

Temperature and the Allocation of Time: Implications for Climate Change

Joshua Graff Zivin, *University of California, San Diego, and the National Bureau of Economic Research*

Matthew Neidell, *Columbia University and the National Bureau of Economic Research*

We estimate the impacts of temperature on time allocation by exploiting plausibly exogenous variation in temperature over time within counties. Temperature increases at the higher end of the distribution reduce hours worked in industries with high exposure to climate and reduce time allocated to outdoor leisure for the non-employed, with this time reallocated to indoor leisure. At the lower end of the distribution, time allocated to labor is nonresponsive to temperature increases, but outdoor leisure increases while indoor leisure decreases as temperature warms. We also find suggestive evidence of short-run adaptation to higher temperatures through temporal substitutions and acclimatization.

I. Introduction

High temperatures cause discomfort, fatigue, and even cognitive impairment depending on the composition of one's activities and the degree to which they are exposed to the elements.¹ As a result, weather may play an

Contact the corresponding author, Matthew Neidell, at mn2191@columbia.edu. Data are available as supplementary material online.

¹ See Nielsen et al. (1993), Galloway and Maughan (1997), and Gonzalez-Alonso et al. (1999) for evidence on endurance and fatigue. See Epstein et al. (1980), Ramsey (1995), Pilcher, Nadler, and Busch (2002), and Hancock, Ross, and Szalma (2007) for evidence on cognitive performance.

[*Journal of Labor Economics*, 2014, vol. 32, no. 1]
© 2014 by The University of Chicago. All rights reserved.
0734-306X/2014/3201-0001\$10.00

important role in individuals' decisions regarding the allocation of their time. Higher temperatures can lead to changes in time allocated to work by altering the marginal productivity of labor (or the marginal cost of supplying labor), especially in climate-exposed industries, such as agriculture, construction, and manufacturing. Higher temperatures may also change the marginal utility of leisure activities, altering the distribution of time allocated to nonwork activities. Each of these responses will, in turn, generate indirect impacts through trade-offs between labor and leisure. Since time is a limited but extremely valuable resource, the welfare implications associated with these weather-induced reallocations of time are potentially quite large.

This article is the first to estimate the impacts of daily temperature shocks on the allocation of time to labor as well as leisure activities. The analysis uses individual-level data from the 2003–6 American Time Use Surveys (ATUS) linked to weather data from the National Climatic Data Center. Our econometric models include year-month and county fixed effects, which enables us to identify the effects of daily temperature using the plausibly exogenous variation in temperature over time within counties and within seasons. We flexibly model temperature by including a series of indicator variables for 5-degree temperature bins, with the highest bin for days over 100 degrees Fahrenheit (°F). One of the tremendous advantages of using the ATUS is that we can exploit data from the 2006 heat wave that produced high temperatures across much of the United States to produce more reliable estimates of behavioral responses at the high end of the temperature distribution.² We also employ a variety of strategies to examine compensatory behavior within and across days as well as by historical climate.

Our results reveal a wide range of impacts. While we find suggestive evidence of a moderate decline in aggregate time allocated to labor at high temperatures, further analysis reveals considerable heterogeneity across industry sectors based on their exposure to climatic elements. At daily maximum temperatures above 85°F, workers in industries with high exposure to climate reduce daily time allocated to labor by as much as 1 hour. Almost all of the decrease in time allocated to labor happens at the end of the day when fatigue from prolonged exposure to heat has likely set in. We find limited evidence consistent with adaptation to higher temperatures, recognizing that demand factors may limit workers' discretion in choosing labor supply.

In terms of leisure activities, we generally find an inverted U-shaped relationship with daily maximum temperature for outdoor leisure and a cor-

² The 2006 heat wave covered most of the United States from mid-July through early August. For example, beginning on July 17, every state except Alaska experienced a high temperature over 90°F for 3 consecutive days based on our weather data (described below).

responding U-shaped relationship for indoor leisure. This relationship is most pronounced for those not currently employed, as they have the most flexibility in their scheduling. Overall, these results suggest that protective behavior in response to warmer temperatures may provide an important channel for minimizing the potential health impacts of heat. Temporal substitutions as well as acclimatization also appear to mute the impacts of extreme heat, although these findings are often not significant at conventional levels, in part because of a limited sample size when the data are disaggregated at the high end of the distribution.

While our analysis of responses to acute temperature changes cannot fully reflect responses to the more gradual and systemic changes in temperature predicted under climate change, our results may help to illuminate a heretofore ignored potential channel through which global warming may affect social welfare. It may also help to shed light on the microfoundations for the macroeconomic literature that has focused on climate and economic growth (Sachs and Warner 1997; Nordhaus 2006; Dell, Jones, and Olken 2012).³ The absence of data suitable for identification at climatic scales makes findings based on weather fluctuations an important, albeit imperfect, input for policy-making processes in the face of this uncertainty.⁴

II. Data

A. The American Time Use Survey

The American Time Use Survey (ATUS) is a nationally representative cross-sectional survey available from 2003 to 2006 describing how and where Americans spend their time. Respondents are individuals over age 15 randomly selected from households that have completed their final month in the Current Population Survey (CPS). Each respondent completes a 24-hour time diary for a preassigned date, providing details of the activity undertaken, the length of time engaged in the activity, and where the activity took place. Each respondent is interviewed the day after the diary date and is contacted for 8 consecutive weeks to obtain an interview.

For simplicity, we categorize time allocated throughout the day into three broad activity categories: work, outdoor leisure, and indoor leisure.⁵ To measure time allocated to labor, we sum the total number of minutes

³ Our analysis also provides a potential test of the assumption embedded in most of the Integrated Assessment Models that time allocated to labor is exogenous; these models are used to simulate the economic impacts of climate change and play a prominent role in the design of climate change policies.

⁴ Note that similar empirical strategies have been employed in other studies examining various aspects of climate change (e.g., Schlenker, Hanemann, and Fisher 2005; Deschenes and Greenstone 2007, 2011).

⁵ In our original specification, we also included sleep because it may be affected through changes in the marginal utility of labor or leisure (Biddle and Hamermesh 1990), but since it proved insensitive to temperature, we focus on the allocation of

in which the activity occurred at the respondent's workplace, noting that this could be driven by both demand and supply factors. Categorizing leisure is less straightforward. Despite information in the ATUS on where the activity took place, there is no single comprehensive indicator of indoor versus outdoor activities. For example, a potential response to where an activity took place is "at the home or yard," so we cannot isolate whether individuals were inside or outside. As a result, we use several steps to construct a measure of time spent outdoors, with all remaining activities coded as indoor activities. First, we code outdoor time if the respondent reported the activity was "outdoors, away from home" or the respondent was "traveling by foot or bicycle." Second, we include activities that do not fall into these categories but that, based on the activity code, were unarguably performed outdoors. For example, if a respondent was "at the home or yard" and conducted "exterior maintenance" or "lawn maintenance," we coded this as an outdoor activity. We classify activities that take place in ambiguous locations, such as "socializing, relaxing, and leisure" that occurred at home, as indoors, so our measurement of total time spent outdoors understates actual outdoor time. Given this categorization, nearly all outdoor activities are somewhat physically demanding, while indoor activities are generally of lower intensity.⁶ While imperfect, this split is particularly attractive for our purposes, since the marginal utility of physically active endeavors, especially those outdoors, is expected to be most responsive to changes in temperature.

