Heat and Learning*

Joshua Goodman
Harvard University
Joshua_Goodman@hks.harvard.edu

Michael Hurwitz
College Board
mhurwitz@collegeboard.org

Jisung Park
University of California, Los Angeles
jisungpark@luskin.ucla.edu

Jonathan Smith
Georgia State University
jsmith500@gsu.edu

April 26, 2018

Abstract

We provide the first evidence that cumulative heat exposure inhibits cognitive skill development and that school air conditioning can mitigate this effect. Student fixed effects models using 10 million PSAT-takers show that hotter school days in the year prior to the test reduce learning, with extreme heat being particularly damaging and larger effects for low income and minority students. Weekend and summer heat has little impact and the effect is not explained by pollution or local economic shocks, suggesting heat directly reduces the productivity of learning inputs. New data providing the first measures of school-level air conditioning penetration across the US suggest such infrastructure almost entirely offsets these effects. Without air conditioning, each 1°F increase in school year temperature reduces the amount learned that year by one percent. Our estimates imply that the benefits of school air conditioning likely outweigh the costs in most of the US, particularly given future predicted climate change.

*Thanks to Lawrence Katz, Joseph Aldy, Alan Barreca, Olivier Deschenes, Sarah Reber, Wes Yin, Manisha Shah, Judson Boomhower, Teevrat Garg, Marc Jacobsen, and Patrick Behrer for valuable comments and feedback. We also thank seminar participants at Harvard, UCSB, UCSD and West Point.

1 Introduction

It is well known that hotter countries tend to be poorer, with each 1°F increase in average temperature associated with 4.5% lower GDP per capita (Dell et al., 2009). A fact less well-established is that hotter countries also tend to score lower on internationally comparable measures of academic achievement. The top panel of Figure 1 shows a strong negative relationship between a country's average annual temperature and its standardized math score from the 2012 Programme for International Student Assessment (PISA). Across the 60 countries' nationally representative samples of 15-year-olds that took this exam, each 1°F increase is associated with a decrease in math scores of 0.02 student-level standard deviations. The bottom panel of the figure shows a similar pattern across counties within the United States, with each 1°F increase associated with a decrease in 3rd to 8th grade math scores of 0.004 student-level standard deviations (Fahle et al., 2017). Our primary contribution is to show that some of this cross-sectional relationship between temperature and academic achievement is causal.

Why might cumulative heat exposure reduce human capital accumulation? Heat may indirectly impact learning through its effect on health and disease burden (Schwartz et al., 2004; Bleakley, 2010; Deschênes and Greenstone, 2011), on agricultural income, nutrition, and the opportunity costs of schooling in developing economies (Schlenker et al., 2006; Garg et al., 2017; Shah and Steinberg, 2017), and on economic activity, institutional norms and political stability more broadly (Acemoglu et al., 2001; Dell et al., 2012, 2014; Hsiang et al., 2013). Heat may also impact learning directly by altering human physiology and cognition. Even moderately elevated temperatures can impair decision-making and cause substantial discomfort (Albouy et al., 2016), and short-term impacts of heat on cognition have been extensively documented in laboratory settings (Seppanen et al., 2006). Hot classrooms may thus reduce the effectiveness of instructional time through physiological impacts on both students and teachers.

To estimate the causal impact of cumulative heat exposure on human capital accumulation, we link local daily weather data to test scores of 10 million American students from the high school classes of 2001-14 who took the PSAT exam multiple times. Student fixed effects regressions identify the impact of heat exposure during the prior school year by leveraging within-student variation in temperature over multiple test takes. Our identification strategy relies on the premise that variation in temperature over successive school years for a given student is uncorrelated with unobserved determinants of learning. We provide evidence consistent with that assumption, showing that selection into test-taking and retaking is not endogenous to temperature, even when controlling for regional trends in warming and secular changes in school quality or student composition.

We then generate three primary findings about the impact of heat on human capital accumulation. First, cumulative heat exposure reduces academic achievement. On average, a 1°F hotter school year prior to the exam lowers scores by 0.002 standard deviations or slightly less than one percent of a year's worth of learning.[1] Relative to school days with temperatures in the 60s (°F), each additional school day with temperature in the 90s (°F) reduces achievement by one-sixth of a percent of year's worth of learning. A day above 100 (°F) has an effect that is up to 50 percent larger. These effects are precisely estimated, are robust to controlling for test-day weather, pollution exposure and state-specific time trends, and are not predicted by heat exposure in the year following the test. These mean effects mask large heterogeneity by income and race. For students living in zip codes in the lowest quintile of average income, a 1°F hotter prior school year is three times as damaging to academic achievement as it is for those living in the highest quintile of income. Similarly, the impact of heat on achievement is three times as large for black and Hispanic students as for white students.

Our second result is that heat most likely reduces academic achievement by decreasing the productivity of instructional time and not through other, more indirect channels, at least in developed economies. We show that only school-day exposure to higher temperatures affects test scores; hot summers and weekends have little impact on achievement and controlling for such exposure does not shrink the magnitude of impact of hot school days. This suggests that heat's disruption of instructional or homework time is responsible for the observed drop in test scores. The magnitude of our estimated impacts of prior year heat exposure on achievement are also unchanged when we control for prior year local economic conditions that might be correlated with hot temperature (e.g. agricultural and labor productivity), suggesting that heat's impact on students does not operate through economic impacts on their families in relatively high income countries. We similarly

[1] On average, students score 0.3 standard deviations higher on their second PSAT take than their first PSAT taken one year prior.

show that heat's association with air pollution does not explain the observed effects.

Our third and final result is that school air conditioning appears to offset nearly all of the damaging impacts of cumulative heat exposure on academic achievement. To show this, we construct the first nationwide measures of school-level air conditioning penetration in the United States by surveying students and guidance counselors across the country about heat-related conditions in their schools. We then document that school air conditioning is less prevalent in cooler parts of the US and that heat is particularly damaging to the achievement of students in these regions. In the cross-section, school air conditioning penetration reported in 2016 mitigates the adverse effect of hot temperatures substantially, such that moving from a school with no air-conditioned classrooms to a school with all air-conditioned classrooms reduces the impact by approximately 78 percent. The strongest causal evidence for the mitigating effect of air conditioning comes from a triple-difference strategy that combines our within-student comparisons with within-school variation in air-conditioning status over time. Cohorts in schools that increased air conditioning penetration experienced decreased impacts of heat exposure over time relative to schools that did not increase air conditioning penetration.

These findings contribute to a growing literature on the impact of heat on individual outcomes and cognition more specifically. Acute exposure to heat stress reduces labor supply and productivity (Hsiang, 2010; Graff Zivin and Neidell, 2014; Behrer and Park, 2017) and increases aggressive behavior (Kenrick and MacFarlane, 1986; Hsiang et al., 2013). Recent work has found that heat exposure on the day of the test reduces performance on standardized graduation exams for New York City high schoolers (Park, 2017), on voluntary math assessments for a national sample of 5- to 14-year-old Americans (Graff Zivin et al., 2017), and on cognitive assessments of primary and secondary school students in India (Garg et al., 2017). Cognitive impacts of acute heat stress may, however, be transitory and not represent permanent reductions in the stock of human capital.

Our work shows that cumulative heat exposure can generate such longer-term reductions in human capital accumulation. This is consistent with a longstanding lab-based literature documenting adverse cognitive impacts of hot temperature (Seppanen et al., 2006; Kjellstrom et al., 2016), the long-run implications of which have not yet been credibly tested in actual learning environments. We believe our study is the first to precisely estimate the direct effect of cumulative heat exposure on learning, building on suggestive but noisy evidence from the United States

(Graff Zivin et al., 2017), suggestive estimates of the impact of summer heat on South Korean students' college entrance exams (Cho, 2017), and the impact of growing season heat on Indian test scores (Garg et al., 2017). Our large sample size allows us to more precisely estimate the impact of cumulative heat exposure in the United States context. We also isolate the primary channel through which cumulative heat exposure operates by ruling out explanations such as changes in income driven by agricultural or labor productivity (Garg et al., 2017; Shah and Steinberg, 2017), contemporaneous impacts of heat or pollution on exam performance (Ebenstein et al., 2016; Graff Zivin et al., 2017; Park, 2017; Marcotte, 2017), and indirect health channels (Bleakley, 2010; Cho, 2017). Moreover, we are the first to empirically assess the effect of school air conditioning on learning, and to document heterogeneity by socioeconomic status, which may be important for understanding persistent racial and economic achievement gaps (Hoxby, 2001; Card and Rothstein, 2007; Dobbie and Fryer Jr, 2011; McFarland et al., 2017). The only other paper that precisely identifies the impact of cumulative heat exposure on human capital accumulation in a developed country context is Isen et al. (2017), which focuses on in utero exposure and thus identifies a very different channel from the learning channel we study here.

Our analysis of the mitigating effects of air conditioning is also one of the first to provide evidence on the efficacy of specific investments in educational infrastructure. Recent evidence suggests that school funding may potentially affect student outcomes (Jackson et al., 2015; Lafortune et al., 2016), but even studies that focus on school infrastructure funding are estimating the impact of broad funding packages and not targeted upgrades to specific school facilities (Cellini et al., 2010). The one paper that finds clear, positive achievement impacts of infrastructure spending studies a comprehensive school construction project in New Haven, CT, one major priority of which seems to have been installing air conditioning in schools that lacked it (Neilson and Zimmerman, 2014). The lack of achievement impact from school capital campaigns in Texas is consistent with our data suggesting nearly all schools in that state were likely air conditioned to begin with (Martorell et al., 2016). Ours is the first study to isolate the impact of this specific form of school infrastructure investment, and to document inequality in their adoption across race and income groups.

More broadly, our results point to the importance of accounting for avoidance behavior and defensive investments when estimating both past and future costs of environmental shocks. That air conditioning offsets a substantial proportion of the effect suggests analyses of human capital accumulation that do not account for avoidance behavior may understate effects, even though the underlying biological and welfare impacts are non-trivial. This is particularly true where levels of air conditioning penetration are high, as in much of the United States, or other wealthy developed economies. Our findings are thus consistent with recent work that finds avoidance behavior and defensive investments contribute substantially to the overall welfare costs associated with an environmental shock such as poor air quality or heat (Graff-Zivin et al., 2011; Deschenes et al., 2017).

