

Temperature and Human Capital in the Short- and Long-Run¹

Joshua Graff Zivin
UC-San Diego and NBER

Solomon M. Hsiang
UC-Berkeley and NBER

Matthew Neidell
Columbia University and NBER

Abstract: We provide the first estimates of the potential impact of climate change on human capital, focusing on the impacts from both short-run weather and long-run climate. Exploiting the longitudinal structure of the NLSY79 and random fluctuations in weather across interviews, we identify the effect of temperature in models with child-specific fixed effects. We find that short-run changes in temperature lead to statistically significant decreases in cognitive performance on math (but not reading) beyond 26C (78.8F). In contrast, our long-run analysis, which relies upon long-difference and rich cross-sectional models, reveals no statistically significant relationship between climate and human capital. This finding is consistent with the notion that adaptation, particularly compensatory behavior, plays a significant role in limiting the long run impacts from short run weather shocks.

¹We thank Eli Berman, Gordon Dahl, and seminar participants at the University of California, San Diego, the University of Illinois – Urbana Champagne, Lausanne, Verona, Royal Holloway University, the Paris Environmental and Energy Economics Seminar, and PERC for numerous useful comments. We are particularly indebted to Wolfram Schlenker for generously sharing data.

1. Introduction

The threat of climate change and its increasing prominence in public discourse has inspired a significant body of economic research that explores the potential consequences of such change on a variety of outcomes.² Inspired by the neurological literature that documents the brain's sensitivity to temperature (Bowler and Tirri, 1974; Schiff, and Somjen, 1985; and Hocking et al., 2001), we provide the first estimates of the potential impacts of climate change on human capital, focusing on the impacts from both short-run weather and long-run climate. Given the importance of human capital as a principal driver of economic growth (e.g., Nelson and Phelps, 1966; Romer, 1986), these relationships represent an important and unexplored channel through which climate change may impact economic well-being.

Our analysis, which focuses on the same study population over both the short- and long-run, is to our knowledge the first of its kind and serves an important purpose.³ Comparisons across the two models provides a framework through which we can examine the potential offsetting effects from adaptive behaviors, which are expected to play a critical role in determining the ultimate impacts of a gradual changing climate in the coming century (IPCC, 2007; Libecap and Steckel, 2011). As such, our analysis has significant implications for the interpretation of other results in the literature, as most economic studies of climate change impacts rely on identification from short-run weather phenomena.⁴

We begin our analysis by focusing on the relationship between weather and cognitive performance. We use assessments of cognitive ability from the children of the National

²These include impacts on such wide ranging outcomes as agriculture, human health, and economic growth. See, for example, Mendelsohn et al., 1994; Schlenker et al., 2006; Deschenes and Greenstone, 2007; Burke et al., 2009; Schlenker and Roberts, 2009; Feng et al., 2010; Hsiang, 2010; Nordhaus, 2010; Deschenes and Greenstone, 2011; Dell et al., 2012; Graff Zivin and Neidell, 2013; Barreca et al., 2013; Sinha and Cropper, 2013.

³ A similar approach has been taken to examine adaptation in the context of agriculture (Burke and Emerick, 2012), though as we describe below our approach differs due to the dynamic accumulation of impacts in the human capital context.

⁴See Dell et al. (2014) for discussion of the key conceptual challenges in translating results from short-run analyses to the long-run.

Longitudinal Survey of Youth (NLSY79) and merge these data with meteorological conditions at the county level on the day of the assessment. We take advantage of the longitudinal nature of the survey to estimate models with child fixed effects, exploiting the exogenous interview date and daily fluctuations in weather across the same children over time to identify the causal effect of temperature on cognitive performance.

Using a flexible specification for temperature, we find that math performance declines linearly above 21C (70F), with the effect statistically significant beyond 26C (79F).⁵ We do not find a statistically significant relationship with the two assessments of reading performance. The disparity across mathematics and other subjects is consistent with differences in the heat-sensitivity of the regions within the brain on which they rely (Hocking et al., 2001; Kiyatkin, 2007). These differential effects across cognitive tasks also generally supports a neurological rather than behavioral explanation for our results, a finding further bolstered by evidence that child's time to completion of each assessment is not related to temperature.

While the negative impacts from idiosyncratic and short-lived weather-shocks have potentially important implications for the optimal scheduling of cognitively-demanding tasks, the key policy question regarding long run human capital impacts under climate change depends on the impact of a permanent shift in the distribution of weather outcomes. This is important because the slow-moving changes associated with climate change provide greater opportunities for adaptation. As such, the second stage of our analysis exploits two approaches to capture the long-run effects: long-difference fixed effects models that examine the impacts of average weather exposure between tests and cross-sectional regressions with extremely rich controls, including parental and grandparental human capital, to examine the impacts of climate exposure from birth until test-taking.

⁵Note that assessments were only conducted during the spring and summer, so we cannot explore the effects of colder temperatures on performance.

The difference between our approaches for the short- and long-run is important because they capture two distinct adaptation channels that have generally been conflated in the literature.⁶ Ex-ante avoidance behavior, such as technological adoption, mobility, and cultural changes designed to buffer against the effects of climate, may limit exposure to temperature extremes.⁷ Our short-run regressions will net out all such avoidance at least insofar as they have been adopted based on historical climate up until the time of the test. Ex-post compensatory behavior occurs when individuals respond to insults on hot days through subsequent investments that partially or fully offset short run effects, thus minimizing their enduring impact. For example, if a child learns less material during a hot day, parents or teachers may invest additional time or the child may increase her effort in the following days, potentially offsetting the effect of lost learning. This compensating behavior encompasses a wide range of potential responses, and is almost certainly costly, but persistent human capital effects may thus be minimized in the long run. Such ex post behaviors will only be captured by our long-run analysis since responses are predicated on the feedback from earlier tests and thus only depend on weather/climate indirectly.

Despite large predicted long-run effects based on our short-run results, we fail to find evidence that climate is significantly related to human capital accumulation in the long run. In both the long-difference and cross-sectional models, our estimated effects are considerably smaller than projections based on short-run estimates. Furthermore, allowing for a flexible functional form for temperature reveals a flat relationship between temperature and human capital over the entire temperature range. These results highlight the importance of compensatory behavior in our setting and complements existing literature on the human

⁶ This distinction is conceptually similar to that made by Graff Zivin and Neidell (2013) with respect to the health effects from pollution. Individual can engage in avoidance behavior by spending more time indoors or ameliorate the impacts of exposure through the use of medical inputs, such as asthma inhalers.

⁷ See Deschenes and Moretti (2009), Deschenes and Greenstone (2011), and Barreca et al. (2013) for evidence on the impacts of adaptation on the relationship between temperature extremes and mortality. An example of a cultural change that reflects adaptation is differences in school schedules throughout the country: schools in southern states typically end in May, a month before schools in northern states.

capital production function for children, an area of emerging interest in economics (e.g. Currie, 2009; Cunha et al., 2010; Almond and Currie, 2011). Our results also highlight the caution needed when using results from short-run weather shocks to project long-run climate impacts.

This paper unfolds as follows. In section 2, we provide some relevant neurological information on temperatures and brain functioning, and provide a conceptual framework for our econometric models. In section 3, we describe our data in more detail. Section 4 discusses the empirical strategy for the short-run analysis and presents results on the relationship between temperature and test scores. Section 5 describes the empirical strategy for the long-run analysis and presents results on the relationship between climate and human capital. Section 6 offers concluding remarks.

2. Background and Conceptual Framework

In order for climate to affect human capital, we need a plausible mechanism that relates brain function to ambient temperature. A particularly important and likely pathway is through the environment's effect on brain temperature. The brain's chemistry, electrical properties, and function are all temperature sensitive (Bowler and Tirri, 1974; Schiff and Somjen, 1985; Deboer, 1998; Yablonskiy et al., 2000; Hocking et al., 2001), with theory suggesting that the brain's performance as a computational network will be influenced by these parameters (Doya et al., 2007; Moore and Cao, 2008; Varshney, 2011). Furthermore, both warm environmental temperatures and cognitive demands can elevate brain temperature. Despite being only 2% of its mass, approximately 20% of the heat released by a human body originates in brain tissue, of which four-fifths is a direct by-product of neuronal signaling (Raichle and Mintun, 2006). Under normal conditions, most excess heat diffuses into the bloodstream and is transported to either the skin or lungs, where it is then transferred to the

environment. When environmental temperatures rise, heat transfer at the skin and lungs slows, reducing the flow of cool blood to the brain, which can temporarily elevate brain temperatures up to 2.5C (Kiyatkin, 2007; Nybo and Secher, 2004).