We define three groups of individuals based on climate exposure and activity choices. Since one of the parameters of interest is the impact of temperature on time allocated to labor, we distinguish between two types of workers based on exposure to climate—those who are generally sheltered from climate (low-risk) and those who are not (high-risk). We separate workers into these risk categories based on National Institute for Occupational Safety and Health (NIOSH) definitions of heat-exposed industries (NIOSH 1986) and industry codes in the ATUS. These include industries where the work is primarily performed outdoors—agriculture, forestry, fishing, and hunting; construction; mining; and transportation and utilities—as well as manufacturing, where facilities are typically not climate-controlled and the production process often generates considerable heat. Individuals from all remaining industries are defined as low-risk. Given potential ambiguities regarding the degree of heat exposure within the manufacturing sector, we also perform sensitivity analyses by classifying these

time over waking hours. Furthermore, for nonwork activities, we do not distinguish between home production and leisure, although for simplicity we refer to it as leisure throughout.

⁶ We can separately identify high-intensity indoor activities that took place at a gym or sports club, but very few activities fell into this category: on average, individuals spend 0.9 minutes per day at a gym.

workers as low-risk, and we find that this makes little difference. The third group consists of those currently unemployed or out of the labor force. This group includes retirees (38%), the unemployed (12%), and students (50%). Those who are on a day off are considered employed, and their work hours are recorded as zero for that day.

To obtain information on the residential location of the individual in order to assign local environmental conditions, we link individuals to the CPS to get their county or MSA (Metropolitan Statistical Area) of residence. County and MSA are only released for individuals from locations with over 100,000 residents to maintain confidentiality, making geographic identifiers available for three-quarters of the sample, though we examine the external validity of this limitation below. Since our weather data are at the county level, we assign individuals with only an MSA reported to the county with the highest population in the MSA. Although spatial variation in weather is unlikely to be substantial within MSAs, we also assessed the sensitivity of this assumption by limiting analyses to individuals with exact county identified, and we found that this had little impact on our estimates.

B. Weather

We obtain historical weather data from the National Climatic Data Center (NCDC) TD 3200/3210 “Surface Summary of the Day” file. This file contains daily weather observations from roughly 8,000 weather stations throughout the United States. The primary data elements we include are daily maximum and minimum temperature, precipitation, snowfall, and relative humidity. Humidity is typically only available from select stations, so we impute humidity from neighboring stations when missing.⁷ Excluding humidity entirely from our regression models had little impact on our results. Furthermore, including county-season fixed effects, which control for average seasonal humidity within an area, also had little impact on our results. The county of each weather station is provided, and we take the mean of weather elements within the county if more than one station is present in the county.⁸

C. Daylight

Daylight is positively correlated with temperature and is likely to influence time allocation, making it a potential confounder. To compute the hours of daylight for every day in each county, we compute daily sunrise

⁷ Unfortunately this limits our ability to explore the joint impacts of heat and humidity, which may also be relevant for affecting time allocation. It is also worth noting that the heat index, which is a nonlinear combination of temperature and humidity, is only valid for temperatures above 80°F and humidity above 40%, so it cannot be calculated for the entire temperature distribution.

⁸ If data from one weather station were missing, we computed means using data from the remaining stations within the county.

and sunset times based on astronomical algorithms (Meus 1991) using the latitude and longitude of the county centroid (obtained from the MABLE/Geocorr2K maintained by the Missouri Census Data Center), adjusting for daylight savings time. The sunrise and sunset results have been verified to be accurate to within a minute for locations between $\pm 72^\circ$ latitude. Since this is an algorithm, we are able to compute these data for every single county and date in our sample.

D. Merged Data

We merge the ATUS and weather data by the county and date, leaving us with a final sample of just over 40,000 individuals with valid weather data. Table 1 presents summary statistics for our final sample. Time allocated to work is just under 3 hours per day, but this includes individuals who report zero hours of work because they are not employed or are interviewed on a day off. Conditional on working, time allocated to labor is 7 hours a day overall, but it is closer to 8 hours a day for high-risk laborers.⁹ In terms of leisure activities, individuals spend just under three-quarters of an hour a day in the defined outdoor activities, recalling that we are likely to understate total outdoor time. Many individuals are identified as spending zero minutes outside; conditional on spending time outside, individuals allocate roughly 2 hours to outdoor leisure. Outdoor leisure is highest for high-risk workers, but it is comparable across the two other groups. Most of the day is spent in indoor activities—nearly 12 hours a day—and nearly everyone spends at least 1 minute a day inside. The nonemployed spend the most time indoors, followed by low-risk workers and then high-risk workers. The remaining 7.5 hours per day is spent sleeping (not shown).¹⁰

Many demographic variables from the CPS are brought forward to the ATUS, providing a large pool of potential covariates for our analysis, which is also shown in table 1. Nearly all demographics are comparable across groups, with one notable exception. The mean age of the nonemployed is 52, compared to 42 and 41 for high- and low-risk workers, respectively. This difference is not surprising given that 38% of the nonemployed in our sample are retired. Nonetheless this difference is important to keep in mind when analyzing responses across groups because while the nonemployed may have more flexibility in their scheduling, they may also be more sensitive to extreme temperatures because of their age (Wagner et al. 1972).

Figure 1 shows the distribution of maximum temperatures from 2003 to 2006 for those county-dates from which we have observations in our

⁹ We also present results from analyses below that explicitly account for the excess zeros.

¹⁰ As previously mentioned, we did not find evidence of a relationship between temperature and sleep.