These findings also suggest substantial scope for adaptation to future climate change. In line with recent work exploring the extent to which air conditioning mitigates the negative effects of heat on mortality and on labor productivity, our results suggest that not taking adaptation into account when estimating the social cost of carbon can be highly misleading (Barreca et al., 2016; Behrer and Park, 2017). Ours is the first study of which we are aware to study the effectiveness of adaptation in the context of human capital accumulation. Finally, our findings highlight the potential distributional consequences of climate change, and suggest that taking avoidance behaviors into account likely exacerbates underlying inequities. Given the correlation between income and temperature noted above, the well-documented links between income and AC ownership, and the potential for Tiebout sorting on local public goods such as school infrastructure, our results suggest that the impacts of additional heat exposure from climate change are likely to be more acute for poorer individuals both across and within countries (Biddle, 2008; Davis and Gertler, 2015; Hallegatte et al., 2018).[2]

Most broadly, evidence that heat exposure affects human capital accumulation points to one potential and understudied channel through which heat may affect macroeconomic growth. A variety of recent papers have documented clear connections between country-level heat and growth but the mechanisms explaining that connection have remained a matter of speculation (Dell et al., 2012; Hsiang et al., 2011; Burke and Emerick, 2016). Hypothesized channels include the negative impacts of heat on physical health and on labor productivity, particularly for physically intensive occupations (Dell et al., 2014; Deschenes, 2014; Heal and Park, 2016). Our evidence suggests that

[2] There is evidence that households in developing countries often face liquidity constraints in the purchasing of energy-intensive appliances (Gertler et al., 2016), which would compound this regressivity.

heat not only interferes with physical capabilities of a nation's workforce but also with its cognitive capacities, and most crucially the rate at which valuable skills are accrued by the workforce over time. In the context of environmental policy, the magnitude of the externality associated with carbon pollution will depend crucially on whether climatic factors affect underlying economic growth rates or simply the level of output in any given year. Existing social cost of carbon estimates may thus be understated given the omission of human capital and growth rate effects from most integrated assessment models (Pindyck, 2013).

We also use our estimates to calculate a variety of previously undocumented quantities, including the contribution of heat to racial achievement gaps in the U.S., the potential impacts of climate change on U.S. student achievement, and the value of school air conditioning. We argue that heat effects account for up to 13 percent of the U.S. racial achievement gap, both because black and Hispanic students live in hotter places than white students and because heat damages minority students' achievement of minority students more than white students' achievement. We show that predicted temperature increases due to global climate change could lower U.S. achievement by 0.1 standard deviations by 2050, assuming no further investment in school air conditioning, but less than 0.05 standard deviations in a world where school air conditioning penetration increases in line with historical trends. Finally, we estimate that school air conditioning would offset over $25,000 per classroom per year in future lost earnings due to temperature increases predicted by climate change models. The magnitude of such benefits appear substantially larger than the costs of installing and operating such infrastructure.

The rest of the paper is organized as follows. Section 2 presents the data and empirical strategy. Sections 3 and 4 describe our estimated impacts of heat on learning and of the mitigating effect of school air conditioning. Section 5 presents policy-relevant calculations based on such estimates. Section 6 concludes.

2 Data and Empirical Strategy

2.1 Data

We combine three primary data sources on test scores, temperature and school air conditioning. Test score data come from the College Board, which administers the PSAT exam to millions of

American high school students annually. The PSAT consists primarily of a reading and a math section.[3] The test is offered once a year roughly during the the third week of October and most students first take it in 10th or 11th grade, though some start as early as 9th grade.[4] Students take and can retake the PSAT for a variety of reasons, including preparation for the SAT college entrance exam, qualification for National Merit Scholarships, and provision of information about their college readiness to them and their schools. The PSAT has multiple advantages over other tests used to study the impact of heat on cognitive skills, including: the test is given once a year on a fixed date with advanced registration required, limiting the scope of endogenous taking or timing decisions; proctoring is nationally harmonized and the test is centrally graded, limiting potential for endogenous score manipulation of any sort (Hallegatte et al., 2018); and the test is designed to assess cumulative progress on skills learned during secondary school, rather than generalized intelligence, making it arguably better-suited for assessing the effects of formal schooling. We have test scores and dates from all takes for the universe of PSAT-takers from the intended high school classes of 2001-2014. Our primary outcome is a student's combined math and reading score, standardized by test administration. The student-level data also contains basic demographic information such as gender, race, parental education, and residential zip code, which we use to assign students zip code level mean incomes as reported by the Internal Revenue Service. High school identifiers allow us to link students to testing locations and to other school-level characteristics.

Daily temperature data come from the National Oceanic and Atmospheric Administration's Daily Global Historical Climatology Network, which includes station-level data for thousands of weather stations across the United States. We focus on the subset of nearly 3,000 weather stations with daily temperature data available for at least 95 percent of the days from 1996 to 2014, the time period covering potential test-taking dates of our sample. Doing so allows us to assign each high school to a single, stable weather station over the entire time period, which avoids endogeneity concerns driven by the possibility that stations coming online or going offline are somehow corre-

[3] A writing section has been added in more recent years. Basic scientific concepts and history are assessed as part of the reading comprehension section.

[4] We use the PSAT rather than the SAT for two reasons. First, the SAT is offered at seven different times of the year, making it harder to assign easily comparable measures of long-term heat exposure to a given exam take. Second, the PSAT is taken by roughly twice as many students as the SAT because the latter is taken in later grades by a more college-oriented and thus selected set of takers.

lated with local population growth, economic conditions or temperatures conditions in ways that might contaminate our estimates (Auffhammer and Mansur, 2014). We impute the small proportion of missing daily observations with those from the nearest stations with non-missing data. We assign each high school to the nearest weather station, resulting in an average distance of 9.7 miles between a student's test site and weather station being used to measure temperature at that site.

We construct two primary measures of cumulative heat exposure experienced by a student: the average daily maximum temperature and the number of days that temperature exceeded a given multiple of 10°F in the 365 days prior to the test.[5] We use daily maximum temperature because schooling occurs during the daytime when such temperatures usually occur.[6] We focus particularly on temperature experienced on school days, treating summer and weekend temperatures as separate sources of variation.[7] We also use the weather stations to construct test date temperature, rain and snowfall, as well as cumulative rain and snowfall exposure over the year prior to the test, which help account for potential independent effects of such precipitation (Goodman, 2014).

School-level air conditioning data do not exist at the national level and very rarely exist at the state and local level. We generate measures of school-level air conditioning penetration through a survey the College Board regularly administers by e-mail to all SAT takers and to high school guidance counselors registered to administer PSAT or other exams. In 2016, we added to the survey the statement "On hot days, classrooms get too hot." Respondents could choose "Never", "Some of the time", "Most of the time", or "All the time".[8] We received valid responses from students in nearly 12,000 schools enrolling 87 percent of our sample's PSAT-takers. To construct school-level measures from individual responses, we assign the four possible responses a value of 0, $\frac{1}{3}$, $\frac{2}{3}$, and 1, then average across all students within a school. Though imperfect, we interpret

[5] Focusing on the prior year implies that the measured heat exposure occurred prior to a given PSAT administration but after the most recent one before that. Because of slight annual variation in the timing of the PSAT, we exclude the third week of October from these measures to guarantee no overlap between PSAT administrations and the constructed measure of heat exposure. The year prior to the test take thus runs from late October of the preceding year to mid-October of the current year.

[6] We note that it is possible for nighttime heat to affect learning through disrupted sleep as well. To the extent that daytime high and nighttime low temperatures are correlated, it is possible that our estimates may include some impacts due to disrupted sleep, though the results on weekend days versus weekdays is suggestive of effects driven primarily by instructional time, as we discuss below.

[7] No comprehensive national data set of school calendars covering this time period exists, so we assign to each student a likely school start and end date based on the calendar of the largest school district in that student's state, as seen in Figure A.1. We then divide the year into three periods: school days, the summer, and weekends or national holidays.

[8] Respondents were also allowed to choose "I don't know." We coded such answers as missing.

this measure as the fraction of a school's classrooms with adequately functioning or sufficiently frequently activated air conditioning. This measure of school air conditioning penetration has the advantage of being reported by the students themselves and of being based on the largest set of responses to our questions about air conditioning. It has the disadvantage of measuring air conditioning penetration at a single point in time and thus may correlate with other school-level factors we can not adequately control for in our analyses.

A different measure of school air conditioning penetration comes from our addition to the survey of two questions posed just to high school guidance counselors: "How many of your school's classrooms have air conditioning?" and "Ten years ago, how many of your school's classrooms had air conditioning?" Guidance counselors could respond with "None", "Fewer than half", "About half", "More than half", and "All". To construct school-level measures, we assign these possible responses a value of 0, $\frac{1}{4}$, $\frac{1}{2}$, $\frac{3}{4}$, and 1, then average across all guidance counselors within a school. This measure has the disadvantage that we received responses from guidance counselors in only about 2,000 high schools. It has the advantage, however, of directly asking about air conditioning and of allowing us to measure air conditioning penetration at two points in time instead of one. Variation in air conditioning penetration over time is more plausibly exogenous than cross-sectional measures, helping us construct stronger causal claims about the impact of school air conditioning in offsetting heat's impacts, though it is possible for AC upgrades to be correlated with other school improvements that protect against hot temperature.

Finally, we supplement our three primary data sources with additional data on residential air conditioning, local economic conditions and pollution levels. We construct county-by-year data on residential air conditioning penetration by combining county-level residential air conditioning penetration estimates from the 1980 decennial census with data on changes in such penetration over time by census region from the Energy Information Agencies Residential Energy Consumption surveys.[9] We estimate county-level economic conditions by constructing the logarithm of annual payroll per capita from the U.S. Census Bureau's County Business Patterns, focusing on sectors that the National Institute of Occupational Safety and Health categorizes as being highly

[9]We use the reported penetration rates in 1980 as a basis and then extrapolate based on the region-level growth rate of total air conditioning penetration recorded by RECS, which provide penetration rates by region from 1980 to 2009 with a two or three-year frequency. We linearly interpolate growth rates for the missing years and assign counties their corresponding regional growth rate. Using this growth rate and the observed penetration rate in 1980 we create a measure of penetration in every county in each year from 1980 to 2011. We top-code penetration at 100 percent.

exposed to weather (namely: construction, mining, transportation, manufacturing, agriculture and utilities). Based on evidence of the adverse impact of hot days on highly exposed sector payroll, we use this measure to control for local economic shocks driven by annual fluctuations in heat (Behrer and Park, 2017). Similarly to our temperature measures, we construct both cumulative and test date measures of exposure to major air pollutants (ozone, sulfur dioxide, nitrogen dioxide, carbon monoxide, and PM10 particulate matter) using data from the Environmental Protection Agency's Ambient Air Monitoring network. Previous research suggests that such pollutants can affect student absenteeism and performance, at least in the short run (Currie et al., 2009; Ebenstein et al., 2016), though there is relatively less evidence regarding the impact of longer-term exposures.