Higher temperatures could have different effects on different subject areas because they use distinct parts of the brain that are differentially affected by temperature. For example, mathematical problem-solving relies on the ability to retain and manipulate abstract numerical information, drawing heavily on the pre-frontal cortex to supply the “working memory” which stores this data in neural circuits. Performing tasks that utilize working memory when core body temperature is elevated increases neuronal activity in the pre-frontal cortex for any given level of performance, suggesting that working memory is less effective when brain temperature is high (Hocking et al., 2001).

That high temperatures could impair cognitive function is also consistent with experimental evidence that documents impaired brain function in a wide range of domains as a result of heat stress. Military research has shown that soldiers executing complex tasks in hot environments make more errors than soldiers in cooler conditions (Fine and Kobrick, 1978; Froom et al., 1993). Exposure to heat has also been shown to diminish attention, memory, information retention and processing, and the performance of psycho-perceptual tasks (e.g. Vasmatazidis et al., 2002; Hocking et al., 2001; Hyde et al., 1997). The impacts of thermal stress on working memory performance are especially relevant as cognitively challenging tasks rely more heavily on the working memory for multi-step processing.

Given the dynamic nature of human capital production, insults from warmer temperatures may accumulate, leading to decreases in human capital attainment levels. If, however, people respond to short-run deficits with compensatory behaviors, long run changes in human capital may be minimized. One of the key questions in this paper is whether sustained exposure to warmer temperatures, as is expected under climate change, results in

accumulated effects on cognitive ability. We outline a framework for conceptualizing these effects in order to facilitate the interpretation of our econometric models.

We begin with a simple 2-period model of cognitive performance. In the first period, performance y is defined as follows:

$$(1) \quad y_1 = f(k_1, [I-a_1(w_1)]*w_1),$$

where k_1 represents human capital endowments at birth, w_1 is weather exposure in period 1 and a_1 is avoidance behavior in period 1. Avoidance behavior is a transient action, such as turning on air conditioning or staying indoors, which depends upon weather. As such, the second argument in the performance production function ($[I-a_1(w_1)]*w_1$) can be viewed as a measure of the effective exposure to ambient temperatures that results from this behavioral response to local weather conditions. Any time-invariant changes in behavior, such as the adoption of air conditioning, are excluded from this model because they will be captured empirically through the use of various fixed effects.

Performance in the second period is defined similarly to first period performance with two key distinctions. Human capital accumulates from earlier periods and individuals now have the opportunity to respond to feedback embodied in their first period performance through compensatory behaviors. As such, second period performance is expressed as follows:

$$(2) \quad y_2 = f(k_2, [I-a_2(w_2)]*w_2, b(y_1)).$$

As with the initial period, performance will depend on human capital levels and exposure to weather conditions, which depends upon ambient weather and avoidance behavior. For simplicity, we assume that $k_2 = k_1 + g(y_1)$ to reflect human capital accumulation between periods, where the function g reflects the growth in human capital, which depends on prior period learning as reflected by test performance. Compensatory behaviors b are an ex post response to performance in period 1. They could include activities such as spending

additional time studying or the devotion of time and resources to a more formal tutoring relationship. As noted earlier, a key feature of this compensatory behavior is that it does not require that individuals understand that their performance depends on weather.

Our short-run analysis focuses on the impact of weather on the day of the assessment on cognitive performance. Since we do not observe avoidance behavior, our short run estimate reflects the total derivative of y_t with respect to w_t , as follows:

$$(3) \quad dy_t/dw_t = (1-a_t)*\delta f/\delta w_t - w_t*\delta f/\delta a_t*\delta a_t/\delta w_t.$$

The first term represents the direct (neurological) effect of temperature on performance. The second term represents the ex ante behavioral effect of temperature, which depends upon the effectiveness of that avoidance behavior in diminishing the impacts on cognitive performance and the extent of that avoidance behavior.⁸ Our empirical estimates of the short-run impacts will capture the direct effect of temperature net of any avoidance.

Our long run analysis is focused on the impacts of climate on test performance. In our simple framework, climate is simply a combination of weather exposure in both periods. If we define climate as $c = w_1 + w_2$, the impacts of climate on test performance in period 2 is the sum of impacts from contemporaneous and past period weather, plus any component of the ex post response to observed first period performance $b(y_1)$ that is affected by first period weather and the dynamic accumulation of human capital. Thus, our long run estimate reflects the total derivative of y_2 with respect to c , which can be expressed as follows:

$$(4) \quad dy_2/dc = (1-a_2)*\delta f/\delta w_2 - w_2*\delta f/\delta a_2*\delta a_2/\delta w_2 + \delta f/\delta g*\delta g/\delta y_1*dy_1/dw_1 + \delta f/\delta b*\delta b/\delta y_1*dy_1/dw_1,$$

where dy_1/dw_1 is as defined in equation (3).

The first two terms are identical to those in equation (3) and reflect the contemporaneous effect of weather on second period performance – both the direct and ex

⁸ This expression is identical for both first and second period performance.

ante avoidance impacts. The third term ($\delta f/\delta g * \delta g/\delta y_1 * dy_1/dw_1$) captures the impacts of first period weather on learning and thus human capital accumulation by period 2. The fourth term ($\delta f/\delta b * \delta b/\delta y_1 * dy_1/dw_1$) captures the impacts of ex post compensatory behavior. It appears in this climate analysis precisely because compensatory behavior responds to prior period performance. Thus, the difference between the short-run estimates characterized in equation (3) and the long-run estimates described by equation (4) will reflect the accumulated impacts of weather extremes on learning plus the impacts of any ex post compensatory behaviors undertaken.

3. Data

The National Longitudinal Survey of Youth (NLSY) is a nationally representative sample of over 12,000 men and women in the United States aged 14-22 in 1979, with participants surveyed annually until 1994 and biannually thereafter. The survey was designed to collect detailed social and economic information for a transitioning demographic. Beginning in 1986, all children of women in the initial sample were surveyed in their homes, with various developmental assessments conducted biannually on a prearranged date. We focus on examinations in mathematics, reading recognition, and reading comprehension, which are derived from the Peabody Individual Achievement Tests (PIAT) and transformed into age-specific standardized scores.⁹ These tests are designed to measure cognitive achievement and capture gains in knowledge over time, making them a popular measure of human capital in the economics literature (e.g. Todd and Wolpin, 2007). All three tests, which were administered to children age five and over, have been found to have high test-retest and concurrent validity (Rodgers et al., 1994). Each child is tested across multiple waves for as long as the child is part of the survey, with test data available as early as 1988

⁹Despite the availability of additional assessments, we focus solely on these three assessments because they were the most frequently administered across the widest age range, thus yielding the largest sample size and greatest opportunity to explore long run outcomes.

and as late as 2006 depending on the age of the child. In our sample of 8,003 children, 80.9% were tested more than once and 41.2% were tested at least four times, enabling us to precisely estimate within-child effects of temperature. Since these tests were predominantly given during the warmer periods of the year,¹⁰ our analysis of short-run temperature effects will only be informative for temperatures in this range.

Using 8 waves of the geocoded version of the NLSY, which contains the child's county of residence at each survey wave, we match data on each child's test scores with the average temperature in their county on the day of their assessment using data from Schlenker and Roberts (2009), who linearly interpolated temperatures at each county centroid using the seven nearest stations with daily temperature data. County temperature is defined as $(\text{maximum temperature} + \text{minimum temperature})/2$, computed daily at the geographic centroid of each county and matched to the county of residence for each child for each wave of the survey. We also assign precipitation, specific humidity, wind speed, and pressure in an analogous fashion. We repeat a similar procedure for assigning climate, except that we match the full history of temperature (and the other meteorological variables) between successive tests and from birth until the date of the test.