Table 1
Summary Statistics

	All (<i>N</i> = 42,280)		High Risk (<i>N</i> = 6,246)		Low Risk (<i>N</i> = 21,151)		Non- employed (<i>N</i> = 14,883)	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Time allocation:								
Labor:								
Hours	2.69	3.99	4.33	4.5	4.02	4.26		
Percent hours = 0	.61	.49	.43	.5	.42	.49		
Hours hours > 0	6.94	3.41	7.66	3.2	6.92	3.34		
Outdoor leisure:								
Hours	.73	1.61	.99	2.09	.65	1.48	.73	1.55
Percent hours = 0	.6	.49	.6	.49	.61	.49	.58	.49
Hours hours > 0	1.82	2.12	2.47	2.7	1.67	1.99	1.75	1.98
Indoor leisure hours	11.7	3.81	10.12	3.81	10.69	3.79	13.80	2.82
Covariates:								
Maximum temperature (°F)	67.48	19.04	67.01	18.92	67.31	19.18	67.91	18.88
Minimum temperature (°F)	47.16	17.57	46.52	17.24	47.03	17.73	47.61	17.47
Precipitation (inches/100)	11.17	30.23	11.09	28.95	10.83	29.5	11.70	31.77
Snowfall (inches/10)	.66	5.17	.69	4.88	.64	4.87	.67	5.66
Maximum relative humidity	84.68	14.22	84.9	13.93	84.67	14.2	84.60	14.38
Age	45.29	17.25	42.32	11.47	41.34	13.39	52.16	21.50
Percent over age 65	.16	.37	.03	.16	.04	.20	.38	.49
Male	.43	.5	.74	.44	.41	.49	.34	.47
No. children < age 18	.92	1.16	1.03	1.17	.98	1.13	.79	1.19
Annual earnings (\$1,000)	46.0	61.2	80.3	62.2	68.2	63.3
Diary day a holiday	.02	.13	.02	.14	.02	.13	.02	.13
Employed	.65	.48
Absent from work	.03	.17	.04	.2	.04	.21
Out of labor force	.31	.4688	.33
Employed full-time	.51	.5	.91	.29	.75	.43
White non-Hispanic	.68	.47	.69	.46	.7	.46	.66	.47
High school dropout	.17	.38	.13	.34	.1	.3	.29	.45
High school graduate	.25	.43	.33	.47	.2	.4	.29	.45
Some college	.26	.44	.28	.45	.28	.45	.22	.42
Spouse/partner in household	.55	.5	.68	.47	.57	.49	.47	.50

NOTE.—All statistics are at the daily level. High risk is defined as those employed in agriculture, forestry, fishing, and hunting; mining; construction; manufacturing; and transportation and utilities industries. Low risk consists of remaining industries. Nonemployed is defined as unemployed or out of the labor force.




FIG. 1.—Historical and forecasted temperature ($^{\circ}$ F) distribution. This figure is based on daily observations for each county included in the final ATUS sample for the years indicated. The 2070–99 forecasted temperatures are based on the Hadley 3 climate model under the highest warming scenario (A1).

sample, along with the forecasted distribution for 2070–99 based on the Hadley 3 climate forecast model under the business-as-usual emissions scenario (A1) for the same counties in our final sample.¹¹ The distribution is predicted to shift almost uniformly to the right, suggesting that while summers may become unpleasantly hot, winters may become more pleasantly temperate. At the high end of the distribution, it is worth noting that the number of days that exceed 100° F is expected to rise from roughly 1% of days in the historic period to more than 15% of days in the period 2070–99. Since these days are concentrated in the summer months, it is expected that greater than 50% of summer days will experience temperatures that exceed 100° F. This dramatic shift underscores the importance of exploring the tails of the distribution.¹²

E. Sample Representativeness

A potential concern with the ATUS is nonresponse—not all individuals selected for the ATUS agree to participate, and this may bias our analysis. While others have assessed the degree of nonresponse bias with respect to sociodemographic factors (Abraham, Maitland, and Bianchi 2006), a particularly relevant concern in this context is that temperature may affect

¹¹ These forecasts were a major input into the Intergovernmental Panel on Climate Change’s third assessment report. Daily values were assigned to counties, as described in Schlenker and Roberts (2008).

¹² See Graff Zivin and Neidell (2010) for a discussion of the impact of these forecasted temperature increases on time allocation.

whether an individual participates in the survey. Because the weather data apply to the universe of observations, we can assess whether temperature is related to survey participation by plotting the distribution of temperature for counties in our final sample for both the days time diaries are available and the days time diaries are unavailable. Shown in appendix figure A1, available in the online version of *Journal of Labor Economics*, the distribution of temperature across the two groups is nearly identical, suggesting that nonresponse bias due to temperature is likely to be minimal in our analysis.¹³

III. Econometric Model

A. Baseline Model

To examine the relationship between temperature and time allocation, we estimate the following econometric model:

$$\text{labor}_i = f_1(\text{temp}_{c(i),t(i)}) + \delta_1 Z_{c(i),t(i)} + \gamma_1 X_i + g_1(t(i)) + \alpha_{1c(i)} + \varepsilon_{1i}, \quad (1)$$

$$\text{outdoor}_i = f_2(\text{temp}_{c(i),t(i)}) + \delta_2 Z_{c(i),t(i)} + \gamma_2 X_i + g_2(t(i)) + \alpha_{2c(i)} + \varepsilon_{2i}, \quad (2)$$

$$\text{indoor}_i = f_3(\text{temp}_{c(i),t(i)}) + \delta_3 Z_{c(i),t(i)} + \gamma_3 X_i + g_3(t(i)) + \alpha_{3c(i)} + \varepsilon_{3i}, \quad (3)$$

where the variable labor is the amount of time allocated to labor market activities for individual i , the variable outdoor is the amount of time allocated to outdoor leisure activities, and the variable indoor is the amount of time allocation to indoor leisure activities. We let $t(i)$ represent the date individual i is observed and $c(i)$ represent the county individual in which i resides.

We include $f(\text{temp})$ to allow for a nonlinear relationship between daily maximum temperature and time allocation: increases in temperature may lead to increases in outdoor leisure at colder temperatures, but beyond a certain point they may lead to decreases (Galloway and Maughan 1997). Our model includes separate indicator variables for every 5°F temperature increment (as displayed in fig. 1), which allows differential shifts in activities for each temperature bin.¹⁴ We omit the 76°F–80°F indicator variable, so we interpret our estimates as the change in minutes allocated to that activity at a certain temperature range relative to 76°F–80°F. We focus

¹³ We note that our results may not generalize to less populated areas because we only observe the county of residence for more populated areas.

¹⁴ Models with 2.5°F-size bins for temperature yield strikingly similar results. We also estimated models with higher-order polynomials in temperature. Our results were sensitive to the polynomial degree (results available upon request), thus persuading us to use a more flexible approach.

on maximum temperature, rather than daily average temperature, because most individuals are indoors for a significant period of time for routine activities, such as sleeping, when minimum temperatures often occur.¹⁵ Maximum temperature is also likely to be highly correlated with other relevant temperature measures throughout the day, so it is likely to be a reasonable proxy for individual exposure.¹⁶

In terms of control variables, $Z_{c(i),t(i)}$ are other county-level environmental attributes potentially correlated with temperature (daylight, precipitation, humidity, and minimum temperature).¹⁷ The X_i are individual-level covariates meant to capture preferences for particular activities, listed in table 1. The $g(t(i))$ includes day-of-week dummy variables to account for differences in schedules throughout the week and year-month dummy variables to control for seasonal and annual time trends in activity choice. The $\alpha_{c(i)}$ are county fixed effects that capture all time-invariant observable and unobservable attributes that affect time allocation decisions. Therefore, our parameters of interest that relate temperature to time are identified from daily variations in weather within a county. We demonstrate below that our results are insensitive to numerous robustness checks, supporting the validity of our model.