2.2 Summary Statistics

As shown in Table 1, the starting sample comprises over 38 million test scores from 27 million students in the high school classes of 2001 to 2014, who took the PSAT as early as October of 1998 and as late as October of 2012.[10] Because our primary identification strategy relies on within-student variation in heat exposure for identification, we focus on the subset of 21 million scores from the nearly 10 million students who took the PSAT at least twice.[11] Those retakers are slightly more advantaged than the general pool of PSAT takers, living in higher income zip codes, more likely to have mothers with college degrees, and with first PSAT scores 0.14 standard deviations about the national average. Importantly, retakers seem geographically similar to the general pool of SAT takers. Both groups experience school days with mean temperature in the mid 60s°F in the year prior to the test, with an average of 12 school days above 90°F. Maps of takers and retakers locations also suggest that both groups have similar geographic distributions in line with the population distribution across the United States.[12] On average, PSAT takers and retakers report that nearly 60 percent of their classrooms are adequately air conditioned, with closer to 80 percent likely to have air conditioning at home. Dividing the retaking sample by race and by income reveals that black and Hispanic students and students in the lowest quintile of zip code income

[10] We exclude a very small number of observations of PSATs taken during October of 12th grade.

[11] In our sample, 85 percent of students take the PSAT twice and 15 percent take it three times, the maximum given testing opportunities in 9th, 10th and 11th grades.

[12] See Figure A.2.

score lower on the PSAT and experience hotter school years than their white and higher income counterparts.

The raw data suggest a strong negative relationship between cumulative heat exposure and academic achievement. Figure 2 maps average PSAT performance by county. On average, Southern counties have substantially lower test scores than do Northern counties. This tracks closely with differences in heat exposure by geography, as seen in Figure 3. In Florida and southern Texas, the average school day experienced by a student is above 80°F, compared to an average school day temperature in the 50s°F in much of the Northern U.S. Southern students in many counties experience 30 or more school days above 90°F, compared to fewer than 10 such days in nearly all Northern counties. The existence of such a strong North-South temperature gradient and test score gradient could be causal or could be driven by other important regional differences. Our goal is disentangle the impact of heat from such other factors.

2.3 Empirical Approach

To do so, we exploit the fact that for students who take the PSAT multiple times, differences across takes in prior year heat exposure are likely uncorrelated with differences in other factors that might affect academic achievement. We thus ask whether students score lower immediately following a hotter year relative to their own score immediately following a cooler school year and, if so, we argue that heat is the only factor that can explain such a difference in outcomes. The main identification assumption is that unobserved determinants of student performance are uncorrelated with year-to-year variation in weather. We implement this identification strategy with student fixed effects regressions of the following form:

$$Score_{iscyn} = \beta Heat_{sy} + \eta_i + \gamma_{cyn} + \epsilon_{isycn} \tag{1}$$

Here, $Score$ denotes the standardized PSAT math and reading score for student i in high school s, high school class c, taking the test in October of year y for the nth time. Inclusion of student fixed effects η implies that identification comes from within-student comparisons of heat exposure and test score differences over multiple takes.[13] High school class by test year by take number

[13]We use student fixed effects rather than high school fixed effects because the latter approach depends on the assumption that selection into PSAT-taking at the school level does not vary over time in ways correlated with heat

fixed effects γ flexibly control for a variety of potential confounds, including differential selection into test-taking across high school classes, differential difficulty of the test across test dates, and differential test performance based on past number of test takes. We cluster standard errors by weather station, the level of variation in our treatment variable.

We most frequently define cumulative heat exposure $Heat$ as the average maximum temperature experienced during school days in the year prior to the test, for all students in high school s taking the test in year y. In that case, the coefficient of interest β can be interpreted as the standard deviation impact on a student's test score of experiencing a one degree F hotter school year on average. We also use a specification that replaces this single heat exposure measure with a vector of counts of the number of school days falling into various temperature bins. One example of such a specification uses four temperature bins: the number of school days below 60, in the 70s, in the 80s, and 90 or above. In this specification, the coefficient on days of 90 or above can be interpreted as the impact of experiencing one more very hot school day, relative to a school day with temperature in the 60s.[14] Identification therefore relies on both spatial variation in heat exposure, as previously shown in Figure 3, as well as temporal variation in heat exposure.[15] To understand the magnitude of such identifying variation, we compute the residual from regressing heat exposure on the aforementioned student and class-year-take fixed effects. The distribution of such residuals suggests that a one standard deviation increase in mean school day temperature is about 1°F, while a one standard deviation increase in the number of schools day above 90°F is about three days.

One potential threat to identification comes from the possibility that cumulative heat exposure drives selection into taking the PSAT for the first time or choosing to retake it. To test for selection into taking the PSAT the first time, we collapse the data by high school and regress the number of first-time test-takers (as well as its logarithm) on high school fixed effects and cohort by test date fixed effects. We see no evidence that within-school fluctuations in annual heat exposure affects the number of test takers and can rule out economically meaningful effect sizes. We then run

exposure. This assumption fails empirically because, over the time period in question, PSAT taking expands to a wider set of students and moreso in regions of the country that are differentially affected by longer-term warming trends. The student fixed effects approach avoids this selection margin entirely.

[14]We do not find strong evidence that cold weather affects learning in our sample, hence our focus on the upper end of the temperature distribution.

[15]See Figure A.3 for annual variation in average school day temperatures.

similar student-level regressions using the demographic characteristics of first-time test-takers as outcomes. The results rule out meaningful impacts of heat exposure on the observable composition of the test-taking population, particularly when controlling for differential heat trends by state.[16] Finally, we ask whether heat in the year prior to the first take or in the year following the first take affects the probability that a student retakes the PSAT. We again find no evidence of such selection, with point estimates suggesting a 1°F hotter school year increases the probability of retaking by 0.05 percentage points and confidence intervals that rule out effects larger than 0.15 percentage points.[17] In total, these results suggest little evidence of endogenous selection into test-taking or retaking as a result of cumulative heat exposure.

3 The Impact of Cumulative Heat Exposure

3.1 Mean Impacts

On average in the U.S., experiencing a 1°F hotter school year lowers academic achievement by 0.002 standard deviations, a result that is very precisely estimated and robust to a variety of controls for potential confounding factors. Table 2 shows these results. The first column of panel A contains the baseline specification described in equation 1, where the test score outcome is measured in hundredths of a standard deviation. The coefficient is highly statistically significant and precise enough to rule out effects smaller than 0.001 standard deviations. The magnitude of these impacts is small enough to have been missed by previous studies with less precision but large enough to imply non-trivial cumulative effects of temperature on learning. For example, the average gain in PSAT score performance between 10th and 11th grade is about 0.3 standard deviations, implying that a 1°F hotter school year reduces learning by close to one percent of the expected gains over that year. As we will see below, learning impacts are even larger for certain demographic subgroups and in the absence of air conditioning.

That students score lower following hotter years relative to their own scores following cooler years does not appear to be driven by other channels potentially correlated with heat in the school year leading up the exam, as seen in the next four columns of Table 2. Controlling for total snow-

[16] See Table A.1 for detailed estimates.

[17] See Table A.2 for detailed estimates.

fall and rainfall in the prior year and for temperature and precipitation on the day of the exam has nearly no effect on the point estimate. This suggests that we are not mistakenly attributing to cumulative heat exposure effects that are actually driven by cumulative precipitation exposure or by contemporaneous heat exposure. Similarly, controlling for both cumulative and contemporaneous pollution exposure leaves our estimate nearly unchanged, implying that we are measuring the direct impact of heat and not of pollutants such as ozone that may be more common on hot days.

Controlling for county-level payroll in industries highly exposed to weather conditions also does little to our point estimate, suggesting that cumulative heat exposure is not operating through the channel of family income or local economic conditions.[18] The robustness of our estimate to controls for state-specific time trends suggests we are not picking up spurious correlations driven by subtle geographic differences in warming trends that may be correlated with other local changes in selection into or preparation for PSAT-taking or retaking.[19] Regardless of which of the aforementioned controls are included, the estimated impact of a 1°F hotter school year, which represents a roughly one standard deviation change in cumulative heat exposure, is never substantially different from 0.002 standard deviations.

Given that the mean distance between weather stations and high schools in our data is a little less than 10 miles, the cumulative heat exposure we assign to each student may be mismeasured, particularly for students farthest away from weather stations. The final column of Table 2 limits the sample to high schools within five miles of a weather station, for which we assume that measurement error is less of an issue. For students whose high schools are particularly close to weather stations, the impact of cumulative heat exposure on academic achievement is about 25 percent larger than for the sample as a whole. This is consistent with the possibility that measurement error in our treatment variable is generating downward bias in our estimates, though it is also consistent with the possibility of heterogeneous treatment effects correlated with factors that make weather stations more likely to be online near schools.

In addition to estimating the impact of generally hotter school years, we also show that very hot days are particularly damaging to student achievement. Panel B of Table 2 shows the specifica-

[18]This is unsurprising given the developed country context but contrasts with research in developing countries that shows agricultural yield shocks driving schooling outcomes (Garg et al., 2017; Shah and Steinberg, 2017).

[19]The estimates are also robust to using quadratic or cubic trends instead of linear trends.

tion in which we measure cumulative heat exposure by counts of school days falling into various temperature bins. Replacing a school day in the 60s with a hotter school day lowers achievement, with the extent of that damage increasing roughly linearly with temperature above 70°F. Consistent with our baseline specification, these estimated magnitudes also imply that a one standard deviation increase in heat exposure, or over three additional days above 90°F, lower achievement by 0.002 standard deviations. Cold days, those below 60°F, appear to have little impact on achievement. Figure 4 shows these point estimates with further disaggregation of colder school days and, consistent with laboratory studies on cognition and recent studies on labor supply and mortality, shows damage that increases with temperature starting around 70°F. Estimated impact of hot days are, like those for mean heat, robust to controls for cumulative precipitation and pollution, test day weather and pollution, local economic conditions and state-specific trends.

