ppb in the 1990's, but increases again to 0.54ppb in the 2000's.

Heterogeneity by Precursor "Limited" Atmospheric Conditions — As detailed in Appendix Section A.2, ozone is formed from precursor pollutants – volatile organic compounds (VOCs) and oxides of nitrogen (NOx) – in the presence of sunlight and heat. Specifically, ozone formation appears to follow a Leontief-like production function, implying that regions where the ambient supply of one of the two precursor pollutants, VOCs or NOx, are limited might be less susceptible to increased ozone formation when faced with increased temperatures.

To examine potential heterogeneity in the temperature/ozone relationship and adaptation along this channel we collected all available data on VOC and NOx emissions for each county in our sample as reported by the EPA. Due to the sparseness of these data, we construct aggregate indicators of whether a county is VOC-limited, NOx-limited, or neither for each 5-year interval of our overall sample.[11]

[11]Because ozone formation follows a Leontief-like production function, a county is "VOC-limited" if the ratio of VOC to NOx is too low, while it would be "NOx-limited" if the ratio is too high, and a middle set of counties would not be limited as they face levels of both precursor emissions closer to the "optimal" mix. Further details on this data can be found in Appendix A.1.

Column (1) of Table B7 presents the results of our main specification when using this restricted sample – approximately 20% of our full sample – finding results that are qualitatively similar, albeit larger in magnitude, to our full sample results for the effects of temperature shock, climate norm, and the resulting measure of adaptation – shown in column (2). The magnitude is likely larger because VOCs may be monitored in places with potentially high concentrations. In column (3) we interact the indicators for VOC- and NOx-limited counties with our other regressors to recover a coarse estimate of the effect that being limited in either precursor has on the relationship between our two measures of temperature and ozone. Both main coefficients, and the resulting measure of adaptation – shown in column (4), remain statistically unchanged for non-limited counties. While the difference from these values is statistically indistinguishable from zero in NOx-limited counties. In VOC-limited counties the effects of temperature shock and climate norm are approximately 31 and 79 percent lower and significant, respectively, although the resulting level of adaptation is not precise enough to conclude that it is statistically different from other counties. This finding appears to corroborate the Leontief-like production function of ozone (e.g., Auffhammer and Kellogg, 2011; Deschenes, Greenstone and Shapiro, 2017); when departing from the balanced mix of ozone precursors, the estimated effects of temperature on ambient ozone concentration decline.

Table B1: Alternative Criteria for Selection of Weather Stations

	Daily Max Ozone Levels (ppb)			
	(1)	(2)	(3)	(4)
Temperature Shock	1.721***	1.700***	1.773***	1.764***
	(0.063)	(0.063)	(0.067)	(0.066)
Climate Norm	1.165***	1.165***	1.156***	1.156***
	(0.051)	(0.051)	(0.050)	(0.050)
Implied Adaptation	0.557***	0.535***	0.617***	0.608***
	(0.041)	(0.042)	(0.044)	(0.043)
Distance Cut-off	30 km	30 km	80 km	80 km
Stations Included	2	2	5	5
Weighting Scheme	Simple Avg	1/Dist	Simple Avg	1/Dist
All Controls	Yes	Yes	Yes	Yes
Observations	5,139,523	5,139,523	5,284,420	5,284,420
R^2	0.484	0.483	0.484	0.485

Notes: This table reports estimates from models using alternative criteria to match weather stations to ozone monitors. These estimates are obtained by our main specification, Equation (13), but using different distance radii, number of weather stations, and weights when matching ozone monitors to weather stations. In our main analysis we use a radius of 30 km, the 2 closest stations, and the inverse squared distance as the weight. In the above columns, we give the same weight to both stations (simple average), or use the inverse distance as an alternative weight. We also increase the radius to 80 km and use the information from the closest 5 weather stations. Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature. The full list of controls are the same as in the main model, depicted in column (1) of Table 1. Standard errors are clustered at the county level. ***, ** and * represent significance at the 1%, 5% and 10%, respectively.