We estimate equations (1)–(3) simultaneously as a generalized method of moments system of equations in order to constrain the net effect from a temperature change on total time to sum to zero (Wooldridge 2002; StataCorp 2011).¹⁸ This procedure also allows us to address autocorrelation and spatial correlation in temperature by clustering standard errors at the state-month level. We estimate these models for all individuals and then separately for those employed in high-risk industries and those employed in low-risk industries. For those not currently employed, we estimate equations (1)–(3), modifying the constraints accordingly.

B. Exploring Adaptation

The above model allows for little adaptation to changes in temperature, and hence at best it describes a partial picture of short-run behavioral

¹⁵ We also control for minimum temperature to allow for potential recovery from higher temperatures, though excluding minimum temperature entirely had minimal affect on our estimates.

¹⁶ In fact, when we used mean daily temperature in place of maximum, we found comparable results. These results are reported in appendix table A2, available in the online version of *Journal of Labor Economics*.

¹⁷ We control for precipitation with a series of indicator variables for no rain, 0–0.1 inches, 0.1–0.2 inches, . . . , 0.8–0.9 inches, and greater than 0.9 inches.

¹⁸ We include this constraint because sleep is not included, which prevents the dependent variables from summing up to 24 hours for each individual. Given that we found little evidence relating temperature to sleep, this restriction had minimal affect on our estimates (shown below in our robustness checks).

responses to temperature. On hot days, individuals may shift activities to cooler moments within the day (intraday substitution) or postpone them until cooler days arrive (interday substitution). In addition to temporal substitutions, individuals may acclimatize to new temperatures through both physiological changes and behavioral changes by adopting various technologies to cope with unpleasant temperatures.

We estimate several alternative models to explore the scope for adaptation, with the mean of the dependent variables for these alternative models shown in table 2. To assess interday substitution, we include (flexibly modeled) lagged temperature in equations (1)–(3) in addition to contemporaneous temperature, and we also place a comparable constraint on lagged temperatures that the net effect on total time sums to zero. Since people may not be able to substitute across immediately adjacent days, we specify lagged temperature as the maximum temperatures across the previous 6 days.¹⁹ If individuals substitute activities across days, then we expect unpleasant lagged temperatures to increase the demand for current activities.

By aggregating responses within a day, any estimated effects are net of intraday substitutions whereby individuals reschedule activities to more pleasant times of the day. To assess intraday substitution, we split the dependent variables in equations (1)–(3) into time spent during daylight versus twilight hours and estimate separate models for each. To define time allocation during twilight, we include activities that began less than 2 hours after sunrise or less than 2 hours before sunset, where sunrise and sunset values vary over both space and time. Since we are interested in comparing daylight versus twilight responses and the mean level of each variable differs (as shown in table 2), we present these results as the percentage change in time allocation by dividing the change in minutes by the mean of the dependent variable. If unpleasantly warm days are cooler during the evening or the morning, then we expect smaller responses to temperatures during twilight hours as compared to daylight hours.

Physiological acclimatization can occur in short periods of time—up to 2 weeks in healthy individuals under controlled training regimens—but longer for unhealthy individuals or those experiencing passive exposure (Wagner et al. 1972).²⁰ We assess the impacts of short-run acclimatization by estimating separate temperature responses for June and August. Since hot days are a relatively new phenomenon in June but quite common by

¹⁹ Consistent with this, we found more moderate evidence of interday substitution using 1-day lag only. We also estimated the impact of future temperatures to assess whether individuals anticipate changes in temperature, and we obtained estimates consistent with anticipatory behavior but statistically insignificant.

²⁰ Physiological acclimatization arises through numerous channels, including changes in skin blood flow, metabolic rate, oxygen consumption, and core temperatures (Armstrong and Maresh 1991).

Table 2
Various Labor and Outdoor Leisure Measures

	All	High Risk	Low Risk	Nonemployed
Intraday substitution:				
Labor, twilight	1.31	2.25	1.92	
Labor, daylight	1.38	2.08	2.10	
Outdoor leisure, twilight	.24	.36	.22	.23
Outdoor leisure, daylight	.48	.63	.43	.51
Number of observations	42,280	6,246	21,151	14,883
By historical July–August temperature:				
Labor, warm	2.68	4.49	4.14	.08
Outdoor leisure, warm	.71	.98	.61	.74
Labor, cold	2.72	4.20	4.01	.10
Outdoor leisure, cold	.76	1.02	.69	.74
Number of observations	15,058	2,364	7,592	5,102
By month:				
Labor, June	2.69	4.58	3.94	.13
Outdoor leisure, June	1.00	1.24	.91	1.03
Labor, August	2.75	4.25	4.16	.08
Outdoor leisure, August	.94	1.19	.86	.98
Number of observations	3,527	477	1,822	1,228

NOTE.—All numbers represent the mean value of each variable for each group. Daylight is defined as the time from 2 hours after sunrise until 2 hours before sunset. Twilight is defined as before 2 hours after sunrise or after 2 hours before sunset. Warm (cool) is defined as counties in the top (bottom) third of the 1980–89 July–August temperature distribution.

August, a diminished response to high temperatures in August should be viewed as evidence of acclimatization.²¹ Since this test greatly reduces our sample size and power to detect differential effects, we modify the minimum temperature bin to under 65°F, a reasonably innocuous change given the months of our focus.

By including county fixed effects, the econometric model identifies short-run behavioral responses to temperature. Although most physiological acclimatization occurs within a short period of time, behavioral acclimatization may require more time to take effect. To assess longer-run adjustments, we explore the affects of temperature separately for historically warmer and cooler areas. In particular, we compare the response function for people who live in places with the warmest third of average July–August temperatures during the 1980s to those that live in the coldest third.²² The presumption is that those who live in hotter climates have had

²¹ We do not want to conduct this test by comparing the impact of temperature across seasons for at least two reasons. First, this only identifies impacts where there is sufficient temperature overlap across seasons, making it unlikely to identify the impact from very hot weather. Second, marginal utility from pleasant weather may diminish at different rates depending on how often such weather is experienced.

²² The colder places predominantly consist of counties in the Northeast and upper Midwest; warmer places in the Southeast and Southwest; and omitted places

longer periods of time to adapt to warmer conditions, through more complete physiological adaptation as well as investments in technologies that make it easier to cope with high temperatures. If people adapt to changes in climate, people in cooler places would show comparable adaptations as they become warmer, suggesting that the short-run response curve of colder places will eventually become like the short-run response curve of hotter places.