Heat's cumulative impact on academic achievement is not driven by one particular subject, in contrast to findings focusing on short-run cognitive impacts (Graff Zivin et al., 2017; Garg et al., 2017). Both math and reading scores drop by similar magnitudes for a given level of additional heat exposure.[20] Heat's impact is also not driven by one particular test take. Heat prior to a first test take has relatively similar negative effects on achievement as heat prior to second take, although there is some evidence that heat's learning impact, if anything, increases with take number.[21] This appears to rule out the possibility that our results are driven somehow by differential selection into retaking based on correlations between heat exposure, first take performance and students' belief about whether their first scores reflect their true abilities. Using future temperature shocks as a placebo test also yields results consistent with our interpretation of these impacts as causal. Controlling for mean school day temperatures in the years one, two and three years after the exam does little to change our estimated impact of cumulative heat exposure and the coefficients on future temperature are much closer to zero than our main effect and never statistically significant.[22] This makes it less likely our results are a statistical artifact driven by spurious correlations between temperature and test scores.

[20]See Table A.3 for details.
[21]See Table A.3 for details.
[22]See Table A.5 for details.

3.2 Productivity of Instructional Time

One indication that cumulative heat exposure affects achievement by directly lowering the productivity of instructional time comes from examining the impact of heat on three mutually exclusive sets of days in the year prior to the test: school days, weekends and national holidays during the school year, and summer days. The first two columns of Table 3 contrast our baseline specification to one in which we control for heat on summer days. Two important facts emerge. First, controlling for summer heat has little effect on our estimated impacts of school day heat. Second, the impact of summer heat on academic achievement is very small and statistically indistinguishable from zero. That summer heat has no impact on academic achievement seems to rule out potential channels such as student health or local economic conditions given that such channels should be affected by summer heat as well as school day heat. Controlling for school year weekend and holiday heat somewhat increases our main effects and we see no evidence that such heat lowers student achievement, again suggesting that the role of time in school is critical to understanding the relationship between heat exposure and human capital accumulation.[23]

One further test yields results consistent with our interpretation that heat exposure interferes with actual learning. In column 4 of Table 3, we break heat exposure into three time periods corresponding to distance in time to the test: post-summer school days (the roughly two months just prior to the PSAT), summer days (roughly three to five months prior to the PSAT), and pre-summer days (roughly six to twelve months prior to the PSAT). Summer heat again has little clear impact on achievement while both pre- and post-summer heat have large negative impacts. That the damage from pre-summer heat is as large as or larger than the impact of post-summer heat suggests that heat operates not just through periods where students might be "cramming" for a test. Instead, heat appears to affect cumulative learning over a longer time frame.

Finally, we test whether heat exposure more than one year prior to the test affects achievement. In column 5 of Table 3, we include controls for heat exposure in school years two and three years prior to taking the test. When measuring heat exposure by average school day temperature, including those lags has no impact on our main estimate and the lags themselves do not appear

[23] The impact of weekend and holiday heat is, if anything, slightly positive. The pattern of coefficients suggests that this may be driven by strong correlations between school day and weekend heat. There is more independent variation between school year and summer heat than there is between weekday and weekend heat within the school year.

to affect achievement. This is consistent with the possibility that either the PSAT tests material taught very close in time to the test (unlikely given test design) or that some of the negative impacts of much earlier heat exposure can be partially mitigated by students and teachers in the intervening time. However, when measuring heat exposure by school days above 90°F, including lags increases our main estimate by 40 percent and the lags themselves are large, negative and statistically significant. This is consistent with the possibility that heat exposure has even longer-lasting impacts than a year and that failing to control for dynamic effects causes our student fixed effects approach to understate the impact of heat (Cunha and Heckman, 2007).[24]

3.3 Heterogeneity by Income and Geography

Cumulative heat exposure has very heterogeneous impacts by various measures of student socioeconomic status, which we show in Table 4 by splitting the sample by student race, high school racial composition, and zip code income.[25] The impact of heat on black and Hispanic students is three times as large as the impact on white students. That disparity is even starker when we categorize students by the fraction of PSAT-takers in their high school who are black or Hispanic. Students in schools with the highest fraction of minority test-takers see impacts of mean temperature five times larger than those with the lowest fraction of minority test-takers. Particularly hot days have no clear impact on students in such low minority schools. A similar pattern appears when we split the sample into students living in zip codes in the lowest and highest quintile of income. The impact of heat on students in the lowest income zip codes is between two and three times the impact on those from the highest income zip codes.

That the achievement of lower income and minority students suffers much more from heat exposure is consistent with poorer households having fewer margins of effective adaptation, ex ante or ex post (Graff Zivin and Neidell, 2013; Kahn, 2016). Wealthier students may be able to compensate for lost learning time by getting additional instruction from their parents or private tutors. Such students may also be more likely to attend schools where teachers have better capacity to compensate for lost learning time by adjusting lesson plans or adding more instructional time to

[24]The student fixed effects approach implicitly assumes complete decay of the effects of heat after the one year period, so that heat before the first take affects only the first take and not the second take itself.

[25]We find no evidence of heterogeneity by student gender, in contrast to some previous work (Ebenstein et al., 2016).

the day.[26] An even simpler explanation for the heterogeneity is that lower income students have less access to school and home air conditioning that might help offset the negative impacts of heat.

We see indirect evidence that air conditioning may play an important role in the last two columns of Table 4, where we divide students by whether their high school's average school year temperature is below or above the median such temperature across all students. Students in cooler areas suffer slightly more from fluctuations in average heat and more than three times as much from particularly hot days. The impact of particularly hot days on students in the hottest areas is small and only marginally statistically significant. This suggests that hotter areas of the country may adapt to heat exposure through strategies like air conditioning adoption that help mitigate the impacts of heat. We now turn to a more direct test of that hypothesis.

4 School Air Conditioning as an Adaptive Response

4.1 Descriptive Analysis

Adaptive responses are important for understanding the welfare implications of environmental shocks, particularly in the long run (Graff-Zivin et al., 2011; Graff Zivin and Neidell, 2013; Kahn, 2016; Deschenes et al., 2017). Responses to pollution are often referred to as "avoidance behaviors" or "defensive investments", while responses to climate variation are more frequently called "adaptations". School air conditioning represents one potential adaptation technology.[27] To explore its role in mitigating the effects of heat exposure, we first provide descriptive analysis of the prevalence of school air conditioning across the U.S.

Figure 5 shows county-level averages of school air conditioning penetration as measured by the extent to which students (panel A) and guidance counselors (panel B) report that "On hot days, classrooms get too hot." The resulting map is roughly the inverse of an average temperature map. Students and counselors are much less likely to report hot classrooms in the hotter regions of the

[26] Both research and media reports suggest teachers are aware of the adverse impacts of heat on student performance and make efforts to offset some of those impacts ex post. Park (2017) finds that New York City teachers selectively boosted grades of students who experienced hot exam sittings and scored just below pass/fail cutoffs.

[27] Teachers and parents seem to highly value school air conditioning, judging by recent labor disputes and community petitions. During a major teacher strike in Chicago in 2012, "Timetable for air conditioning" was one of four major contract demands, with an agreement to provide universal air conditioning eventually reached in 2016. Parents and teachers in a number of major school districts such as New York City, Los Angeles, and Denver have have recently signed petitions asking districts to upgrade air conditioning equipment. See: http://www.denverpost.com/2011/09/08/heat-related-illnesses-spur-petition-for-sept-school-start-in-denver/.

country and much more likely to report hot classrooms in cooler regions. Students in the Northeast, for example, report that heat interferes with learning on the majority of hot days. Students in the South report heat interfering with learning on only about one-fourth of hot days.

Because this reporting may partly reflect the extent to which students and counselors are accustomed to heat, rather than actual air conditioning status, we ask guidance counselors directly about the fraction of classrooms with air conditioning. That map, in panel A of Figure 6, looks quite similar to the map of hot classrooms. According to guidance counselors, nearly all classrooms in the South have air conditioning while the majority of classrooms in the Northeast lack it. We therefore interpret student reports of hot classrooms as measure of air conditioning penetration. These various measures are the first we know of to provide national school-level estimates of air conditioning status. We also note, as seen in panel B of Figure 6, that both home and school AC seem to vary substantially by region, with lower penetration rates particularly in more mountainous regions of the country.

We next document racial and income gaps in school air conditioning penetration rates. To do so, we regress the student-generated measure of school air conditioning on race and zip code income quintile indicators, controlling for a quartic in school-level mean school year temperature experienced over the sample time period. As Table 5 shows, students in predominantly black and Hispanic high schools report 3.5 percentage point lower school air conditioning penetration rates relative to students in high schools with the fewest black and Hispanic students, a gap that does not vary by location. Students living in the lowest income zip codes report five percentage points lower school air conditioning penetration rates relative to those living in the highest income zip codes, a gap that is larger in cooler areas than in hotter areas. Figure 7 further shows that racial gaps in air conditioning access are not explained purely by income differences. If school air conditioning does mitigate the impacts of heat, differential access to that technology may partly explain the observation that heat's impacts are much greater for racial minorities and low income students.

4.2 Causal Analysis

We use two approaches to show that school air conditioning offsets heat's adverse impact on learning. First, we interact the cross-sectional measure of air conditioning penetration reported by students with our heat exposure measure and add it to our baseline specification. That regression has the form:

$$Score_{iscyn} = \alpha Heat_{sy} + \beta Heat_{sy} * SchoolAC_s + \eta_i + \gamma_{cyn} + \epsilon_{isycn} \tag{2}$$

The coefficient α now has the interpretation of the impact of heat on a school with no air conditioning, while β represents the predicted difference between that impact and the impact on a fully air conditioned school.

This analysis suggests that school air conditioning almost fully offsets the impacts of cumulative heat exposure. In column 1 of Table 6, the main coefficient implies that, for students in schools with no air conditioning, a 1°F hotter school year lowers achievement by 0.0032 standard deviations. The interaction coefficient suggests that, in fully air conditioned schools, this effect is 0.0025 standard deviations smaller. For the average student, school air conditioning thus appears to offset 73 percent of the learning impact of hot school days. That interaction coefficient may represent the causal impact of school air conditioning but it may also be picking up effects of other factors correlated in the cross-section with a school's air conditioning status.