Table B2: Comparison to Alternative Estimation Methods (Semi-Balanced Panel)

	Daily Max Ozone Levels (ppb)		
	Unifying	Fixed-Effects	Cross-Section
	(1)	(2)	(3)
Temperature Shock	2.028***		
	(0.109)		
Climate Norm	1.492***		
	(0.084)		
Max Temperature		2.009***	
		(0.109)	
Average Max Temperature			0.904
			(0.950)
Implied Adaptation	0.536***	1.105	
	(0.082)	(0.773)	
Fixed Effects:			
Monitor-by-Season-by-Year	Yes		
Monitor-by-Month-by-Year		Yes	
State			Yes
Precipitation Controls	Yes	Yes	Yes
Latitude & Longitude			Yes
Non-Attainment Control			Yes
Observations	520,670	520,670	89
R^2	0.475	0.534	0.545

Notes: This table reports our main climate impact results using a semi-balanced panel including only those monitors that exist in every year of our data. Column (1) reports the estimates of our unifying approach, in which we decompose daily maximum temperature into climate norms and weather shocks, and exploit variation in both components in the same estimating equation – our Equation (13). Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year to allow for economic agents to potentially adapt, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature. Column (2) reports the effect of daily maximum temperature on ambient ozone from the panel fixed-effects approach, exploiting day-to-day variation in temperature, hence capturing the effect of a change in weather. Column (3) reports cross-sectional estimates using average maximum temperature and ambient ozone concentrations for each ozone monitor in the sample. Having averaged the variables over all the years from 1980-2013, this estimate captures the effect of a change in climate. Note that while estimates in column (3) must additionally control for whether a county is in violation of the CAA ozone standards, this is implicitly controlled for via the fixed-effects in columns (1) and (2). Standard errors are clustered at the county level in columns (1) and (2), while column (3) uses standard heteroskedastic robust errors. ***, ** and * represent significance at the 1%, 5% and 10%, respectively.

Table B3: Further Robustness Checks

	Daily Max Ozone Levels (ppb)		
	Daily Moving Average	Monthly Aggregation	Meteorological Controls
	(1)	(2)	(3)
Temperature Shock	1.684***	1.806***	1.749***
	(0.064)	(0.062)	(0.078)
Climate Norm	1.207***	1.171***	1.126***
	(0.050)	(0.050)	(0.070)
Average Wind Speed			−2.325***
			(0.309)
Total Daily Sunlight			0.015***
			(0.001)
Implied Adaptation	0.477***	0.636***	0.624***
	(0.040)	(0.041)	(0.064)
All Controls	Yes	Yes	Yes
Observations	5,139,454	178,175	453,829
R^2	0.480	0.859	0.479

Notes: This table reports estimates, obtained by Equation (13), from models that replace our monthly moving average with a daily one in Column (1), aggregate our high-frequency daily data to monthly averages in Column (2), and include additional meteorological controls in Column (3). Specifically, for Column (1) we first decompose contemporaneous maximum temperature into an alternative climate norm, represented by the 30-year *daily* moving average, and the respective temperature shock, represented by the difference between this value and the contemporaneous maximum temperature. We then proceed to estimate our main specification as normal, following Equation (13). For Column (2), we first aggregate our final sample to the monthly level for each ozone monitor before estimating Equation (13) in order to simulate the application of our model to contexts with less granular data. This reduces our sample from 5,139,523 observations to 178,175. Despite this reduction, our results remain qualitatively similar. In Column (3) we augment our main specification by including further meteorological controls, for daily average windspeed and total daily sunlight, in our matrix of additional regressors. While both coefficients are strongly significant, they do not meaningfully affect our coefficients of interest, but drastically restrict our total sample size. Recall that, except for in column (1), the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature. The full list of controls are the same as in the main model, depicted in column (1) of Table 1. Standard errors are clustered at the county level. ***, ** and * represent significance at the 1%, 5% and 10%, respectively.

Table B4: Alternative Clustering and Bootstrapped Standard Errors

	Daily Max Ozone Levels (ppb)
	(1)
Temperature Shock	1.678***
(County Cluster)	(0.063)
(State Cluster)	(0.134)
(Bootstrapped)	(0.063)
Climate Norm	1.164***
(County Cluster)	(0.051)
(State Cluster)	(0.091)
(Bootstrapped)	(0.051)
Implied Adaptation	0.514***
(County Cluster)	(0.041)
(State Cluster)	(0.106)
(Bootstrapped)	(0.042)
All Controls	Yes
R^2	0.481
Observations	5,139,523

Notes: This table compares the standard errors of our main estimates with ones obtained by clustering at the state- rather than county-level, and by bootstrap (block method clustered at the county level, 1000 iterations). The latter addresses the potential concern that because our temperature shocks and norm are constructed, they could be seen as generated regressors. Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature. The full list of controls are the same as in the main model, depicted in column (1) of Table 1. Standard errors are clustered at the county level. ***, ** and * represent significance at the 1%, 5% and 10%, respectively.