Since temperature is likely to have a non-linear relationship with our outcomes of interest, we use various definitions in our analysis. In the short run analysis, we use both the number of degree days above 21C ($DD > 21$) and below 21C ($DD < 21$), as well as a nonparametric specification with a full set of indicator variables for every 2C. As will become clear, our choice of 21C for the degree day model was chosen based on the nonparametric analysis that revealed a kink at that level. This degree day measure is useful in studies of temperature impacts when (1) a response to daily temperature is roughly constant across days but changes nonlinearly in temperature and (2) the response to daily temperature

¹⁰Assessments were conducted between May and October, except for 1986, which conducted between February and April. To ensure common overlap across seasons and years, however, we exclude the 1986 wave.

can be well approximated by a piece-wise linear function, with kinks at the specified cut-off temperatures.¹¹ The use of indicator variables is even more flexible, allowing for a non-parametric relationship between temperature and performance.

For the long run analysis, we use 3 measures of climate for the between test and lifetime exposure models. First, we take the average of the number of degree days above 21C over the relevant time period. Second, we take the average of the 2C indicator bins for temperature, which amounts to the percent of days in each bin. Third, we calculate the mean January-February and July-August temperatures over the relevant time period to provide estimates with a more intuitive interpretation. We also compute the same time-period averages for humidity, precipitation, wind speed, and pressure.

Table 1 contains summary statistics for the data used in this study. Our final sample includes 8,003 children across 951 counties in 48 states that received exams during multiple survey waves. Children's test scores are at roughly the national median. Since assessments were conducted in warmer months, average temperature exposure on the day of the test is relatively warm, at 22.8°C (73F), reflecting the fact the assessments were conducted during the warmer months. Although children were given the PIAT assessments for all 3 subjects, discrepancies in sample sizes largely reflect differences in the ability to convert raw scores into standardized and percentile scores (Baker and Mott, 1989). Since weather is unrelated to the probability of a test score being available (shown below), we are not concerned that these differences induce a sample selection bias.

4. The Short-Run: Temperature and Cognitive Performance

¹¹Degree days are defined as the number of degrees by which the average daily temperature exceeds 21C, with values below 21C assigned a value of 0. The degree day approach has been widely used to study the non-linear impact of temperature on crop yields (e.g. Schlenker and Roberts, 2009), electricity demand (e.g. Auffhammer and Aroonruengsawat, 2011) and GDP growth (Hsiang, 2010).

To explore the short run relationship between temperature and cognitive performance, we estimate linear fixed effects regression models of the following form:

$$(5) \quad y_{i,t} = f(\beta^{\text{SR}}, T_{c(i),t}) + \eta_1 X_{it} + \eta_2 Z_{c(i),t} + g(t,s(i)) + \alpha_{c(i)} + \varepsilon_{i,t}$$

The test score (y) of child i on date t is regressed on the temperature faced by that child in county c on the same day ($T_{c(i),t}$).¹² β^{SR} reflects dy_t/dw_t from equation (3). We include the child's age ($X_{i,t}$) and other meteorological variables ($Z_{c(i),t}$) that may confound the relationship between temperature and test scores. Our regression models also control for the month and weekday of the assessment and state-specific non-linear time trends in order to capture time-varying factors that influence performance ($g(t,s(i))$). Importantly, these time trends will capture any changes in time-invariant adaptive behaviors during our period, such as air-conditioning penetration or other avoidance 'technologies', to the extent they are common to families in each state.¹³ The longitudinal nature of the survey enables us to specify child fixed effects ($\alpha_{c(i)}$), which controls for all time invariant characteristics of a child. The disturbance term ($\varepsilon_{i,t}$) consists of an individual idiosyncratic component and a clustered component by state-week, which serves three purposes: to allow for arbitrary spatial correlation across counties within a state, to allow for autocorrelation in test scores within a week, and to account for the fact that the same temperature measure can be assigned to multiple children. Since the date the child has the assessment is prearranged, it is unlikely to respond to short-run changes in temperature and thus plausibly exogenous, allowing us to identify the causal effect of temperature on performance.¹⁴

¹² We also explore lagged temperatures as well, shown below.

¹³ 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 (Graff Zivin and Neidell, 2014). We directly probe the role of air conditioning in blunting the short-run impacts of temperature on cognition in the Appendix. While the results are statistically insignificant, our coefficients change in the expected direction, i.e. air conditioning appears to be protective against cognitive harm, though small in magnitude.

¹⁴ We test our exogeneity assumption by separately regressing the probability that a child is male, Hispanic, black, the child's age in months, the child's height in inches, the mothers age at the child's birth, and the birth order of the child, on our full suite of temperature dummies as well as county and state-by-year fixed effects (results available upon request). We find no systematic or significant patterns of selection by these observables with respect to the temperature on the day of the interview and exam.

As described earlier, temperature is included in our model using two distinct approaches to explore its potentially non-linear relationship with performance: (1) a series of indicator variables for temperatures in 2°C bins from 12 to 32°C, with 20-22°C as the reference category; and (2) a linear function in heating and cooling degree days with a cutoff at 21°C, so chosen because the temperature bin at 20-22°C was the local maximum in the first approach.

Results

Table 2 presents the core short-run results for our three test outcomes of interest: math, reading comprehension, and reading recognition. Given our interest in temperature extremes at the high end, we begin with a specification that only includes degree days above 21C. We then add degree days below 21C to capture any effects that might occur at lower temperatures. Columns 1 and 3 present results with mother fixed effects (since siblings are in the sample). Columns 2 and 4 present results with child fixed effects. The math results, shown in Panel A, indicate that warmer temperatures lead to a statistically significant decrease in performance. The results are insensitive to whether we include degree days below 21 and to the type of fixed effect used.¹⁵ The estimate of -0.219 in the first row of Column 4 implies that each degree day above 21C lowers the math score by 0.219 of a percentile.

In contrast, we find that temperature does not have a statistically significant relationship with reading recognition or reading comprehension, regardless of the specification, shown in Panels B and C. As described earlier, one potential explanation for the discrepancy in impacts by subject is that mathematical problem solving utilizes functions

¹⁵ The insignificance of degree days below 21 should be interpreted with caution since very few exams occur on cold days. As such the coefficient on DD<21 largely reflects the impacts of moderate temperatures on test performance.

of the brain that are distinct from the other subject areas, and different parts of the brain are differentially affected by temperature.¹⁶

Figure 1 plots estimates for each of our three outcome variables using the more flexible specification for temperature. Shown in Panel A, we find that child performance in mathematics shows a monotonic decline in outdoor temperatures above 22°C (71.6°F) but is relatively flat and statistically insignificant for temperatures below this point. Furthermore, two of the estimates in the four highest bins are individually statistically significant at the 5% level, with the other two at the 10% level. This monotonic relationship at the high-end reassures us that the significant estimates for math in Table 2 are not simply the result of a Type I error. We interpret the magnitude of our estimates as follows: changing the temperature of the outdoor environment from 20-22°C (68-71.6°F) to 30-32°C (86-89.6°F) decreases a child's mathematics score by 1.6 percentile points, which is a sizeable 0.12 of a standard deviation. The predicted effect using estimates from our degree days specification matches these results quite closely, suggesting that our math estimates are largely insensitive to how we specify our temperature variable.

The nonparametric results for both reading outcomes, shown in Panels B and C of Figure 1, are consistent with the results in Table 2. The coefficients are small, statistically insignificant, and relatively flat across the entire temperature range, providing additional support for the conclusion that performance on these measures is unaffected by temperature.

One concern with interpreting these results is that weather may affect the child's value from alternative activities, and this affects the child's effort on the exam. For example, warmer weather makes playing outside more attractive, and a child may rush through the assessment in order to play outside, which lowers her performance.¹⁷ Given the differential effect by subject, such a mechanism seems unlikely in this setting. Nonetheless, we probe this

¹⁶We note that math is always the first of the three exams, so increased fatigue cannot explain this pattern.

¹⁷In the conceptual framework, this would amount to $\delta b_i/\delta w_i < 0$.

channel using data on the speed at which children complete tests. If children hurry to finish the assessment in warmer weather, then time to completion will fall as temperatures rise. Such data was collected in the 1994, 1996, and 1998 waves, thus providing a useful test for detecting changes in the child's effort on the day of the test.