To test this, we add to equation 2 four more terms in which heat exposure is interacted with county-level home air conditioning rates, as well as mean school year temperature, zip code income, and the racial composition of a school's PSAT takers.[28] The results of this augmented specification are shown in column 2 of Table 6. The topmost coefficient now suggests that, for a student with neither school nor home air conditioning, a 1°F hotter school year lowers achievement by 0.0057 standard deviations. For that student, school air conditioning and home air conditioning respectively offset 40 and 57 percent of this effect, implying that a student with access to both would see little negative impact of heat exposure. That the magnitude of the school air conditioning coefficient does not change substantially with the addition of these controls implies that

[28]We use sample-demeaned versions of the latter three variables so that coefficients can be interpreted as impacts on schools with average temperatures, income and racial composition.

omitted variable bias from such sources is unlikely to explain the observed effect.

We make one further attempt to isolate the impact of school air conditioning. To do so, we use the change over time in penetration rates implied by differences in counselors' reports about their schools' air conditioning status in 2016 versus 10 years before that. We assign to students a variable $SchoolACchange$, which represents their high school cohort's change in air conditioning penetration rate relative to 2006 implied by the counselors' answers. Cohorts from 2006 and earlier are assigned a value of zero and more recent cohorts are assigned a change linearly interpolated from the counselors two responses. We then run the following specification:

$$\begin{aligned} Score_{iscyn} = \beta Heat_{sy} * SchoolACchange_{sc} * HSClass_i + \delta Heat_{sy} * SchoolACchange_{sc} \\ + \mu Heat_{sy} * HSClass_i + \nu Heat_{sy} + \eta_i + \gamma_{cyn} + \epsilon_{isycn} \end{aligned} \quad (3)$$

Here, $HSClass$ is a continuous measure of a student's cohort. In effect, the coefficient β from this triple-difference approach estimates whether schools that have installed additional air conditioning over time see the impact of heat shrink across cohorts, relative to schools that have not added air conditioning.

The results in column 3 of Table 6 suggest that this is the case. The coefficient on the triple interaction term is positive and highly statistically significant, implying that later cohorts do see smaller impacts of heat in schools that improved air conditioning penetration, relative to schools that did not. That conclusion is unchanged when we control for additional interactions with home air conditioning, local temperature, local income, and school racial composition. This provides an additional piece of compelling evidence that school air conditioning itself is mitigating the impacts of heat exposure. The main threat to validity here is that a school's adoption of air conditioning correlates with other unobserved changes over time in that school that might also mitigate the impacts of heat. Though possible, it seems likely that changes in school air conditioning penetration are more exogenous than cross-sectional variation in such penetration. That both approaches yield consistent results suggests school air conditioning mitigates a very substantial portion of the learning impacts of heat exposure.

5 Policy Simulations

In this section, we use our estimates to calculate the contribution of heat to racial achievement gaps in the U.S., the potential impacts of climate change on U.S. student achievement, and the implied value of school air conditioning.

5.1 Heat and Racial Achievement Gaps

Understanding the factors that contribute to racial achievement gaps is of particular importance in the context of U.S. education policy. We estimate the contribution of heat to such achievement gaps as follows. We first use our data to estimate that black and Hispanic students experience, on average, school years that are 5°F hotter (with 9 additional days above 90°F) than white students typically experience. The estimates from Table 4 suggest that each 1°F hotter school year lowers black and Hispanic students' achievement by 0.002 standard deviations more than white students' achievement. Multiplying this by the 5°F difference between minority and white students implies that, in a single school year, the differential heat experienced by these two groups lower minority achievement by 0.01 standard deviations relative to white students. If such heat effects accumulate over the roughly 10 years of schooling leading to the PSAT, heat could thus be lowering black and Hispanic students' achievement by 0.1 standard deviations relative to white students. Finally, our data show that black and Hispanic students score, on average, 0.8 standard deviations lower on the PSAT than do white students. These estimates thus imply that heat effects account for between one percent (0.01 standard deviations) and 13 percent (0.1 standard deviations) of the U.S. racial achievement gap.

5.2 Estimating Potential Impacts due to Climate Change

Based on current estimates of projected warming in the U.S., we engage in the following thought experiment: By 2050, how much heat-related learning disruption can we expect for the average high school student, relative to a student attending high school in 2010? Median climate change scenarios for the contiguous United States predict average warming of roughly 5°F by 2050. To generate a summary measure of the impact of climate change on future learning, we take the average treatment effect in degree F terms from above (0.002 standard deviations) and multiply by the extent of predicted mid-century warming (5°F), yielding an estimate of 0.01 standard deviation lower achievement. If the impact of heat is cumulative over the roughly 10 years of schooling leading to the PSAT, climate change would lower achievement by 0.1 standard deviations by 2050.

These effects would be larger for students in the Northeast and other cooler parts of the US, where the per degree impact is larger today. These calculations assume that both school and home air conditioning rates do not change and do not account for potential non-linearities in the damage function for temperatures outside the range of historical experience, such as days above 110°F. If we assume "secular adaptation", whereby the observed cross-sectional pattern of air conditioning penetration reflects the expected future time-path of adaptation as places warm, the estimated impact of heat on learning would be slightly less than half of our estimates that assume no adaptation. If all schools have air conditioning in all classrooms, the effect is less than one quarter of the no adaptation estimates. It is important to note that these back-of-the-envelope estimates do not account for heterogeneity in warming or air conditioning penetration patterns across regions, or for the costs of installing air conditioning, both of which are likely to be important determinants of net welfare impacts.

5.3 The Value of Air Conditioning

To monetize the value of air conditioning, we apply previous estimates of the relationship between test scores and later life earnings from Chetty et al. (2011). That paper finds that having a teacher who raises test scores by 0.1 standard deviations brings a net present value of $8,500 in future increased earnings for current 16 year olds.[29] Our estimates in Table 6 suggest that school air conditioning offsets about 0.0025 standard deviations in learning damage for each 1°F increase in temperature. This translates into a net present value of $212 in recovered future earnings per student, per 1°F increase in temperature, and per year that heat is potentially damaging learning. In a city like Houston, where the average school day is approximately 80°F and thus 10°F above the point where heat's impact on learning first appears, the present value of school air con-

[29] Chetty et al. (2011) compute a $7,000 net present value in increased earnings for the typical 12 year old student in their sample. We apply their five percent discount rate to generate the $8,500 figure for the typical 16 year old student taking the PSAT. An important assumption we make by using these estimates is that the achievement gains due to a better teacher result in later life impacts that are equivalent in magnitude to those associated with having a climate-controlled learning environment. For instance, if teachers impart valuable non-cognitive skills but air conditioning does not, these estimates would be overstated.

ditioning is approximately \$2,120 per year for each student, \$53,000 per year for each 25 student classroom, or \$2.1 million per year for each 1,000 student high school. Put differently, the extent to which school air conditioning would offset the earnings loss driven by the 5°F increase predicted by climate change models is \$1,060 per student, \$26,500 per classroom, or just over \$1 million per high school. Given the effect heterogeneity documented above, the estimated present value could be substantially larger for minority or low income students. Although these are rough estimates, benefit values of this order of magnitude imply that school infrastructure improvements may more than justify their costs. For example, in 2017, New York City public schools allocated \$28 million to install air conditioning in 11,000 classrooms, which comes to approximately \$2,500 per classroom.

6 Conclusion

We provide the first evidence that cumulative heat exposure inhibits cognitive skill development and that school air conditioning can mitigate this effect. Student fixed effects models using 10 million PSAT-takers show that hotter school days in the year prior to the test reduce learning, with extreme heat being particularly damaging and larger effects for low income and minority students. Weekend and summer heat has little impact and the effect is not explained by pollution or local economic shocks, suggesting heat directly reduces instructional time productivity. New data providing the first measures of school-level air conditioning penetration across the US suggest such infrastructure almost entirely offsets these effects. Without air conditioning, each 1°F increase in school year temperature reduces the amount learned that year by one percent.

These findings contribute to a growing literature on the impact of heat on individual outcomes by showing that cumulative heat exposure generates long-term reductions in human capital accumulation. We are also the first to isolate the impact of one specific form of school infrastructure investment, namely school air conditioning. Consistent with recent research in the context of environmental determinants of health (Barreca et al., 2016; Deschenes et al., 2017), our analyses suggest that accounting for avoidance behavior and defensive investments when estimating both the costs of environmental shocks is quantitatively important. Ours is the first study of which we are aware to estimate the effectiveness of such adaptation in the context of human capital accumulation, a particularly important fact in the context of likely future climate change.

Understanding the causal relationship between cumulative heat exposure and learning is of heightened policy relevance in light of accelerating warming in most parts of the world, and given that the overwhelming majority of the world's population does not yet have access to air conditioning (Chambwera et al., 2014; Davis and Gertler, 2015). In particular, these findings have important implications for the ongoing policy discussion over the benefits and costs of reducing greenhouse gas emissions. The true social cost of carbon depends in part on whether a hotter climate affects only levels of contemporaneous economic activity or instead the growth rates of macroeconomies (Pindyck, 2013). It also depends on the extent to which economic agents can adapt to a hotter climate and, relatedly, whether the effects of climate documented in lower income countries can be extrapolated into a future when such economies will be increasingly industrialized (Kahn, 2016). An ongoing debate spurred in large part by Dell et al. (2012) remains inconclusive on both of these issues. We believe that this study provides important empirical evidence on these two dimensions.

Our findings suggest that, even in highly industrialized economies, heat exposure can reduce the rate of learning and skill formation, thus potentially reducing the rate of economic growth. We take these findings to be complementary to the growing literature based in developing economies which find substantially larger indirect effects on cognitive skills operating through agricultural and health channels Cho (2017); Garg et al. (2017); Shah and Steinberg (2017). Based on median climate change projections for the United States, our estimates suggest that students in 2050 could be 0.01-0.1 standard deviations less cognitively skilled due to hotter temperatures. This implies that climate impacts may indeed manifest as growth rate effects, which would lead to larger cumulative impacts over time and suggest that existing social cost of carbon estimates may understate the urgency of climate mitigation. At the same time, our data suggests that adaptation, particularly in the form of air conditioning, may be an important component of assessing the costs and benefits of various climate policies. Damage projections with and without incorporating potential changes in AC penetration over time vary substantially, suggesting that studies which do not account for adaptation may overstate the impacts of climate change on future generations.

Further questions about the impact of heat on learning remain. What portion of the achievement gap between hot and cool countries is explained by the direct impact of heat exposure on learning documented here? Does the impact of heat on learning we document have longer-term impacts on students' economic outcomes? Other than school air conditioning, what other investments or actions can be taken to mitigate the impacts of heat on learning? We hope future work addresses such questions.

References

Acemoglu, Daron, Simon Johnson, and James A Robinson (2001), "The colonial origins of comparative development: An empirical investigation." *The American Economic Review*, 91, 1369–1369.