IV. Results

A. Baseline Results

We begin with a focus on the impacts of temperature on time allocation for all individuals and then focus on the impacts for the groups defined in table 1.²³ In figure 2, we find some evidence of a downward trend in time allocated to labor from higher temperatures, shown in the first panel. The estimates, however, are not large in magnitude—the response at daily maximum temperature 100°F is 19 minutes—and are not statistically significant at conventional levels. This suggests that, consistent with recent findings (Connolly 2008), time allocated to labor on net is not responsive to changes in temperature.

Turning to leisure time, we find an asymmetric relationship between daily maximum temperature and outdoor leisure. Time outside at 25°F is 37 minutes less than at 76°F – 80°F , and it steadily climbs until 76°F – 80°F . It remains fairly stable until 100°F and falls after that, though the impact at the highest temperature bin is not statistically significant. While this pattern is consistent with physiological evidence suggesting fatigue from exposure at temperature extremes (Galloway and Maughan 1997), the lack of significance at high temperatures and the high inflection point suggests external factors may play an important role in individual responses.

Indoor leisure shows a highly asymmetric U-shaped pattern. Indoor leisure increases by roughly 30 minutes at 25°F compared to 76°F – 80°F , and it then steadily decreases until 76°F – 80°F . It remains stable until roughly 95°F and then increases considerably after that. At daily maximum temperatures over 100°F , indoor leisure increases by 27 minutes relative to 76°F – 80°F , with this estimate statistically significant at conventional levels.

The analysis in figure 2, however, masks potentially important heterogeneity due to differential occupational exposure to temperature. In figure 3, we focus on time allocations for individuals employed in industries

in the mid-Atlantic, mountain states, and lower Midwest. California was almost evenly split among the three categories. We also perform this analysis for those in the warmest/coldest quartile or quintile, and we found comparable results.

²³ We present all results graphically, but we also provide coefficient estimates for figs. 2–5 in appendix table A1, available in the online version of *Journal of Labor Economics*.

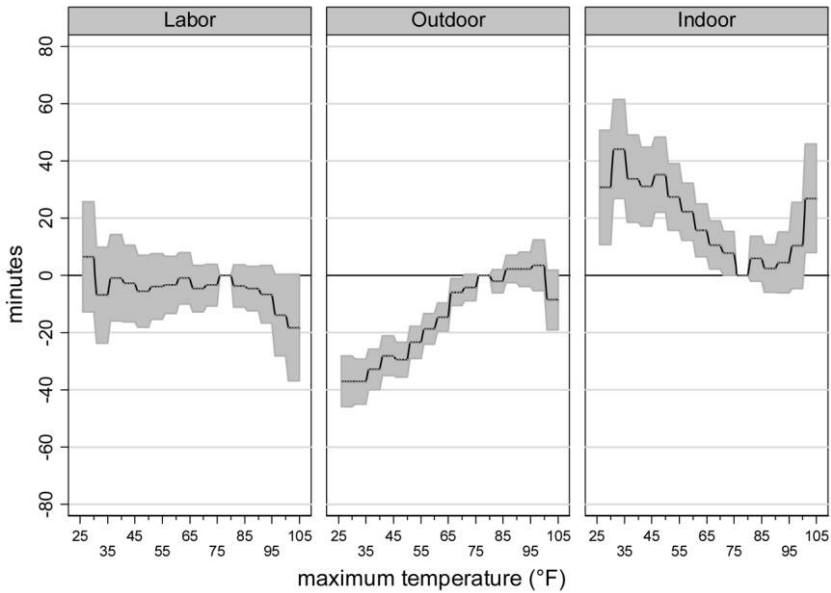


FIG. 2.—Relationship between temperature and time allocation for all individuals. $N = 42,280$ in all regressions. The 95% confidence interval is shaded in gray. Each figure displays the estimated impact of temperature on time allocation based on equations (1)–(3) in the text. Covariates include age, gender, number of children, earnings, employment status, race, education, marital status, family income, day-of-week dummies, minimum temperature, precipitation, humidity, sunrise, sunset, year-month dummies, and county fixed effects.

with a high risk of climate exposure. For time allocated to labor, we continue to find little response to temperatures below 80°F but monotonic declines in time allocated to labor above 85°F. At daily maximum temperatures over 100°F, time allocated to labor drops by a statistically significant 59 minutes as compared to 76°F–80°F. Thus, as hypothesized, the marginal productivity of labor for these workers appears to be significantly affected by temperatures at the high end of the temperature spectrum.

In terms of leisure activities, the results are comparable to the patterns found for all workers, with a slightly higher increase in indoor leisure to accommodate the decrease in time allocated to labor at higher temperatures. At high temperatures, workers appear to substitute their time allocated to labor for indoor leisure, with surprisingly no decline in outdoor leisure. This suggests that, while the marginal utility from outdoor leisure may be declining, the marginal utility of indoor leisure is decreasing at a faster rate over this temperature range.




FIG. 3.—Relationship between temperature and time allocation for high-risk industries. See the legend of figure 2 for the estimating equations and the list of covariates included. $N = 6,246$ in all regressions. The 95% confidence interval is shaded in gray. High-risk industry is defined as agriculture, forestry, fishing, and hunting; mining; construction; manufacturing; and transportation and utilities industries.

In figure 4, we focus on time allocations for those in low-risk industries. For time allocated to labor, we again see little response to colder temperature. While we see a decrease in time allocated to labor at daily maximum temperatures above 95°F, this effect is modest and not statistically significant. The high fraction of workers in these industries explains why we see no net effect on time allocated to labor from higher temperatures. In terms of leisure activities, we see comparable responses as above for colder temperatures but more muted responses at hotter temperatures, which is consistent with the smaller labor response for this group.

In figure 5, we present results for those not employed. Consistent with expectations, we find outdoor and indoor leisure more responsive to temperature changes, particularly at hotter temperatures. Outdoor leisure begins decreasing at lower temperatures when compared to employed individuals, with declines beginning around 90°F. Furthermore, the impacts at higher temperatures are larger and statistically significant. Daily maximum temperatures over 100°F lead to a statistically significant decrease in outdoor leisure of 22 minutes compared to 76°F–80°F. Consistent with Deschênes and Greenstone (2011), such responses at high temperatures are




FIG. 4.—See the legend of figure 2 for the estimating equations and the list of covariates included. $N = 21,151$ in all regressions. 95% confidence interval shaded in gray. Low-risk industry is defined as remaining industries not listed in figure 3.

supportive of short-run adaptation whereby individuals protect themselves from the heat by spending more time inside, which may lessen the health impacts from higher temperatures (Alberini, Mastrangelo, and Pitcher 2008).

B. Robustness Checks

In figure 6, we display results from models that assess the sensitivity of our results to several specification checks, though our results are robust to additional checks not shown. We focus solely on time allocated to labor for high-risk workers and outdoor leisure for the nonemployed because this is where we find the largest and most significant effects, though results are similar for the other activities and groups shown in figures 2–5. We include in this figure the confidence intervals from our baseline results to facilitate interpretation.