Consistent with the differential effects by subject, we fail to find evidence to support this channel. Column 1 of Table 3 fails to find a statistically significant relationship between temperature and assessment completion time for all 3 assessments. Allowing for a more flexible specification for temperature, shown for math in Figure 2, we see a generally flat relationship between temperature and time to completion, with only a decline at the highest temperature bin, though it is not statistically significant. This drop at the highest bin, while large in magnitude, does not coincide with the general pattern of temperature effects on performance, suggesting completion time is unlikely to explain our short-run findings.

Table 3 also provides two additional robustness checks. First, the different sample sizes across the subjects (as seen in Table 1), which indicates that some scores are unavailable for children, is a potential concern. In particular, one might worry about sample selection bias if the missing test scores correlate with warmer temperatures, perhaps because families cancel the visit or the child scored below a certain value, making a standardized score infeasible. To assess this, we regress our weather variables on score availability. Shown in column 3 of Table 4, we find that probability of completing the assessment is unrelated to warmer weather, suggesting that sample composition across subjects is unlikely to bias our results. Second, one might be concerned that exams are shifted to cooler times of the day to avoid peak exposure, thereby minimizing the effect on performance. In column 4, we show results using the start time of the assessment as the dependent variable, and find the start time is unrelated to the temperature on the day of the test.¹⁸

¹⁸ We only show results for math because it is always the first test given.

Our analysis has thus far focused solely on the effects of weather on the day of the assessment, thereby ignoring potential lagged effects. While the neurological mechanisms discussed in Section 2 suggest a rather immediate effect from exposure, lagged exposure has been shown to affect health and thus might also affect performance.¹⁹ Figure 3 presents results when we add 3 lags of temperature, and also 1 lead of temperature as a falsification test. The coefficient on contemporaneous temperature is largely unchanged (though precision is comprised) and the coefficients on lagged temperature are considerably smaller than our main estimate.²⁰ The absence of an effect for lead temperature offers further assurance that our results are not driven by unobserved confounding. In the end, and despite moderate to high-levels of air-conditioning penetration, these results provide strong evidence for a contemporaneous and negative effect of warmer temperatures on mathematical performance.²¹

Before turning to our long-run analysis, we simulate the potential long run effects based on our short run estimates to give a sense for the potential magnitude of the effects we might find. To do so, we need to make some assumptions about the human capital accumulation process (the function $g(\cdot)$ in our conceptual model). First, we assume children's percentile performance on the test is equivalent to a ranking in human capital. Absent negative shocks, children accumulate human capital at comparable rates so that their rank remains unchanged. Second, exposure to a simulated distribution of weather shocks, measured in 2C bins, leads to a reduction in rank. We assume this reduction is permanent. Further, we assume that a change in performance, absent compensatory behaviors, amounts

¹⁹Furthermore, since we do not know the exact time assessments were given for all years, we may be assigning weather with error. A lagged specification may better capture exposure for those with, for example, morning exams. We note, however, that for the sample years in which we observe assessment times, the average start time is 2:41pm.

²⁰The effect of the prior day's temperature is roughly half the size of our main estimate, suggesting that temperature impacts on cognitive function might accumulate over short periods or that hot nights might interfere with sleep, although this effect is not significant.

²¹As noted earlier, approximately 80% of U.S. households had some form of air conditioning in 2001, with the rate of ownership much higher in warmer regions (Graff Zivin and Neidell, 2014).

to a change in learning. Although our short run estimates do not directly test for learning effects, the ergonomics literature reviewed above demonstrates effects on memory, information retention, and information processing, suggesting the plausibility of such an effect. Given uncertainties about the magnitude of this learning effect, we scale the change in performance by λ , under a variety of parameter values $\lambda = \{0.1, 0.3, 0.5\}$. For example, $\lambda=0.5$ implies that a 10 percentage point decrease in test performance translates into a 5 percentage point decrease in learning. We then accumulate the shocks between tests based on the realized weather exposure for each child to compute the change in human capital. Our simulation based on our short-run estimates under the quite conservative assumption that $\lambda=0.1$ implies that children’s exposure between tests would reduce performance by 6.2 percentile points on average. At $\lambda=0.3$ and 0.5 , this rises to 18.6 and 31.0 percentile points, respectively, thereby implying quite large long run effects.

5. The Long-Run: Climate and Human Capital

The analysis thus far focused on the contemporaneous impacts of temperature on performance. In this section, we turn our attention to the long run impacts of climate on human capital, where additional adaptation strategies are expected to play a significant role. While our short-run results are net of any avoidance behavior undertaken to minimize exposure to temperature extremes, individuals can respond to outcomes (as opposed to weather) over the longer-run. To capture these compensatory behaviors, we estimate two distinct models.

In our first approach, we estimate a “long difference” model of the following form:

$$(6) \quad y_{i,t} - y_{i,t-1} = f(\beta^{LD}, \sum_{t-1}^t T_{c(i)}) + \eta_1(X_{it} - X_{it-1}) + \eta_2 \sum_{t-1}^t Z_{c(i)} + g(t, s(i)) + (\varepsilon_{i,t} - \varepsilon_{i,t-1})$$

The dependent variable is the change in performance over time, which reflects the accumulation of human capital between tests. The variable $\sum_{t-1}^t T_{c(i)}$ reflects our measure of

climate, which is a summary measure of temperature between successive tests. We continue to define $T_{c(i)}$ in degree days and indicator bins as before. Given the different structure of this model, the interpretation of β^{LD} now takes a slightly different form. When we use degree days, we interpret β^{LD} as the increase in human capital from a 1C degree day increase in temperature across all days between tests. When we use indicator bins, we interpret β^{LD} as the increase in human capital from a 1 percent increase in the number of days that the temperature falls in a certain bin (relative to 20-22C) between tests. For example, we would interpret the coefficient on the 30-32C bin as the effect from shifting 1 percent of all days between successive tests from 20-22C to 30-32C. To better align with intuition we also use seasonal average temperature (separately for January-February and July-August). In this case, the coefficients reflect the impacts of a 1C increase in the mean July-August (or January-February) temperature between tests on human capital accumulation. The other meteorological variables (Z) are defined analogously, while the variables X_{it} and $g(t,s(i))$ remain unchanged from equation (3).

By defining the model in long differences (LD), the model captures a wider range of adaptive responses (Dell et al., 2014; Burke and Emerick, 2012), where the coefficient $\beta^{LD} = dy_2/dc$ from equation (4). For example, if parents respond to poor performance in school with compensatory investments, regardless of whether they know the source of the poor performance, our estimate for β^{LD} is net of this investment. The model also remains well-identified because we are controlling for all time invariant characteristics of the child.

In our second approach, we assign climate as the accumulated temperature from birth until the date of the test, hence providing an even longer-term measure of climate exposure. This necessitates the use of cross-sectional models, which leaves greater scope for omitted variable bias since parents choose where to raise their children and thus climate exposure may be correlated with other characteristics that influence human capital attainment. To

address this concern, we exploit the unusual richness of the NLSY to control for a wide range of background factors in the human capital production function (Black et al., 2005). In particular, we estimate the following regression specification:

$$(7) \quad y_{i,t} = f(\beta^{CS}, \sum_0^t T_{c(i)}) + \eta_1 X_{it} + \eta_2 \sum_0^t Z_{c(i)} + \eta_3 X_{m(i)} + g(t,s(i)) + \varepsilon_{i,t}$$

Climate ($\sum_{t-1}^t T_{c(i)}$) is now measured as lifetime exposure from birth until the time of the test, and we continue to specify this in terms of degree days, indicator bins, and seasonal averages. The interpretation of coefficients is similar to the “between-test” model except they now reflect the effect from birth until the time of the test. The other meteorological variables (Z) are defined analogously.