Albouy, David, Walter Graf, Ryan Kellogg, and Hendrik Wolff (2016), "Climate amenities, climate change, and american quality of life." *Journal of the Association of Environmental and Resource Economists*, 3, 205–246.

Auffhammer, Maximilian and Erin T Mansur (2014), "Measuring climatic impacts on energy consumption: A review of the empirical literature." *Energy Economics*, 46, 522–530.

Barreca, Alan, Karen Clay, Olivier Deschenes, Michael Greenstone, and Joseph S Shapiro (2016), "Adapting to climate change: The remarkable decline in the us temperature-mortality relationship over the twentieth century." *Journal of Political Economy*, 124, 105–159.

Behrer, A Patrick and Jisung Park (2017), "Will we adapt? temperature, labor and adaptation to climate change."

Biddle, Jeff (2008), "Explaining the spread of residential air conditioning, 1955–1980." *Explorations in Economic History*, 45, 402–423.

Bleakley, Hoyt (2010), "Malaria eradication in the americas: A retrospective analysis of childhood exposure." *American Economic Journal: Applied Economics*, 2, 1–45.

Burke, Marshall and Kyle Emerick (2016), "Adaptation to climate change: Evidence from us agriculture." *American Economic Journal: Economic Policy*, 8, 106–40.

Card, David and Jesse Rothstein (2007), "Racial segregation and the black–white test score gap." *Journal of Public Economics*, 91, 2158–2184.

Cellini, Stephanie Riegg, Fernando Ferreira, and Jesse Rothstein (2010), "The value of school facility investments: Evidence from a dynamic regression discontinuity design." *The Quarterly Journal of Economics*, 125, 215–261.

Chambwera, Met al, Geoffrey Heal, C Dubeux, S Hallegatte, L Leclerc, A Markandya, BA McCarl, R Mechler, and JE Neumann (2014), "Economics of adaptation."

Chetty, Raj, John N Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan (2011), "How does your kindergarten classroom affect your earnings? evidence from project star." *The Quarterly Journal of Economics*, 126, 1593–1660.

Cho, H (2017), "Effect of summer heat on test scores: A cohort analysis." *Journal of Environmental Economics and Management*, 83, 185–196.

Cunha, Flavio and James Heckman (2007), "The technology of skill formation." *American Economic Review*, 97, 31–47.

Currie, Janet, Eric A Hanushek, E Megan Kahn, Matthew Neidell, and Steven G Rivkin (2009), "Does pollution increase school absences?" *The Review of Economics and Statistics*, 91, 682–694.

Davis, Lucas W and Paul J Gertler (2015), "Contribution of air conditioning adoption to future energy use under global warming." *Proceedings of the National Academy of Sciences*, 112, 5962–5967.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken (2009), "Temperature and income: reconciling new cross-sectional and panel estimates." *American Economic Review*, 99, 198–204.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken (2012), "Temperature shocks and economic growth: Evidence from the last half century." *American Economic Journal: Macroeconomics*, 4, 66–95.

Dell, Melissa, Benjamin F Jones, and Benjamin A Olken (2014), "What do we learn from the weather? the new climate–economy literature." *Journal of Economic Literature*, 52, 740–798.

Deschenes, Olivier (2014), "Temperature, human health, and adaptation: A review of the empirical literature." *Energy Economics*, 46, 606–619.

Deschênes, Olivier and Michael Greenstone (2011), "Climate change, mortality, and adaptation: Evidence from annual fluctuations in weather in the us." *American Economic Journal: Applied Economics*, 3, 152–185.

Deschenes, Olivier, Michael Greenstone, and Joseph S Shapiro (2017), "Defensive investments and the demand for air quality: Evidence from the nox budget program." *American Economic Review*, 107, 2958–2989.

Dobbie, Will and Roland G Fryer Jr (2011), "Are high-quality schools enough to increase achievement among the poor? evidence from the harlem children's zone." *American Economic Journal: Applied Economics*, 3, 158–87.

Ebenstein, Avraham, Victor Lavy, and Sefi Roth (2016), "The long-run economic consequences of high-stakes examinations: evidence from transitory variation in pollution." *American Economic Journal: Applied Economics*, 8, 36–65.

Fahle, Erin M, Benjamin R Shear, Demetra Kalogrides, Sean F Reardon, Richard DiSalvo, and Andrew D Ho (2017), "Stanford education data archive."

Garg, Teevrat, Maulik Jagnani, and Vis Taraz (2017), "Human capital costs of climate change: Evidence from test scores in india."

Gertler, Paul J, Orie Shelef, Catherine D Wolfram, Alan Fuchs, et al. (2016), "The demand for energy-using assets among the world's rising middle classes." *American Economic Review*, 106, 1366–1401.

Goodman, Joshua (2014), "Flaking out: Student absences and snow days as disruptions of instructional time." *NBER Working Paper*.

Graff Zivin, Joshua, Solomon M Hsiang, and Matthew Neidell (2017), "Temperature and human capital in the short-and long-run." *Journal of the Association of Environmental and Resource Economists*.

Graff Zivin, Joshua and Matthew Neidell (2013), "Environment, health, and human capital." *Journal of Economic Literature*, 51, 689–730.

Graff Zivin, Joshua and Matthew Neidell (2014), "Temperature and the allocation of time: Implications for climate change." *Journal of Labor Economics*, 32, 1–26.

Graff-Zivin, Joshua, Matthew Neidell, and Wolfram Schlenker (2011), "Water quality violations and avoidance behavior: Evidence from bottled water consumption." *The American Economic Review: Papers and Proceedings*, 101, 448–453.

Hallegatte, Stephane, Jisung Park, Mook Bangalore, and Evan Sandhoefner (2018), "Households and heat stress: estimating the distributional consequences of climate change." *Environment and Development Economics*.

Heal, Geoffrey and Jisung Park (2016), "Temperature stress and the direct impact of climate change: A review of an emerging literature." *Review of Environmental Economics and Policy*, 10, 347–362.

Hoxby, Caroline M (2001), "All school finance equalizations are not created equal." *The Quarterly Journal of Economics*, 116, 1189–1231.

Hsiang, Solomon M (2010), "Temperatures and cyclones strongly associated with economic production in the caribbean and central america." *Proceedings of the National Academy of sciences*, 107, 15367–15372.

Hsiang, Solomon M, Marshall Burke, and Edward Miguel (2013), "Quantifying the influence of climate on human conflict." *Science*, 341, 1235367.

Hsiang, Solomon M, Kyle C Meng, and Mark A Cane (2011), "Civil conflicts are associated with the global climate." *Nature*, 476, 438.

Isen, Adam, Maya Rossin-Slater, and Reed Walker (2017), "Relationship between season of birth, temperature exposure, and later life wellbeing." *Proceedings of the National Academy of Sciences*, 201702436.

Jackson, C Kirabo, Rucker C Johnson, and Claudia Persico (2015), "The effects of school spending on educational and economic outcomes: Evidence from school finance reforms." *The Quarterly Journal of Economics*, 131, 157–218.

Kahn, Matthew E (2016), "The climate change adaptation literature." *Review of Environmental Economics and Policy*, 10, 166–178.

Kenrick, Douglas T and Steven W MacFarlane (1986), "Ambient temperature and horn honking: A field study of the heat/aggression relationship." *Environment and behavior*, 18, 179–191.

Kjellstrom, Tord, David Briggs, Chris Freyberg, Bruno Lemke, Matthias Otto, and Olivia Hyatt (2016), "Heat, human performance, and occupational health: a key issue for the assessment of global climate change impacts." *Annual review of public health*, 37, 97–112.

Lafortune, Julien, Jesse Rothstein, and Diane Whitmore Schanzenbach (2016), "School finance reform and the distribution of student achievement." Technical report, National Bureau of Economic Research.

Marcotte, Dave E (2017), "Something in the air? air quality and children's educational outcomes." *Economics of Education Review*, 56, 141–151.

Martorell, Paco, Kevin Stange, and Isaac McFarlin (2016), "Investing in schools: capital spending, facility conditions, and student achievement." *Journal of Public Economics*, 140, 13–29.

McFarland, Joel, Bill Hussar, Cristobal de Brey, Tom Snyder, Xiaolei Wang, Sidney Wilkinson-Flicker, Semhar Gebrekristos, Jijun Zhang, Amy Rathbun, Amy Barmer, et al. (2017), "The condition of education 2017. nces 2017-144." *National Center for Education Statistics*.

Neilson, Christopher A and Seth D Zimmerman (2014), "The effect of school construction on test scores, school enrollment, and home prices." *Journal of Public Economics*, 120, 18–31.

Park, Jisung (2017), "Hot temperature and high stakes exams: Evidence from new york city public schools." *Harvard University Working Paper*.

Pindyck, Robert S (2013), "Climate change policy: What do the models tell us?" *Journal of Economic Literature*, 51, 860–872.

Schlenker, Wolfram, W Michael Hanemann, and Anthony C Fisher (2006), "The impact of global warming on us agriculture: an econometric analysis of optimal growing conditions." *The Review of Economics and Statistics*, 88, 113–125.

Schwartz, Joel, Jonathan M Samet, and Jonathan A Patz (2004), "Hospital admissions for heart disease: the effects of temperature and humidity." *Epidemiology*, 15, 755–761.

Seppanen, Olli, William J Fisk, and QH Lei (2006), "Effect of temperature on task performance in office environment." *Lawrence Berkeley National Laboratory*.

Shah, Manisha and Bryce Millett Steinberg (2017), "Drought of opportunities: Contemporaneous and long-term impacts of rainfall shocks on human capital." *Journal of Political Economy*, 125, 527–561.