Since those employed in the manufacturing industry may in fact work in low-risk industries if the manufacturing plant is climate-controlled, we may have erroneously classified exposure risk for some workers. Our first robustness check shifts individuals from the manufacturing industry into low risk.²⁴ Despite the nearly 50% decrease in sample size in the high-risk

²⁴ This test is irrelevant for the nonemployed group.




FIG. 5.—Relationship between temperature and time allocation for the nonemployed. See the legend of figure 2 for the estimating equations and the list of covariates included. $N = 14,883$ in all regressions. The 95% confidence interval based on standard errors clustered on state-month is shaded in gray. Nonemployed is defined as unemployed or out of labor force.

group, our estimates are largely unaffected by this change. If anything, we find a slightly larger reduction in time allocated to labor at higher temperatures, which is consistent with this misclassification, though the difference is minimal.

In the next two checks, we assess potential omitted-variable bias. First, we exclude all individual-level covariates to assess whether county fixed effects capture sorting into locations based on temperature. Second, we include county-season fixed effects, which allows for seasonal factors specific to each county, such as differences in seasonal activities and humidity (to the extent it is not captured in our imputed humidity variable).²⁵ Figure 6 confirms that these modifications have minimal impact on our estimates, suggesting that confounding is unlikely to be a major concern.²⁶

²⁵ We define seasons as the 3-month periods from December–February, March–June, July–August, and September–November. As an additional approach for controlling for area-time specific shocks, we also ran specifications with state-months fixed effects, and this too had little effect on our estimates.

²⁶ When we include county-season fixed effects, although the estimates for work in high-risk industries are not statistically different from using county fixed effects, according to a Hausman test, the estimates at the highest temperature bins are no

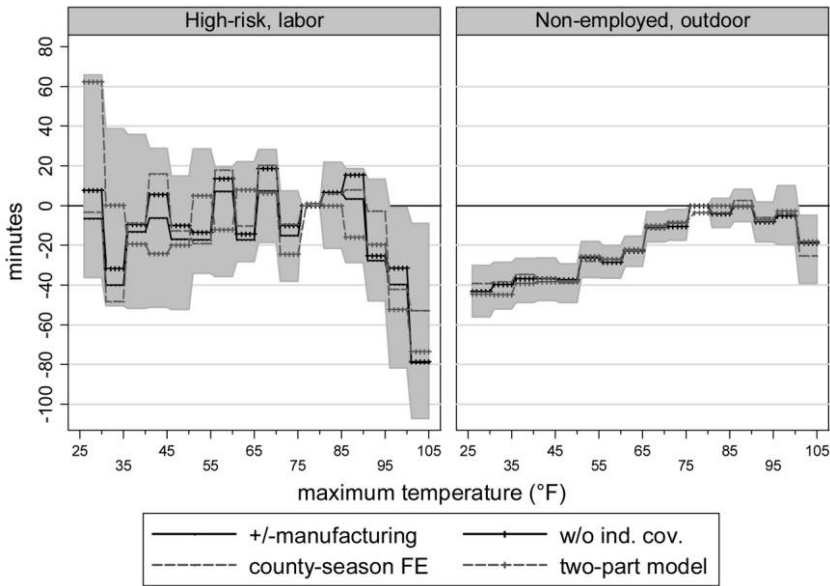


FIG. 6.—Robustness checks. See the legend of figure 2 for the estimating equations and the list of covariates included. The 95% confidence interval for baseline estimates is shaded in gray. “+/- manufacturing” moves those employed in the manufacturing industry from high to low risk; “w/o ind. cov.” excludes all individual level covariates; “county-season FE” includes a county-season fixed effect; “two-part model” presents marginal effects from models that separately estimate the extensive and intensive margins to account for excess zeros and does not constrain coefficients across activities to sum to zero.

As previously mentioned, we have a large mass of observations at zero. Since these zeros represent corner solutions rather than a negative latent value, linear models should produce consistent estimates of the partial effects of interest near its mean value. We further probe this by estimating two-part models, which formally accommodate the mass at zero when it represents a corner solution by separately modeling the extensive and intensive margins.²⁷ In estimating this model, we also relax the constraint that the coefficients across activities sum to zero, so it also tests this restriction.²⁸ Shown in figure 6, the results from the two-part model are

longer statistically significant. Given the loss in precision and the minimal change in point estimates, we use the more efficient county fixed effects as our baseline specification.

²⁷ More specifically, based on laws of probability, $E(y|x) = P(y > 0|x) \times E(y|y > 0, x)$. We estimate $P(y > 0|x)$ using a probit model and $E(y|y > 0, x)$ by OLS, and we compute marginal effects by taking the derivative of $P(y > 0|x) \times E(y|y > 0, x)$.

²⁸ In addition to two-part models, we also estimated Tobit fixed effect models by brute force and semi-parametric censored regression models with fixed effects

quite comparable to the linear estimates. Taken together, the results from figure 6 document a robust relationship between temperature and time allocation.

C. Adaptation

Our static, short-run model may conceal important responses that minimize the impact of temperature shocks. In this section, we probe potential behavioral substitutions and acclimatization as described in the econometric section. As with the robustness checks, we focus solely on time allocated to labor for high-risk workers and outdoor leisure for the nonemployed because this is where we find the largest effects and hence have the largest scope for adaptation. It is important to keep in mind that some of these tests rely on considerably smaller sample sizes, particularly at the upper tail of the distribution, and thus they are underpowered to produce statistical significance at conventional levels. Given the importance of this topic and the inherently limited data availability under current climatic conditions, these results should be viewed as suggestive of the types of adaptation we may see in the future.

We begin by exploring the interday effects of temperature whereby individuals may compensate for unpleasant weather by shifting their activities across days, suggesting that the estimates we have shown thus far may overstate the impacts from warmer temperatures. In figure 7, we present estimates from regressions that include the same indicator variables for lagged temperature (recalling that lagged temperature is defined as the maximum temperature over the previous 6 days) as well as indicators for current temperature. Given that we find a decrease in time allocated to labor for high-risk workers, if interday substitution exists, we expect to see an increase in time allocated to labor from high lagged temperatures. This does not appear to be the case, suggesting little or no role for interday substitution in the workplace. In contrast, outdoor leisure for the nonemployed appears responsive to rescheduling. The two highest temperature bins for lagged temperature are positive, with the estimate of an increase of 15 minutes at 100°F (compared to 76°F – 80°F) statistically significant at the 10% level.

In figure 8, we present results exploring the potential for intraday substitution by estimating whether individuals shift the timing of activities within the day. For time allocated to labor, we find that hours worked during daylight is largely unaffected by warmer temperatures. However, hours worked during twilight is highly responsive to warmer temperatures, and hence this appears to be the driving force behind the labor response found in our base analyses. Furthermore, the difference in responses for temperatures above 85°F is statistically significant at con-

(Honoré 1992), and we found quite comparable results for the marginal effect of temperature on observed time allocation.