Given the greater concern for omitted variable bias in this specification, we also add several measures that reflect the child’s potential human capital endowment at birth. X_i now includes the child’s birth weight, an important measure of intellectual endowments (Black et al., 2007), which we control for flexibly by including a series of indicator variables for each pound. It also includes the child’s sex, birth order indicators, and maternal age at birth. $X_{m(i)}$ includes the mother’s scores on the armed forces qualifying test (AFQT), completed years of schooling, a measure of self-esteem, height, weight, race/ethnicity, foreign language, the religion she was raised, and her spouse's level of education. We also include flexible controls by allowing for all two-way interactions between these variables and 3rd order polynomials for all continuous variables. Including grandparent characteristics (grandmother and grandfather's years of schooling, Duncan SEI, foreign born) in $X_{m(i)}$ further helps to minimize concerns about omitted variable bias. Z is also extended to include numerous county level characteristics, including the age of housing stock, birth rate, death rate, infant mortality rate, physicians per capita, hospital beds per capita, education per capita, household income per capita, fraction below poverty, geographic size of the county, maximum elevation, and whether it borders an ocean or great lake.

Results

Table 4 presents our long-difference results based on the “between test” specification. We focus solely on mathematical performance since this is the only outcome where we find an effect in the short run. In contrast to the short run results, however, we do not find a statistically significant relationship between climate and human capital. Column 1, which only includes degree days above 21C (DD>21), reveals a statistically insignificant estimate of -0.630. This estimate indicates that a 1 degree day increase in temperature across all days between two tests, a rather substantial change, decreases math performance by only 0.630 percentile points. Under the conservative case where $\lambda = 0.1$, this result implies that parent are offsetting more than a 6 percentile point accumulated decrease in human capital due to warmer temperature exposures. Since the absence of a long-term impact is somewhat imprecise, it may also be instructive to look at the lower bound of the 95% confidence interval, which corresponds to an estimate of -1.304. Even at this tail of the parameter distribution, our results imply considerable adaptive behaviors or a λ no larger than an implausibly small value of 0.007. When we add degree days below 21C (DD<21), shown in column 2, our estimate for warmer temperatures rises slightly to -0.250 but remains statistically insignificant. Focusing on mean winter and summer temperatures yields estimates that are again statistically insignificant and considerably smaller than the simulated long run estimates under even highly conservative assumptions. For example, our estimate of -0.196 for July-August suggests that a 1C increase for every day in those two months decrease math performance by 0.196 percentile points.

In Table 5, we show results using lifetime temperature exposure. Given that these estimates rely on cross-sectional models, we assess the sensitivity of results to slowly adding more controls, continuing to use different assumptions about the functional form for

temperature. In the first panel, which only uses degree days above 21C, we see in column 1 that temperature is negatively associated with math performance. Adding simple controls for maternal human capital, a strong predictor of children's human capital attainment (Black et al., 2005), raises that coefficient to -0.463 as shown in Column 2. As we include more control variables, this estimate remains fairly stable and far from statistically significant. Panel B adds degree days below 21, while Panel C uses winter and summer temperatures. In both cases, the estimates show the same general pattern: statistically insignificant estimates that are much smaller than those implied by the simulation exercise.

In Figure 4 we show results allowing for the flexible specification in temperature for the “between test” and “from birth” models that matches the indicators used in the short run model. As with the previous long run results, we do not find statistically significant estimates. Moreover, we do not find a pattern in the estimates that comports with the short run results – estimates are relatively flat over the entire temperature distribution. The consistently insignificant and small coefficients across our long-run models are consistent with the notion that individuals engage in non-trivial amounts of adaptation to minimize the effects of high temperature days on the human capital accumulation of children.

6. Conclusion

In this paper, we merge rich data from the NLSY with meteorological data to provide the first economic analysis of the relationship between temperature/climate and human capital. We find that short-run changes in temperature lead to statistically significant decreases in cognitive performance on math (but not reading) beyond 26C (78.8F). Notably, these results obtain despite quite high levels of air conditioning penetration in our study region, suggesting that in the short run, individuals do not completely insulate themselves from climatic factors. In contrast, our long-run analysis fails to find a statistically significant

relationship between climate and human capital, and one that is much smaller in magnitude than that suggested by the short run estimates. These results are consistent with the notion that individuals compensate for short-run setbacks with additional investments in subsequent periods.

This set of results is important for several reasons. Our short-run results indicate that analytical thinking is compromised at modest temperatures well below our popular conventions regarding a very hot day. Cognitive performance of this sort is the lifeblood of homo economicus and critical for decision making in a wide range of domains. That this temperature range is a regular occurrence in summer across much of the globe and all year long in parts of the tropics portends potentially sizable impacts on economic well-being. These findings also appear to have strong implications for the optimal timing of cognitively demanding tasks, such as financial decision-making and significant health choices.

While cognitive performance and decision making may be compromised by warmer weather, our long-run results demonstrate that these insults have no demonstrable effect on human capital attainment in the long-run. Since permanent adaptation strategies are largely held fixed in our comparisons across our short and long run specifications, we argue that the difference between these results are driven by compensatory behavior. An interesting feature of this compensatory behavior is that it requires no knowledge of the ‘harmful’ effects of temperature since it is an ex-post adaptive strategy. The feedback from poor test performance may be sufficient to induce individuals to increase investments in learning.

It is important to note, however, that there may be an alternative explanation for the absence of a long-run effect. Test scores are a composite measure of knowledge and performance, and the inter-temporal dependencies of one on the other are largely unknown. Thus, it is possible that short-run changes in performance simply do not add up to sizable long-run changes in learning. Given our parameter estimates and simulations, the plausibility

of this explanation hinges on a near-zero relationship between the two. Evidence from randomized laboratory studies which find that learning is compromised from higher temperatures further bolsters the claim that the difference between our short- and long-run estimates reflects the offsetting effects of compensatory behaviors.

Finally, our results suggest a cautionary tale regarding what we can learn about policy-relevant climate change impacts from studying the well-identified influences from weather. Whether the same pattern applies to other settings is an open question. Regardless, the absence of long-run impacts in our setting should not be interpreted as the absence of welfare impacts. Compensatory behavior is not costless, and represents a potentially important, and heretofore largely ignored, element in the optimal design of climate change policies.

References

- Currie, Janet, and Douglas Almond. "Human capital development before age five." *Handbook of Labor Economics* 4 (2011): 1315-1486.
- Auffhammer, Maximilian, and Anin Aroonruengsawat. "Simulating the impacts of climate change, prices and population on California's residential electricity consumption." *Climatic change* 109, no. 1 (2011): 191-210.
- Baker, Paula C., and Frank Mott. *NLSY child handbook 1989: A guide and resource document for the National Longitudinal Survey of Youth 1986 Child Data*. Center for Human Resource Research, The Ohio State University, 1989.
- Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S. Shapiro. "Adapting to climate change: The remarkable decline in the US temperature-mortality relationship over the 20th century." NBER Working Paper 18692, 2013.
- Black, Sandra E., Paul J. Devereux, and Kjell G. Salvanes. "Why the Apple Doesn't Fall Far: Understanding the Intergenerational Transmission of Education." *American Economic Review* 95(2005): 437-449.
- Black, Sandra E., Paul J. Devereux, and Kjell G. Salvanes. "From the Cradle to the Labor Market? The Effect of Birth Weight on Adult Outcomes." *The Quarterly Journal of Economics* 122, no. 1 (2007): 409-439.
- Bowler, K., and R. Tirri. "The temperature characteristics of synaptic membrane ATPases from immature and adult rat brain." *Journal of neurochemistry* 23, no. 3 (1974): 611-613.
- Burke, Marshall B., Edward Miguel, Shanker Satyanath, John A. Dykema, and David B. Lobell. "Warming increases the risk of civil war in Africa." *Proceedings of the National Academy of Sciences* 106, no. 49 (2009): 20670-20674.
- Burke, Marshall, and Kyle Emerick. "Adaptation to climate change: Evidence from US agriculture." Available at SSRN 2144928 (2012).
- Cunha, Flavio, James J. Heckman, and Susanne M. Schennach. "Estimating the technology of cognitive and noncognitive skill formation." *Econometrica* 78(3) (2010): 883-931.
- Currie, Janet. "Healthy, Wealthy, and Wise: Socioeconomic Status, Poor Health in Childhood, and Human Capital Development." *Journal of Economic Literature* 47, no. 1 (2009): 87-122.
- Deboer, Tom. "Brain temperature dependent changes in the electroencephalogram power spectrum of humans and animals." *Journal of sleep research* 7, no. 4 (1998): 254-262.
- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. "What do we learn from the weather? The new climate-economy literature." *Journal of Economic Literature* 52, no 3 (2014): 740-798.