Figure 1: Climate and Achievement Across Geographies

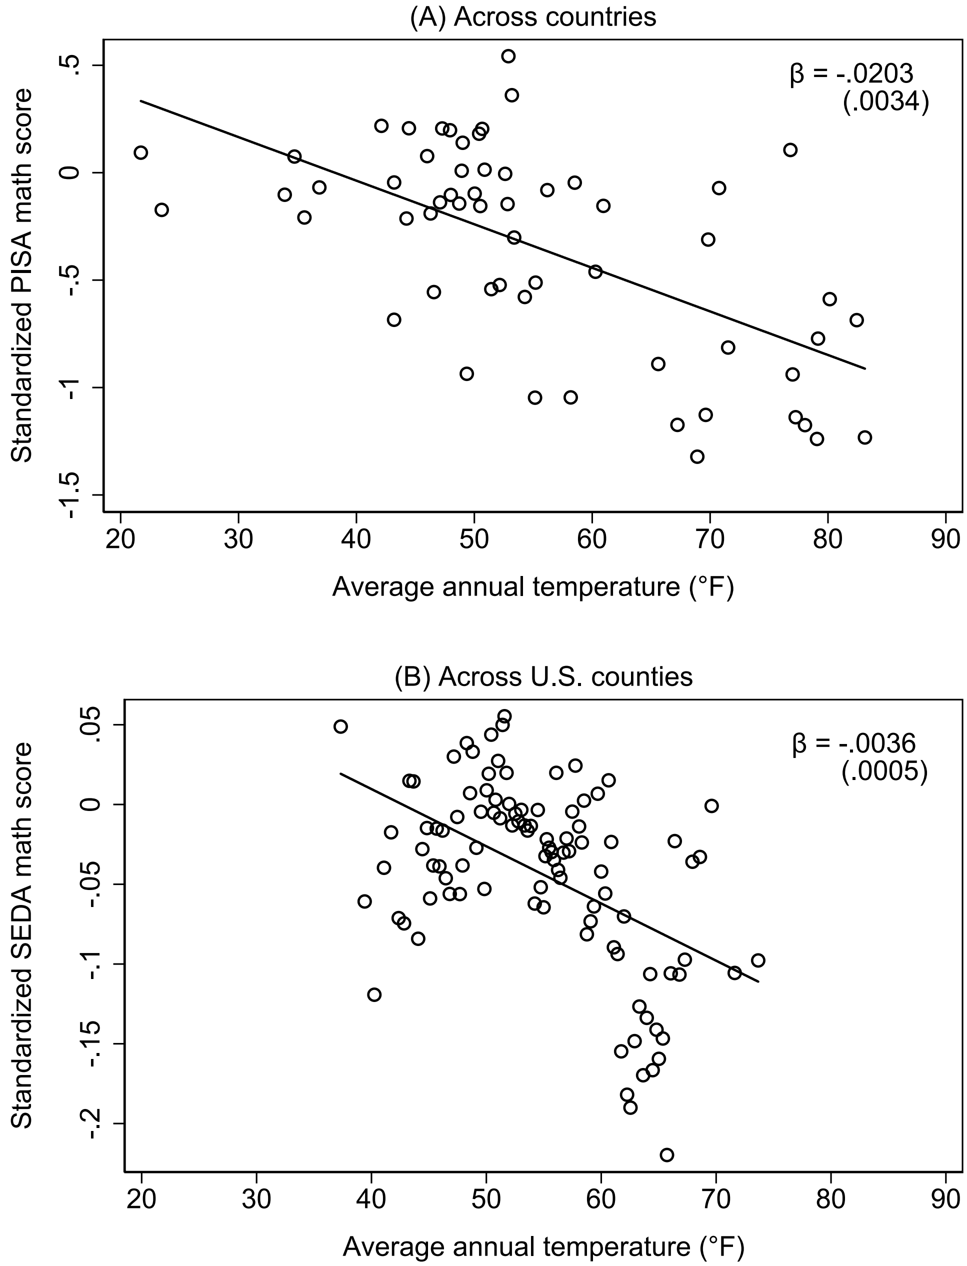

Notes: The above figure shows a scatterplot of mean 2012 PISA (panel A) or SEDA (panel B) math scores and average annual temperature by country or U.S. county. Average annual temperatures are measured over the period 1980-2011. Panel B shows a binned percentile plot of standardized 3rd-8th grade math scores (2009-2013) by percentile of the county-level average temperature distribution, with scores standardized by subject, grade and year as in Fahle et al. (2017). Also shown is a fitted line and slope coefficient from a bivariate regression of scores on temperatures, using heteroskedasticity robust standard errors.

Figure 2: Spatial Distribution of PSAT Z-Scores

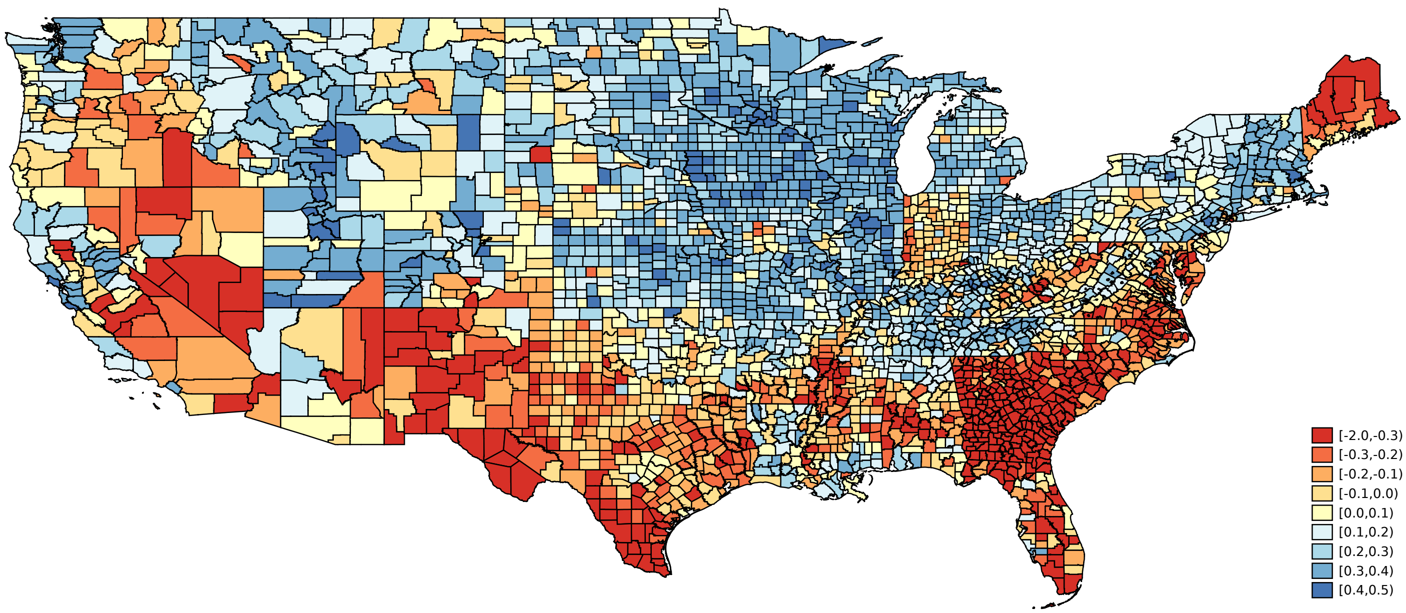

Notes: The above figure shows county-level average standardized PSAT scores from the high school classes of 2001-14.

Figure 3: Spatial Variation in Prior Year Temperature

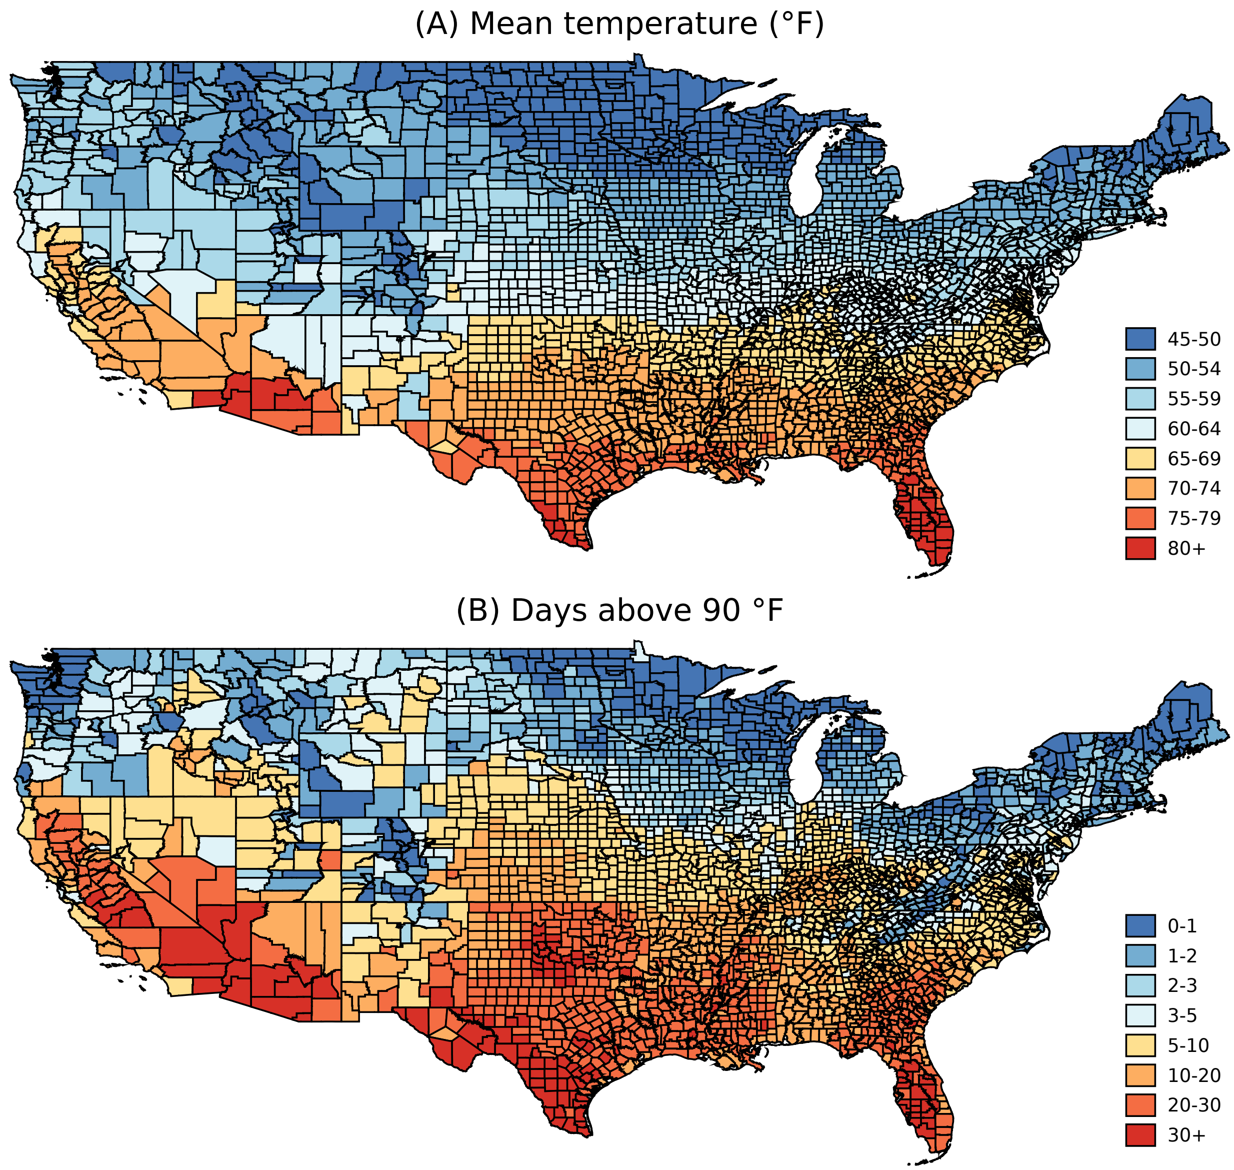

Notes: The above figure shows the mean daily maximum temperature (panel A) and number of days above 90°F (panel B) experienced by students on school days in the 365 days prior to taking the PSAT, by county. The sample consists of all PSAT-takers from the high school classes of 2001-14, whose PSATs were taken between 1997 and 2012.