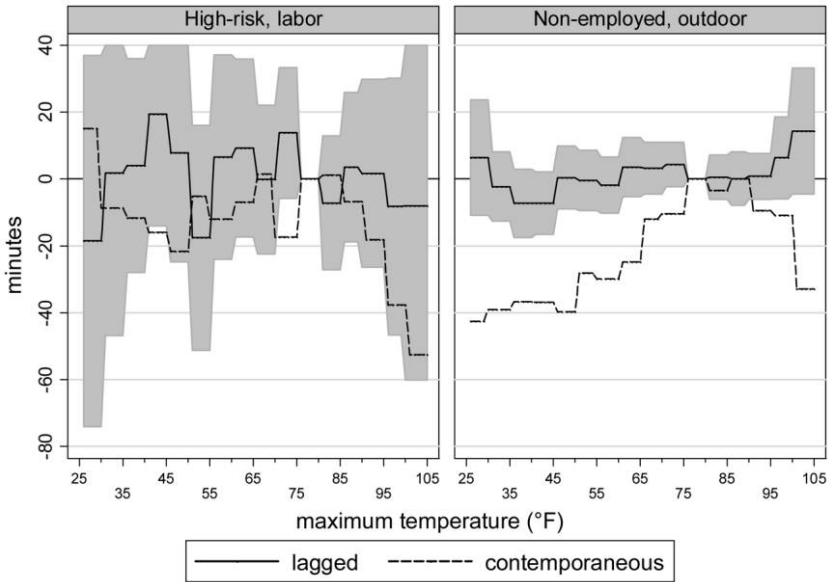


FIG. 7.—Interday substitution. See the legend of figure 2 for the estimating equations and the list of covariates included. Estimates include indicator variables for both contemporaneous temperature and lagged temperature, where lagged temperature is defined as the maximum of the 6 previous days' temperature. The 95% confidence interval for lagged estimates is shaded in gray.

ventional levels.²⁹ If we separate twilight time into the beginning versus end of the day (not shown), we also find that nearly all of the decrease during twilight hours comes from the end of the day. This pattern is consistent with the idea that workers have little discretion over labor supply during core business hours but as fatigue sets in from accumulated exposure to higher temperatures and marginal productivity declines, time allocated to labor becomes responsive.

Turning to outdoor leisure for the nonemployed, we find patterns consistent with individuals shifting activities to more favorable times of the day, though the differences are not statistically significant. For example, we find the turning point for twilight activities occurs at higher temperatures. Furthermore, the drop-off from daily maximum temperatures above 100°F is smaller during twilight hours, representing a 26% decrease as opposed to a 58% decrease during daylight hours (compared to 76°F–80°F).

²⁹ It is also worth noting that, consistent with the notion that these effects are being driven by exposure to high temperature, we find no such twilight effect for those employed in low-risk industries.




FIG. 8.—Intraday substitution. See the legend of figure 2 for the estimating equations and the list of covariates included. Daylight is defined as the time from 2 hours after sunrise until 2 hours before sunset. Twilight is defined as before 2 hours after sunrise or after 2 hours before sunset.

As a test of short-run acclimatization, we explore whether individuals are less sensitive to warmer temperatures as they become more common by estimating the impact of temperatures separately in June versus August. As shown in figure 9, while our estimate for the highest temperature bin is consistent with acclimatization for labor, the overall pattern is less well behaved. For outdoor leisure, we find a pattern highly consistent with short-run acclimatization. Responses in August compared to June are smaller at high temperatures but larger at unseasonably cold temperatures. Given the dramatic drop in sample size, it is unsurprising that these differences are not statistically significant. The differences at high temperatures, however, are large in magnitude. For example, at days with a maximum temperature over 100°F, the nonemployed spend 30 more minutes outside in August than in June.

Our final test for adaptation allows for heterogeneous responses to temperature based on historical climates by grouping counties into those in the highest third of historical July–August temperatures and the coldest third. As shown in figure 10, although we continue to see declines in both time allocated to labor and outdoor leisure at high temperatures in the historically warmer places, the response to high temperatures, particularly for outdoor leisure, is noticeably smaller than the response in colder

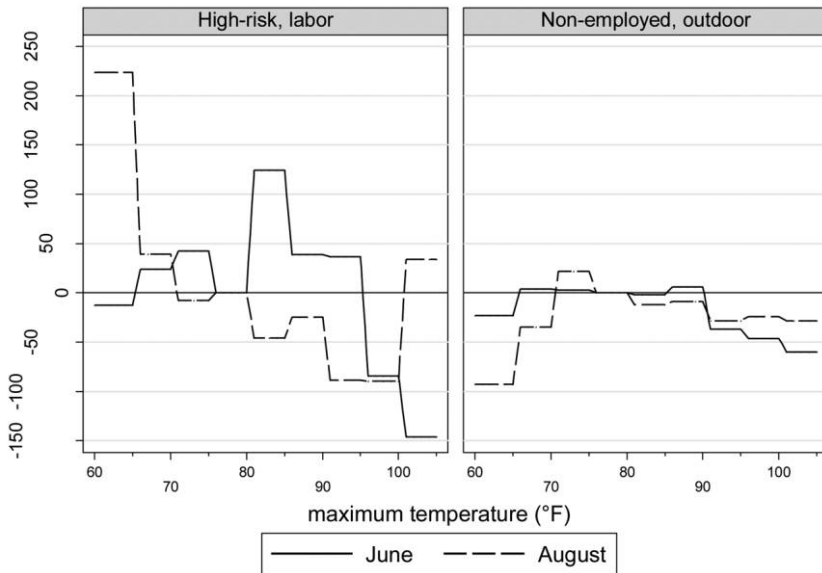


FIG. 9.—Short-run acclimatization. See the legend of figure 2 for the estimating equations and the list of covariates included. Results from this figure are based on regressions stratified by month. $N = 483$ for June estimates and $N = 477$ for August estimates for high risk. $N = 1,173$ for June estimates and $N = 1,228$ for August estimates for the nonemployed.

places. Here again the difference in estimates is not statistically significant but the point estimates are quite large.

V. Conclusion

In this article, we examine the impacts of temperature on individual's allocation of time within the United States. We find large reductions in time allocated to labor in climate-exposed industries as daily maximum temperatures increase beyond 85°F , most of which is reallocated to indoor leisure. Thus, at high temperatures, the marginal productivity of labor in these sectors appears to fall. The near omnipresence of air conditioning in the United States ensures that labor is reallocated to indoor activities since the marginal utility from outdoor leisure is presumed to also fall at temperature extremes.³⁰ For outdoor leisure activities, we generally find an inverted U-shaped relationship with temperature, which is most pronounced for those not currently employed. The absence of any labor-market impacts for

³⁰ According to the 2001 American Housing Survey, 79.5% of all households had some form of air conditioning, with the rate of ownership much higher in warmer regions.