Table B5: Non-Linear Effects of Temperature

	Panel A. Below 20°C	
	Daily Max Ozone Levels (ppb)	Adaptation
	(1)	(2)
Temperature Shock	0.691***	
	(0.017)	
Climate Norm	0.142***	0.550***
	(0.034)	(0.030)
	Panel B. 20-25°C	
Temperature Shock	1.694***	
	(0.072)	
Climate Norm	1.278***	0.417***
	(0.069)	(0.031)
	Panel C. 25-30°C	
Temperature Shock	2.017***	
	(0.087)	
Climate Norm	1.826***	0.191***
	(0.092)	(0.041)
	Panel D. 30-35°C	
Temperature Shock	2.196***	
	(0.096)	
Climate Norm	1.496***	0.700***
	(0.128)	(0.070)
	Panel E. Above 35°C	
Temperature Shock	2.049***	
	(0.135)	
Climate Norm	0.901***	1.148***
	(0.180)	(0.136)
All Controls	Yes	
Observations	5,139,523	
R^2	0.494	

Notes: This table reports the average marginal effect of a 1°Celsius increase in the temperature shock and climate norm on the daily maximum ambient ozone concentration (ppb) for days in which the daily maximum temperature fell within different temperature bins. We categorize temperature into 5 bins from < 20°C to > 35°C with 5°C intervals in between. Estimates in column (1) correspond to Equation (13), while estimates in column (2) report the implied measure of adaptation. Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature. The full list of controls are the same as in the main model, depicted in column (1) of Table 1. Standard errors are clustered at the county level. ***, ** and * represent significance at the 1%, 5% and 10%, respectively.

Table B6: Results by Decade

	Panel A. 1980's	
	Daily Max Ozone Levels (ppb)	Adaptation
	(1)	(2)
Temperature Shock	2.264***	
	(0.142)	
Climate Norm	1.726***	0.539***
	(0.086)	(0.088)
	Panel B. 1990's	
Temperature Shock	1.768***	
	(0.051)	
Climate Norm	1.339***	0.428***
	(0.049)	(0.037)
	Panel C. 2000's	
Temperature Shock	1.280***	
	(0.030)	
Climate Norm	0.743***	0.537***
	(0.034)	(0.030)
All Controls	Yes	
Observations	5,139,523	
R^2	0.490	

Notes: This table reports our main estimates disaggregated by the three in our sample: 1980-1990; 1991-2001 and 2002-2013. Estimates in column (1) correspond to Equation (13), while estimates in column (2) report the implied measure of adaptation. Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature. The full list of controls are the same as in the main model, depicted in column (1) of Table 1. Standard errors are clustered at the county level. ***, ** and * represent significance at the 1%, 5% and 10%, respectively.

Table B7: Adaptation by VOC- or NOx-limited Atmosphere

	Daily Max Ozone Levels (ppb)			
	Main Specification		VOC/NOx-Limited	
	Restricted Sample	Adaptation	*Restricted Sample*	Adaptation
	(1)	(2)	(3)	(4)
Temperature Shock	2.135***		2.185***	
	(0.165)		(0.206)	
x VOC-limited			−0.674**	
			(0.281)	
x NOx-limited			−0.082	
			(0.115)	
Climate Norm	1.378***	0.757***	1.347***	0.838***
	(0.140)	(0.119)	(0.139)	(0.151)
x VOC-limited			−1.070***	0.397
			(0.335)	(0.284)
x NOx-limited			0.176	−0.258*
			(0.108)	(0.139)
All Controls	Yes		Yes	
Observations	1,007,563		1,007,563	
R^2	0.505		0.506	

Notes: This table reports estimates of temperature shock and climate norm interacted with an indicator of whether the county was, on average, more VOC-limited, NOx-limited, or non-limited. Using 5-year bins (1980-1984, 1985-1989, etc.) a county is designated as VOC-limited, NOx-limited, or neither for each bin based on whichever of these three categories the county observed the most days of during the 5-year period. We restrict our sample to only those counties for which data on these precursor pollutants is available (approximately 20% of our full sample), and depict the results of our main specification under this restricted sample in column (1) for comparison, with the implied measure of adaption in column (2). In column (3), the main effect reflects the result for non-limited counties, while the interaction term depicts the relative difference in the effect of shocks and trends in pre-cursor limited counties. Column (4) reports the implied measure of adaptation in non-limited counties, and the differential effect in limited ones. Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature. The full list of controls are the same as in the main model, depicted in column (1) of Table 1. Standard errors are clustered at the county level. ***, ** and * represent significance at the 1%, 5% and 10%, respectively.