- Dell, Melissa, Benjamin F. Jones, and Benjamin A. Olken. "Temperature shocks and economic growth: Evidence from the last half century." *American Economic Journal: Macroeconomics* 4, no. 3 (2012): 66-95.
- Deschenes, Olivier, and Michael Greenstone. "The economic impacts of climate change: evidence from agricultural output and random fluctuations in weather." *The American Economic Review* (2007): 354-385.
- Deschênes, Olivier, and Michael Greenstone. "Climate Change, Mortality, and Adaptation: Evidence from Annual Fluctuations in Weather in the US." *American Economic Journal: Applied Economics* (2011): 152-185.
- Deschenes, Olivier, and Enrico Moretti. "Extreme weather events, mortality, and migration." *The Review of Economics and Statistics* 91, no. 4 (2009): 659-681.
- Doya, Kenji, Shin Ishii, Alexandre Pouget, and Rajesh PN Rao. "Bayesian Brain: Probabilistic Approaches to Neural Coding (Computational Neuroscience)." (2007).
- Graff Zivin, Joshua and Matthew Neidell, "Temperature and the Allocation of Time: Implications for Climate Change," *Journal of Labor Economics*, 32(2014): 1-26.
- Graff Zivin, Joshua, and Matthew Neidell. "Environment, Health, and Human Capital." *Journal of Economic Literature* 51, no. 3 (2013): 689-730.
- Feng, Shuaizhang, Alan B. Krueger, and Michael Oppenheimer. "Linkages among climate change, crop yields and Mexico–US cross-border migration." *Proceedings of the National Academy of Sciences* 107, no. 32 (2010): 14257-14262.
- Fine, Bernard J., and John L. Kobrick. "Effects of altitude and heat on complex cognitive tasks." *Human Factors: The Journal of the Human Factors and Ergonomics Society* 20, no. 1 (1978): 115-122.
- Froom, Paul, Yeheskial Caine, Igal Shochat, and Joseph Ribak. "Heat stress and helicopter pilot errors." *Journal of Occupational and Environmental Medicine* 35, no. 7 (1993): 720-732.
- Hocking, Chris, Richard B. Silberstein, Wai Man Lau, Con Stough, and Warren Roberts. "Evaluation of cognitive performance in the heat by functional brain imaging and psychometric testing." *Comparative Biochemistry and Physiology Part A: Molecular & Integrative Physiology* 128, no. 4 (2001): 719-734.
- Hsiang, Solomon M. "Temperatures and cyclones strongly associated with economic production in the Caribbean and Central America." *Proceedings of the National Academy of Sciences* 107, no. 35 (2010): 15367-15372.
- Hyde, Dale, John R. Thomas, John Schrot, and W. F. Taylor. *Quantification of special operations mission-related performance: Operational evaluation of physical measures*. No. NMRI-97-01. Naval Medical Research Institute, Bethesda MD, 1997.

IPCC, 2007: Climate Change 2007: The Physical Science Basis. Contribution of Working Group I to the Fourth Assessment Report of the Intergovernmental Panel on Climate Change [Solomon, S., D. Qin, M. Manning, Z. Chen, M. Marquis, K.B. Averyt, M. Tignor and H.L. Miller (eds.)]. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA.

Kiyatkin, Eugene A. "Brain temperature fluctuations during physiological and pathological conditions." *European Journal of Applied Physiology* 101, no. 1 (2007): 3-17.

Libecap, Gary D., and Richard H. Steckel, eds. *The economics of climate change: Adaptations past and present*. University of Chicago Press, 2011.

Mendelsohn, Robert, William D. Nordhaus, and Daigee Shaw. "The impact of global warming on agriculture: A Ricardian analysis." *The American Economic Review* (1994): 753-771.

Moore, Christopher I., and Rosa Cao. "The hemo-neural hypothesis: On the role of blood flow in information processing." *Journal of Neurophysiology* 99, no. 5 (2008): 2035-2047.

Nelson, Richard R., and Edmund S. Phelps. "Investment in humans, technological diffusion, and economic growth." *The American Economic Review* (1966): 69-75.

Nordhaus, William D. "Geography and macroeconomics: New data and new findings." *Proceedings of the National Academy of Sciences of the United States of America* 103, no. 10 (2006): 3510-3517.

Nordhaus, William D. "The economics of hurricanes and implications of global warming." *Climate Change Economics* 1, no. 01 (2010): 1-20.

Nybo, Lars, and Niels H. Secher. "Cerebral perturbations provoked by prolonged exercise." *Progress in Neurobiology* 72, no. 4 (2004): 223-261.

Raichle, Marcus E., and Mark A. Mintun. "Brain work and brain imaging." *Annual Review of Neuroscience* 29 (2006): 449-476.

Romer, Paul M. "Increasing returns and long-run growth." *The Journal of Political Economy* (1986): 1002-1037.

Rodgers, Joseph Lee, David C. Rowe, and Kim May. "DF analysis of NLSY IQ/achievement data: Nonshared environmental influences." *Intelligence* 19, no. 2 (1994): 157-177.

Sachs, Jeffrey D., and Andrew M. Warner. "Sources of slow growth in African economies." *Journal of African Economies* 6, no. 3 (1997): 335-376.

Schlenker, Wolfram, W. Michael Hanemann, and Anthony C. Fisher. "The impact of global warming on US agriculture: an econometric analysis of optimal growing conditions." *Review of Economics and Statistics* 88, no. 1 (2006): 113-125.

Schlenker, Wolfram, and Michael J. Roberts. "Nonlinear temperature effects indicate severe damages to US crop yields under climate change." *Proceedings of the National Academy of Sciences* 106, no. 37 (2009): 15594-15598.

Schiff, Steven J., and George G. Somjen. "The effects of temperature on synaptic transmission in hippocampal tissue slices." *Brain Research* 345, no. 2 (1985): 279-284.

Sinha, Paramita, and Maureen L. Cropper. "The value of climate amenities: Evidence from us migration decisions." NBER Working Paper 18756, 2013.

Todd, Petra E., and Kenneth I. Wolpin. "The production of cognitive achievement in children: Home, school, and racial test score gaps." *Journal of Human capital* 1.1 (2007): 91-136.

Varshney, Lav R. "Performance of LDPC codes under faulty iterative decoding." *Information Theory, IEEE Transactions on* 57, no. 7 (2011): 4427-4444.

Vasmatazidis, Ioannis, Robert E. Schlegel, and Peter A. Hancock. "An investigation of heat stress effects on time-sharing performance." *Ergonomics* 45, no. 3 (2002): 218-239.

Yablonskiy, Dmitriy A., Joseph JH Ackerman, and Marcus E. Raichle. "Coupling between changes in human brain temperature and oxidative metabolism during prolonged visual stimulation." *Proceedings of the National Academy of Sciences* 97, no. 13 (2000): 7603-7608.

Table 1. Summary Statistics

A. Cognitive outcomes

		N	Mean	SD	within SD
Math	percentile	24361	49.70	27.45	14.37
	length	10389	9.98	4.91	3.14
	energy	24260	0.11	0.31	0.25
Reading comprehension	percentile	20439	51.60	27.74	14.34
	length	8557	10.19	5.26	3.15
	energy	20041	0.16	0.36	0.28
Reading recognition	percentile	24229	56.52	28.50	13.00
	length	10367	3.59	1.92	1.25
	energy	22814	0.15	0.35	0.28
# of tests per child		24361	3.66	1.20	-
# of years between tests		16304	2.15	0.64	-

B. Temperature measures

	N	°C	SD	within SD	°F
<u>Day of test</u>					
Temperature	24361	22.77	4.96	3.12	72.99
Degree days \geq 21	24361	3.05	3.12	1.85	5.49
Degree days < 21	24361	1.28	2.65	1.92	2.30
<u>Between tests</u>					
Degree days \geq 21	16304	1.21	0.99	0.22	2.18
Degree days < 21	16304	13.07	3.75	0.71	23.52
January-February	16304	4.61	6.86	1.33	40.30
July-August	16304	24.57	3.03	0.63	76.23
<u>From birth</u>					
Degree days \geq 21	24294	1.23	0.98	-	2.22
Degree days < 21	24294	13.08	3.68	-	23.54
January-February	24294	4.12	6.85	-	39.42
July-August	24294	24.53	2.95	-	76.15

Notes: Length is the number of minutes for the child to complete the evaluation. Energy is the percent of children reported by the interviewer as acting lethargic. Temperature measures are reported for the math test only. 'Day of test' reflects temperature the day the test is taken. 'Between tests' reflects the mean temperature between successive tests. 'From birth' reflects the mean temperature from birth until the time of the test. Degree days \geq 21 (<21) is the number of degrees above (below) 21°C. January-February and July-August are the mean temperatures for those months.