FIG. 10.—Medium-run acclimatization. See the legend of figure 2 for the estimating equations and the list of covariates included. Results from this figure are based on regressions stratified by historical climate. Warm (cool) is defined as counties in the top (bottom) third of the 1980–89 July–August temperature distribution. $N = 2,066$ (2,364) for warm (cool) July–August in high-risk industries. $N = 5,365$ (5,102) for warm (cool) July–August for the nonemployed.

climate-insulated industries underscores the importance of exposure and thus of climate control technologies.

While our evidence from temperature shocks cannot adequately characterize the behavioral responses that could arise under the more gradual and systemic temperature changes expected under climate change, as the only estimates available they provide a unique opportunity to explore the potential implications of a warmer climate on time allocation decisions. If responses to high-frequency variation in temperature are indicative of the sorts of responses we may see under low-frequency variation, our results imply that climate change could lead to a substantial transfer of income from mostly blue-collar sectors of the economy to more white-collar sectors. Moreover, the restructuring of leisure time could have substantial welfare implications; summer reductions in outdoor time would represent a direct utility loss while winter increases could bring gains, both of which could also affect population health through changes in physical activity.

Of course all societies make fixed cost investments in technologies—both physical and social—that balance climatic expectations and adaptability to short-run deviations in weather. As these investments evolve, the scope for both short- and long-run responses to temperature extremes

will likely differ from those found here. As such, all inference with regard to climate change should be undertaken with extreme care. Nonetheless, our results underscore the important role played by environmental factors in shaping labor markets,³¹ as well as the demand for leisure. This represents a fruitful area for future research.

References

- Abraham, Katharine, Aaron Maitland, and Suzanne Bianchi. 2006. Non-response in the American Time Use Survey: Who is missing from the data and how much does it matter? *Public Opinion Quarterly* 70, no. 5: 676–703.
- Alberini, Anna, Erin Mastrangelo, and Hugh Pitcher. 2008. Climate change and human health: Assessing the effectiveness of adaptation to heat waves. Unpublished manuscript, Department of Agricultural and Resource Economics, University of Maryland, College Park.
- Armstrong, Lawrence E., and Carl M. Maresh. 1991. The induction and decay of heat acclimatisation in trained athletes. *Sports Medicine* 12, no. 5:2–12.
- Biddle, Jeff, and Daniel Hamermesh. 1990. Sleep and the allocation of time. *Journal of Political Economy* 98, no. 5:922–43.
- Connolly, Marie. 2008. Here comes the rain again: Weather and the intertemporal substitution of leisure. *Journal of Labor Economics* 26, no. 1: 73–100.
- Dell, Melissa, Benjamin Jones, and Benjamin Olken. 2012. Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics* 4, no. 3:66–95.
- Deschênes, Olivier, and Michael Greenstone. 2007. The economic impacts of climate change: Evidence from agricultural output and random fluctuations in weather. *American Economic Review* 97, no. 1:354–85.
- . 2011. Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the U.S. *American Economic Journal: Applied Economics* 3, no. 4:152–85.
- Epstein Y., G. Keren, J. Moisseiev, O. Gasko, and S. Yachin. 1980. Psychomotor deterioration during exposure to heat. *Aviation, Space, and Environmental Medicine* 51, no. 6:607–10.
- Galloway, S. D., and R. J. Maughan. 1997. Effects of ambient temperature on the capacity to perform prolonged cycle exercise in man. *Medicine and Science in Sports and Exercise* 29, no. 9:1240–49.

³¹ For evidence on the effect of air pollution on worker productivity, see Graff Zivin and Neidell (2012).

- González-Alonso, José, Christina Teller, Signe L. Andersen, Frank B. Jensen, Tino Hyldig, and Bodil Nielsen. 1999. Influence of body temperature on the development of fatigue during prolonged exercise in the heat. *Journal of Applied Physiology* 86, no. 3:1032–39.
- Graff Zivin, Joshua, and Matthew Neidell. 2010. Temperature and the allocation of time: Implications for climate change. NBER Working Paper no. 15717, National Bureau of Economic Research, Cambridge, MA.
- . 2012. The impact of pollution on worker productivity. *American Economic Review* 102, no. 7:3652–73.
- Hancock, P. A., Jennifer M. Ross, and James L. Szalma. 2007. A meta-analysis of performance response under thermal stressors. *Human Factors* 49, no. 5:851–77.
- Honoré, Bo. 1992. Trimmed LAD and least squares estimation of truncated and censored regression models with fixed effects. *Econometrica* 60, no. 3:533–65.
- Meuss, Jean. 1991. *Astronomical algorithms*. Richmond, VA: Willmann-Bell.
- Nielsen, B., J. R. Hales, S. Strange, N. J. Christensen, J. Warberg, and B. Saltin. 1993. Human circulatory and thermoregulatory adaptations with heat acclimation and exercise in a hot, dry environment. *Journal of Physiology* 460 (January): 467–85.
- NIOSH (National Institute for Occupational Safety and Health). 1986. Working in hot environments. NIOSH Publication no. 86–112, National Institute for Occupational Safety and Health, Washington, DC.
- Nordhaus, William. 2006. Geography and macroeconomics: New data and findings. *Proceedings of the National Academy of Science* 103:3510–17.
- Pilcher, June J., Eric Nadler, and Caroline Busch. 2002. Effects of hot and cold temperature exposure on performance: A meta-analytic review. *Ergonomics* 45, no. 10:682–98.
- Ramsey, J. D. 1995. Task performance in heat: A review. *Ergonomics* 38, no. 1:154–65.
- Sachs, Jeffrey D., and Andrew M. Warner. 1997. Sources of slow growth in African economies. *Journal of African Economies* 6, no. 3:335–76.
- Schlenker, Wolfram, W. Michael Hanemann, and Anthony Fisher. 2005. Will U.S. agriculture really benefit from global warming? Accounting for irrigation in the hedonic approach. *American Economic Review* 95, no. 1:395–406.
- Schlenker, Wolfram, and Michael Roberts. 2008. Estimating the impact of climate change on crop yields: The importance of nonlinear temperature effects. NBER Working Paper no. 13799, National Bureau of Economic Research, Cambridge, MA.
- StataCorp. 2011. *Stata 12 base reference manual*. College Station, TX: Stata.

- Wagner, J. S., A. S. Robinson, S. P. Tzankoff, and R. P. Marino. 1972. Heat tolerance and acclimatization in the heat in relation to age. *Journal of Applied Physiology* 33, no. 5:616–22.
- Wooldridge, Jeffrey. 2002. *Econometric analysis of cross section and panel data*. Cambridge, MA: MIT Press.