Table B8: Adaptation by Belief in Climate Change *Regulation*

	Daily Max Ozone Levels (ppb)	Adaptation
	(1)	(2)
Temperature Shock	1.507***	
	(0.046)	
x Low Belief	−0.397***	
	(0.063)	
x High Belief	0.483***	
	(0.118)	
Climate Norm	1.115***	0.392***
	(0.061)	(0.047)
x Low Belief	−0.344***	−0.053
	(0.085)	(0.068)
x High Belief	0.210**	0.273***
	(0.104)	(0.084)
All Controls	Yes	
Observations	5,139,523	
R^2	0.486	

Notes: This table reports estimates of temperature shock and climate norm interacted with an indicator of whether the residents of the county generally believed in the use of regulations on carbon emissions to combat climate change or not. Specifically, all counties in the sample were split into terciles based on the results of a survey conducted on climate change beliefs (Howe et al., 2015). In column (1) the main effect reflects the result for the median tercile of counties, while the interacted effects reflect the difference from this value observed in the lower and higher tercile counties. Column (2) reports the implied measure of adaptation for the median counties along with the differential effect in the low and high belief counties. Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature. The full list of controls are the same as in the main model, depicted in column (1) of Table 1. Standard errors are clustered at the county level. ***, ** and * represent significance at the 1%, 5% and 10%, respectively.

Table B9: Adaptation by Political Leaning

	Daily Max Ozone Levels (ppb)	Adaptation
	(1)	(2)
Temperature Shock	1.325***	
	(0.047)	
x Democrat	0.558***	
	(0.100)	
Climate Norm	0.975***	0.349***
	(0.043)	(0.042)
x Democrat	0.302***	0.256***
	(0.085)	(0.071)
All Controls	Yes	
Observations	5,139,523	
R^2	0.484	

Notes: This table reports estimate of temperature shock and climate norm interacted with an indicator of whether the county voted Democrat in the 2008 presidential election. Column (1) follows Equation (13), with an additional interaction term for Democrat political preference depicting the differential effect of shocks and norms in these counties compared to baseline Republican voting counties. Similarly, column (2) reports the implied measure of adaptation for Republican leaning counties, with the differential effect in Democrat leaning counties noted by the interaction effect. Recall that the climate norm represents the 30-year monthly moving average of the maximum temperature, lagged by one year, while the temperature shock represents the difference between this value and the contemporaneous maximum temperature. The full list of controls are the same as in the main model, depicted in column (1) of Table 1. Standard errors are clustered at the county level. ***, ** and * represent significance at the 1%, 5% and 10%, respectively.

Appendix C. Sources of Variation to Identify Climate Impacts

This appendix provides further elaboration of the two sources of variation used to identify β_C, the coefficient on $\bar{x}_{i\bar{p}}$ in Equation (13) in our empirical application. That estimating equation employs monitor-by-season-by-year fixed effects, a more flexible way to control for a number of unobserved time-varying factors in our ambient ozone setting.

Many least-squares estimators weight heterogeneity with factors that depend on group sizes and the within and between variances of explanatory variables. A univariate regression coefficient, for example, equals an average of coefficients in mutually exclusive (and demeaned) subsamples weighted by size and subsample x-variance (see Goodman-Bacon, 2018, footnote 11):

$$\begin{aligned}
\hat{\beta} &= \frac{\sum_i (y_i - \bar{y})(x_i - \bar{x})}{\sum_i (x_i - \bar{x})^2} \\
&= \frac{\sum_A (y_i - \bar{y})(x_i - \bar{x}) + \sum_B (y_i - \bar{y})(x_i - \bar{x})}{\sum_i (x_i - \bar{x})^2} \\
&= \frac{n_A s^A_{xy} + n_B s^B_{xy}}{s^2_{xx}} \\
&= \frac{n_A s^{2,A}_{xx}}{s^2_{xx}} \hat{\beta}_A + \frac{n_B s^{2,B}_{xx}}{s^2_{xx}} \hat{\beta}_B,
\end{aligned} \tag{C.1}$$