Table 2. Fixed effect estimates of relationship between short-run temperature and cognitive performance

	1	2	3	4
A. Math				
Degree days \geq 21	-0.211* [0.0903]	-0.205* [0.0960]	-0.240** [0.0925]	-0.219* [0.0984]
Degree days < 21			-0.151 [0.0899]	-0.0749 [0.0934]
Fixed effect	mother	child	mother	child
Observations	24,361	24,361	24,361	24,361
R-squared	0.551	0.737	0.551	0.737
B. Reading comprehension				
Degree days \geq 21	0.0607 [0.102]	0.0611 [0.102]	0.0524 [0.104]	0.0711 [0.104]
Degree days < 21			-0.0434 [0.0942]	0.0509 [0.0985]
Fixed effect	mother	child	mother	child
Observations	20,439	20,439	20,439	20,439
R-squared	0.601	0.779	0.601	0.779
B. Reading recognition				
Degree days \geq 21	-0.027 [0.0899]	0.0441 [0.0875]	-0.0325 [0.0919]	0.0461 [0.0896]
Degree days < 21			-0.0286 [0.0856]	0.0101 [0.0828]
Fixed effect	mother	child	mother	child
Observations	24,229	24,229	24,229	24,229
R-squared	0.587	0.802	0.587	0.802

Notes: The above coefficients reflect estimates of the relationship between temperature on the day of the test and cognitive performance. Standard errors clustered on state-week in brackets. ** $p < 0.01$, * $p < 0.05$. All regression models control for precipitation, pressure, wind speed, humidity, and dummy variables for day of week, month, year, and state-by-year. Regressions with mother fixed effects also control for child sex, birth order dummies, age of mother at birth of child, and child birth weight dummies.

Table 3. Robustness checks for estimates of short-run temperature relationship

	1	2	3
	Time to completion	Assessment completed	Start time
A. Math			
Degree days \geq 21	-0.011 [0.054]	-0.0012 [0.0012]	0.0138 [0.0447]
Degree days < 21	0.004 [0.052]	-0.0015 [0.0013]	-0.0282 [0.0418]
Observations	8,620	26,091	8,621
R-squared	0.663	0.566	0.707
B. Reading comprehension			
Degree days \geq 21	-0.038 [0.064]	-0.0008 [0.0019]	
Degree days < 21	-0.025 [0.069]	-0.0038* [0.0019]	
Observations	7,092	26,062	
R-squared	0.725	0.52	
C. Reading recognition			
Degree days \geq 21	-0.011 [0.023]	-0.0010 [0.0013]	
Degree days < 21	-0.051* [0.024]	-0.0015 [0.0014]	
Observations	8,597	26,089	
R-squared	0.65	0.55	

Notes: The above coefficients reflect estimates of the relationship between temperature on the day of the test and the following outcomes: column 1 is the number of minutes it takes the child to complete the assessment; column 2 is whether the assessment is completed, and column 3 is the start time of the assessment, only relevant for math since it is always first. Standard errors clustered on state-week in brackets. ** $p < 0.01$, * $p < 0.05$. All regression models control for precipitation, pressure, wind speed, humidity, dummy variables for day of week, month, year, and state-by-year, and child fixed effects.

Table 4. Long difference estimates of relationship between temperature between tests and math performance

	1	2	3
Degree days \geq 21	-0.630 [0.344]	-0.250 [0.466]	
Degree days < 21		-0.414* [0.203]	
July-August			-0.196 [0.134]
January-February			-0.0905 [0.135]
Observations	16,304	16,304	16,304
R-squared	0.034	0.035	0.035

Notes: The above coefficients reflect estimates of the relationship between temperature exposure between tests and the change in math performance. Standard errors clustered on state-week in brackets. ** $p < 0.01$, * $p < 0.05$. All regression models control for between test measures of precipitation, pressure, wind speed, and humidity; dummy variables for day of week, month, year, and state-by-year.

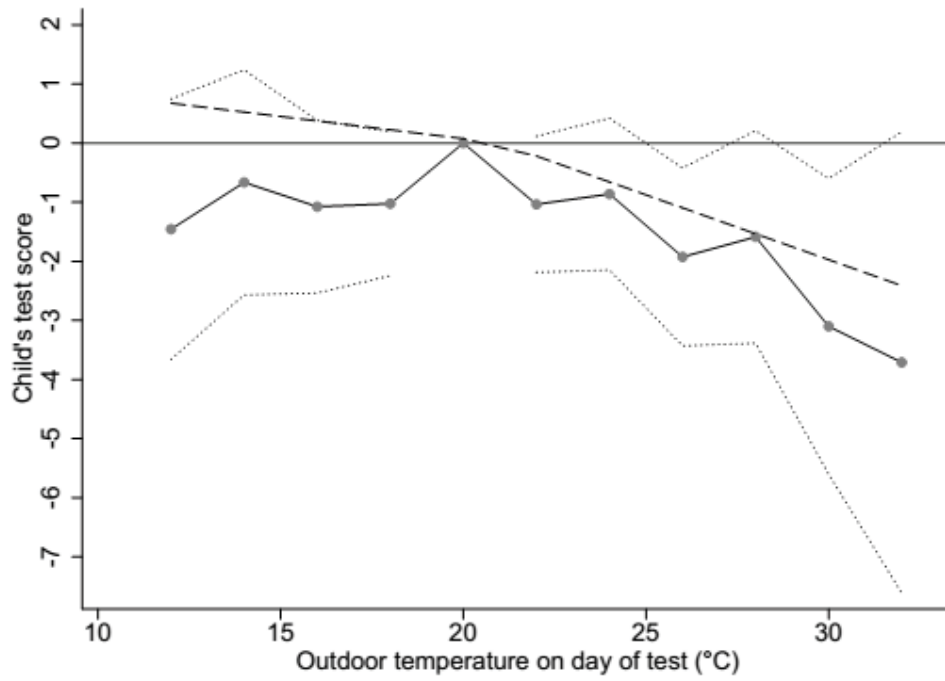
Table 5. Cross-sectional estimates of relationship between lifetime temperature exposure and math performance

	1	2	3	4
A. Cooling degree days				
Degree days \geq 21	-1.211 [0.977]	-0.463 [1.116]	-0.440 [1.087]	-0.329 [0.999]
Observations	24,294	24,294	24,294	24,294
R-squared	0.151	0.269	0.271	0.281
B. Cooling and heating degree days				
Degree days \geq 21	0.347 [0.962]	0.0368 [1.009]	0.102 [0.970]	0.158 [0.918]
Degree days $<$ 21	-1.694** [0.548]	-0.550 [0.430]	-0.600 [0.440]	-0.543 [0.420]
Observations	24,294	24,294	24,294	24,294
R-squared	0.152	0.269	0.271	0.281
C. Seasonal temperatures				
July-August	-0.691* [0.297]	-0.496 [0.291]	-0.422 [0.312]	-0.349 [0.302]
January-February	-0.285 [0.281]	0.237 [0.237]	0.139 [0.235]	0.125 [0.230]
Observations	24,294	24,294	24,294	24,294
R-squared	0.155	0.270	0.273	0.282
County chars.	Y	Y	Y	Y
Geography	Y	Y	Y	Y
Child chars.	Y	Y	Y	Y
Maternal HC	N	Y	Y	Y
Grandparent HC	N	N	Y	Y
f(Maternal HC)	N	N	N	Y


Notes: The above coefficients reflect estimates of the relationship between temperature exposure from birth until the time of test and math performance. Standard errors clustered on state-week in brackets. ** $p < 0.01$, * $p < 0.05$. All regression models control for from-birth measures of precipitation, pressure, wind speed, humidity, and dummy variables for day of week, month, year, and state-by-year. County chars. includes county level measures of age of housing stock, birth rate, death rate, infant mortality rate, physicians per capita, hospital beds per capita, education per capita, household income per capita, and fraction below poverty. Geography includes county size, max. elevation, and borders ocean or great lake. Child chars. includes sex, birth order dummies, maternal age at birth, and age of child at time of test. Maternal HC includes mother's years of schooling, AFQT, self-esteem, height, weight, race, foreign language, religion, and dummy variables to indicate when schooling, AFQT and self-esteem were imputed. Grandparent HC includes grandmother and grandfather's years of schooling, Duncan SEI, foreign born, and dummies if schooling missing. f(maternal HC) includes 3rd order polynomial for all continuous maternal HC variables and all 2-way interactions.