where A and B are mutually-exclusive subsamples, and as usual,

$$\hat{\beta}_j = \frac{s^j_{xy}}{s^{2,j}_{xx}}, \quad j = A, B. \tag{C.2}$$

A simpler version of the estimating equation in our empirical application – Equation (13) – focusing on the time-varying temperature norm, $\bar{x}_{i\bar{p}}$, is:

$$y_{it} = \beta_C \bar{x}_{i\bar{p}} + \phi_{is} + \epsilon_{it}, \tag{C.3}$$

where y represents ambient ozone concentrations, i monitor, t day, and s *season-of-the-sample*, and $\bar{p}$ refers to the aggregation of time – in our case, month – used to construct the climate normals from past observations – the 30-year monthly moving averages of temperature.

Applying a fixed-effects transformation in Equation (C.3), we obtain

$$(y_{it} - \bar{y}_j) = \beta_C (\bar{x}_{i\bar{p}} - \bar{x}_j) + (\epsilon_{it} - \bar{\epsilon}_j), \tag{C.4}$$

where j represents subsamples defined by the trio monitor-by-season-by-year is. Using an

alternative notation, we can express the transformed equation as

$$\tilde{y}_{it} = \beta_C \tilde{x}_{i\bar{p}} + \tilde{\epsilon}_{it}. \tag{C.5}$$

By running OLS on the transformed equation, analogous to the decomposition in Equation (C.1), we can express $\hat{\beta}_C$ as:

$$\hat{\beta}_C = \sum_j \frac{n_j s^{2,j}_{\tilde{x}\tilde{x}}}{s^2_{\tilde{x}\tilde{x}}} \hat{\beta}_{C,j}. \tag{C.6}$$

Equation (C.6) helps us understand why we are leveraging both variation across months within the subsample defined by the monitor-by-season-by-year trio, and variation over time to identify β_C. Indeed, $\hat{\beta}_C$ incorporates variation in temperature norms within each subsample j, $s^{2,j}_{\tilde{x}\tilde{x}}$, and variation in how economic agents located around the same monitor respond to temperature norms in different points in time, which is captured by $\hat{\beta}_{C,j}$.

References

Auffhammer, Maximilian, and Ryan Kellogg. 2011. "Clearing the Air? The Effects of Gasoline Content Regulation on Air Quality." *American Economic Review*, 101(6): 2687–2722.

Dawson, John P., Peter J. Adams, and Spyros N. Pandisa. 2007. "Sensitivity of Ozone to Summertime Climate in the Eastern USA: A Modeling Case Study." *Atmospheric Environment*, 41(7): 1494–1511.

Deschenes, Olivier, Michael Greenstone, and Joseph S. Shapiro. 2017. "Defensive Investments and the Demand for Air Quality: Evidence from the NOx Budget Program." *American Economic Review*, 107(10): 2958–89.

Goodman-Bacon, Andrew. 2018. "Difference-in-Differences with Variation in Treatment Timing." *NBER Working Paper #25018*.

Howe, Peter D., Matto Mildenberger, 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(6): 596–603.

MIT, Election Data & Science Lab. 2018. "County Presidential Election Returns 2000-2016." *https://doi.org/10.7910/DVN/VOQCHQ*, accessed on June 3, 2019.

Muller, Nicholas Z., and Paul A. Ruud. 2018. "What Forces Dictate the Design of Pollution Monitoring Networks?" *Environmental Modeling & Assessment*, 23(1): 1–14.

NOAA, National Oceanic & Atmospheric Administration. 2014. "National Oceanic and Atmospheric Administration (NOAA), Global Historical Climatology Network." *ftp://ftp.ncdc.noaa.gov/pub/data/ghcn/daily/by_year/*, accessed on November 30, 2014.

USEPA, U.S. Environmental Protection Agency. 2006. "Air Quality Criteria for Ozone and Related Photochemical Oxidants - Volume II." Available at epa.gov/ttn/naaqs/aqmguide/collection/cp2/20060228_ord_epa600_r05-004bf_ozone_criteria_document_vol2.pdf, accessed on July 23, 2017.

Zhang, Yuzhong, and Yuhang Wang. 2016. "Climate-driven Ground-level Ozone Extreme in the Fall Over the Southeast United States." *Proceedings of the National Academy of Sciences*, 113(36): 10025–10030.