Figure 1. Relationship between short-run temperature and cognitive performance


A. Mathematics



B. Reading comprehension




C. Reading recognition




Notes: The solid line shows coefficient estimates of the effect of temperature on the day of the test of cognitive performance, with 95% confidence intervals in dotted lines. The regression includes indicators for each 2°C, and also controls for precipitation, pressure, wind speed, humidity, dummy variables for day of week, month, year, and state-by-year, and child fixed effects. The predicted effect using degree days above and below 21°C is shown in the dashed line, and is based on a regression using the same set of controls.

Figure 2. Relationship between short-run temperature and time to completion for math assessment



Notes: The solid line shows coefficient estimates of the effect of temperature on the day of the test, with 95% confidence intervals in dotted lines. The regression includes indicators for each 2°C, and also controls for precipitation, pressure, wind speed, humidity, dummy variables for day of week, month, year, and state-by-year, and child fixed effects. The predicted effect using degree days above and below 21°C is shown in the dashed line, and is based on a regression using the same set of controls.


Figure 3. Timing of exposure and math performance



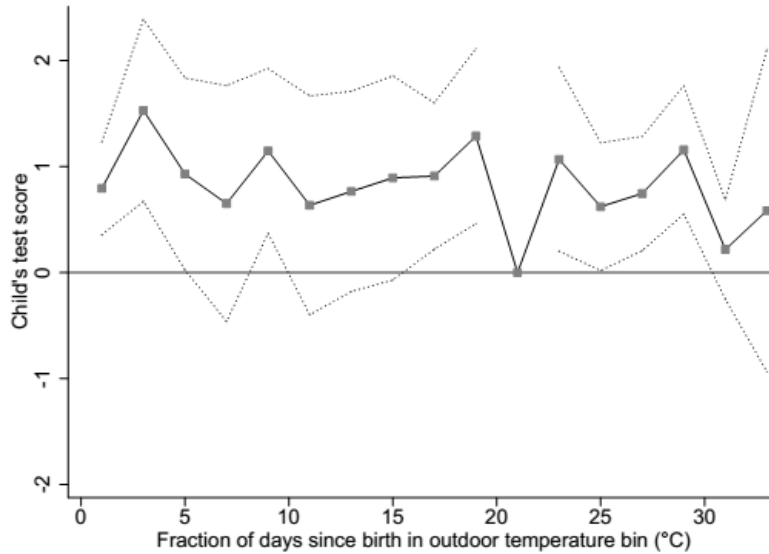
Notes: The solid line shows coefficient estimates of the effect of degree days ≥ 21 on cognitive performance for 3 days before, the day of, and one day after the test. The regression also controls for degree days < 21 3 days before, the day of, and one day after the test. Other controls include precipitation, pressure, wind speed, humidity, dummy variables for day of week, month, year, and state-by-year, and child fixed effects. The estimate for the contemporaneous only model (Table 2, column 4) is shown with the gray square, along with the 95% confidence interval.

Figure 4. Relationship between long-run temperature and math performance

A. Temperature exposure between successive tests



B. Temperature exposure from birth until time of test



Notes: The solid line shows coefficient estimates with 95% confidence intervals in dotted lines. Panel A focuses on measures of temperature between successive tests. Panel B focuses on measures of temperature from birth until the time of the test. The regressions include indicators for the fraction of days the temperatures was in each 2°C bin, and also controls for precipitation, pressure, wind speed, humidity (measured analogously), dummy variables for day of week, month, year, and state-by-year. Panel B includes the full set of controls as used in column 4 of Table 5.

Appendix – Exploring the role of air conditioning

While our short-run results are net of any protective effects from the availability of air conditioning at the time children take their exam, we use this Appendix to directly explore the role of air conditioning as a moderator of the short-run relationship between temperature and cognitive performance. A core challenge in this analysis is the limited availability of data on air conditioning (AC) penetration, particularly over the relevant spatial and temporal scales. The best data available comes from the American Housing Survey (AHS), and even this data is rather incomplete for our purposes. As a result, our analysis relies on imputed estimates of AC ownership based on the following procedure. First, we create a balanced panel of 131,904 observations from 21,984 housing units from biannual observations of the 2001-2011 AHS surveys.²² From this survey, we create an indicator for whether the housing unit has a window or central AC unit. We then perform a logit regression of AC ownership on indicators for race and age of the head of household, SMSA and year. We then map the NLSY counties into SMSAs, though 1/3 of NLSY respondents could not be mapped into an SMSA since not all counties reside in an SMSA. Using the coefficients from this equation, we then predict the probability of AC ownership in the NLSY. Given data limitations and the strong assumptions required to construct this measure, the results that follow should be viewed as suggestive and treated with considerable caution.

Using our imputed measure of AC ownership, we interact AC ownership with our temperature variables as given in equation (5). The results for degree days are shown in Appendix Table 1. Re-estimating equation (5) for this reduced sample gives an estimate of -0.133, which is somewhat smaller than the corresponding estimate in Table 2 of -.205. When we include interactions with AC, the level effect of $DD > 21$, which should be interpreted as

²² Although the NLSY test score data we use goes back to 1988, 2001 is the first year of the AHS survey. We are unaware of any national sources of AC data with geographic identifiers that covers the years between 1988 and 2000.


the effect if someone does not own AC, drops to $-.464$. The interaction term is 0.415 , which implies that the effect of temperature for those with AC is $-.049$. While imprecise, this pattern is as expected: owning AC decreases the effect of temperature on cognitive performance. We find a similar pattern when we include $DD < 21$, although we obtain the surprising result that $DD < 21$ is also affected by AC ownership. The results using the more flexible specification for temperature are shown in Appendix Figure 1. While estimates are generally decreasing in temperature, the interaction with AC ownership has little impact on our estimates.

Appendix Table 1. Relationship between short-run temperature, air conditioning, and cognitive performance

	1	2	3	4
Degree days \geq 21	-0.133 [0.118]	-0.464 [0.378]	-0.148 [0.121]	-0.379 [0.396]
Degree days < 21			-0.0874 [0.127]	-0.314 [0.291]
AC*Degree days \geq 21		0.415 [0.453]		0.287 [0.480]
AC*Degree days < 21				0.345 [0.388]
Observations	15,719	15,719	15,719	15,719
R-squared	0.745	0.745	0.745	0.745

Notes: The above coefficients reflect estimates of the relationship between temperature on the day of the test and math performance. Standard errors clustered on state-week in brackets. ** $p < 0.01$, * $p < 0.05$. All regression models control for precipitation, pressure, wind speed, and humidity; dummy variables for day of week, month, year, and state-by-year; and child fixed effects. 'AC' is imputed air conditioning ownership.

Appendix Figure 1. Relationship between short-run temperature, air conditioning, and cognitive performance



Notes: The solid line shows coefficient estimates of the effect of temperature on the day of the test of cognitive performance, with 95% confidence intervals in dotted lines. The dashed line adds to this the coefficient estimates from temperature interacted with imputed air conditioning ownership. The regression includes indicators for each 2°C, and also controls for precipitation, pressure, wind speed, humidity, dummy variables for day of week, month, year, and state-by-year, and child fixed effects.