Figure 4: Cumulative Hot Days and Test Performance

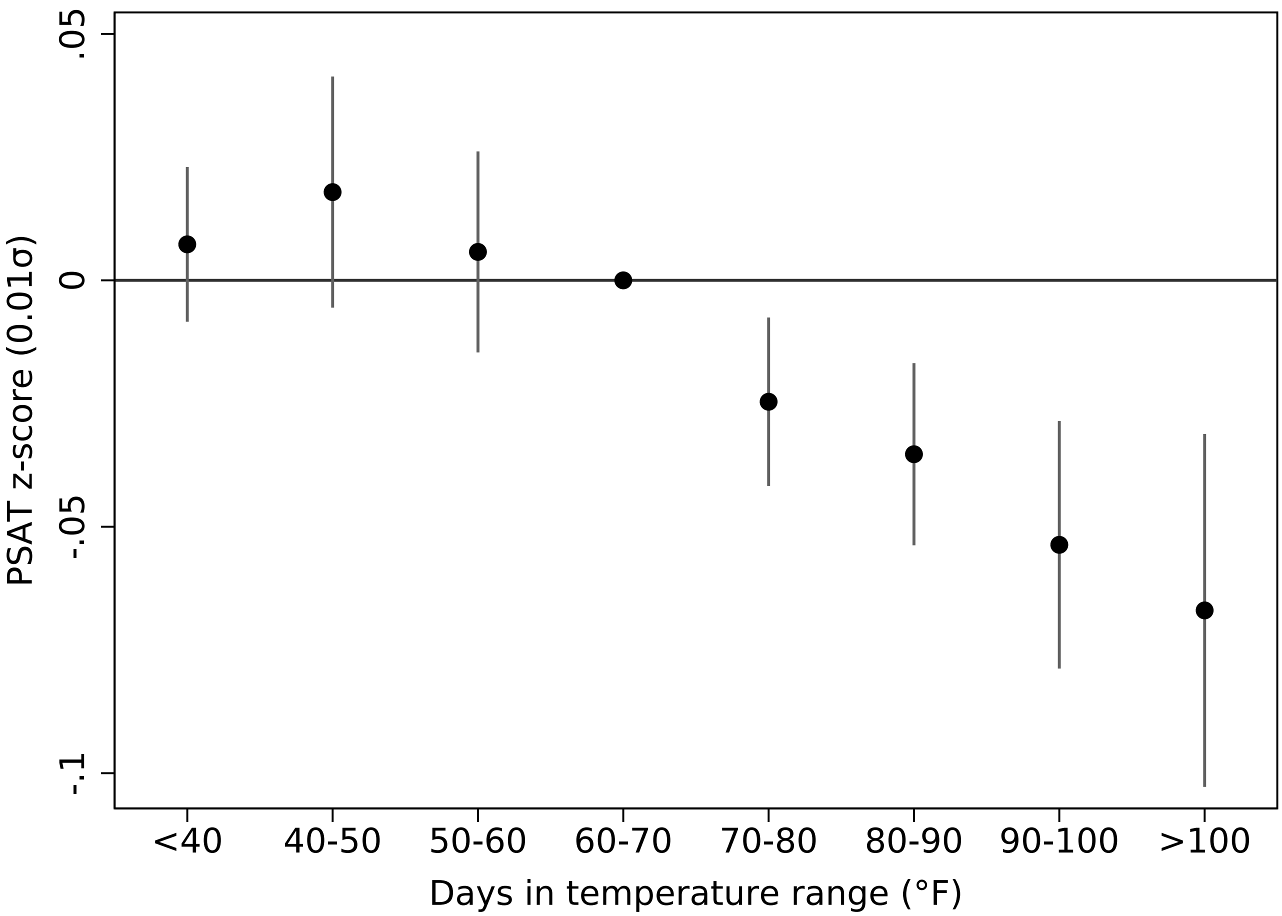

Notes: Shown above are coefficients from a regression of hundredths of a standard deviation in PSAT total (math plus reading) scores on the number of school days within a given temperature range during the 365 days preceding the PSAT take. The regression includes student fixed effects and fixed effects for each combination of cohort, test date and take number. Heteroskedasticity robust standard errors clustered by weather sensor are in parentheses (* p<.10 ** p<.05 *** p<.01). The sample comprises all students from the high school classes of 2001-14 who took the PSAT more than once.

Figure 5: Hot Classrooms

(A) Fraction of hot days when classrooms get too hot (students)

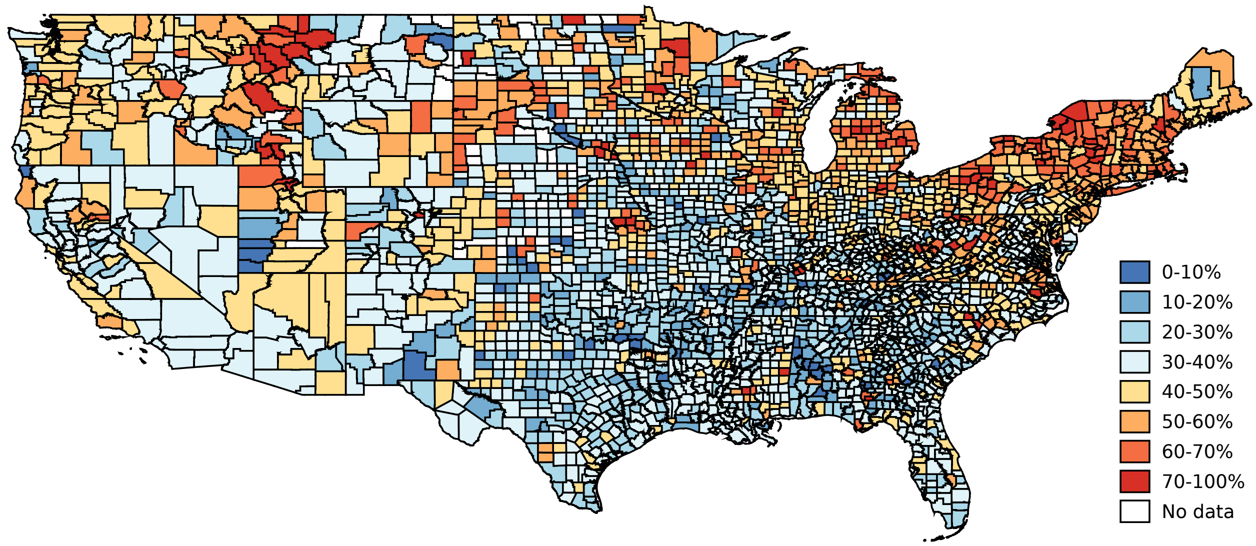

(B) Fraction of hot days when classrooms get too hot (counselors)

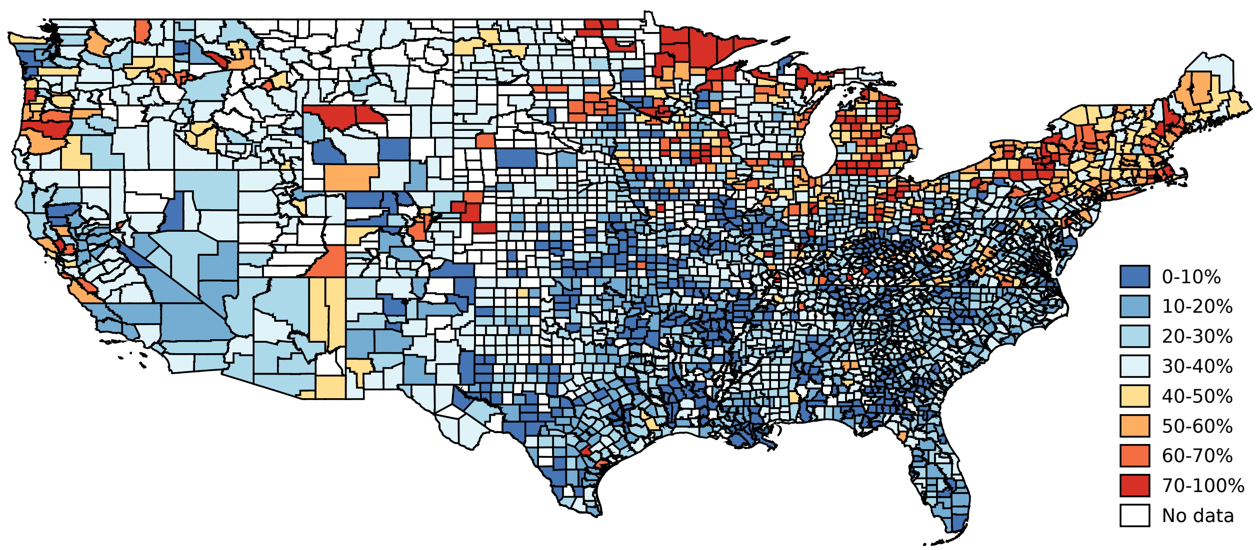

Notes: The above figure shows by county the mean fraction of classrooms reported as “too hot” on hot days by students (panel A) and guidance counselors (panel B). Both measures are derived from student or counselor responses to a survey administered by the College Board. The sample consists of all PSAT-takers from the high school classes of 2001-14, whose PSATs were taken between 1997 and 2012.

Figure 6: School and Home Air Conditioning

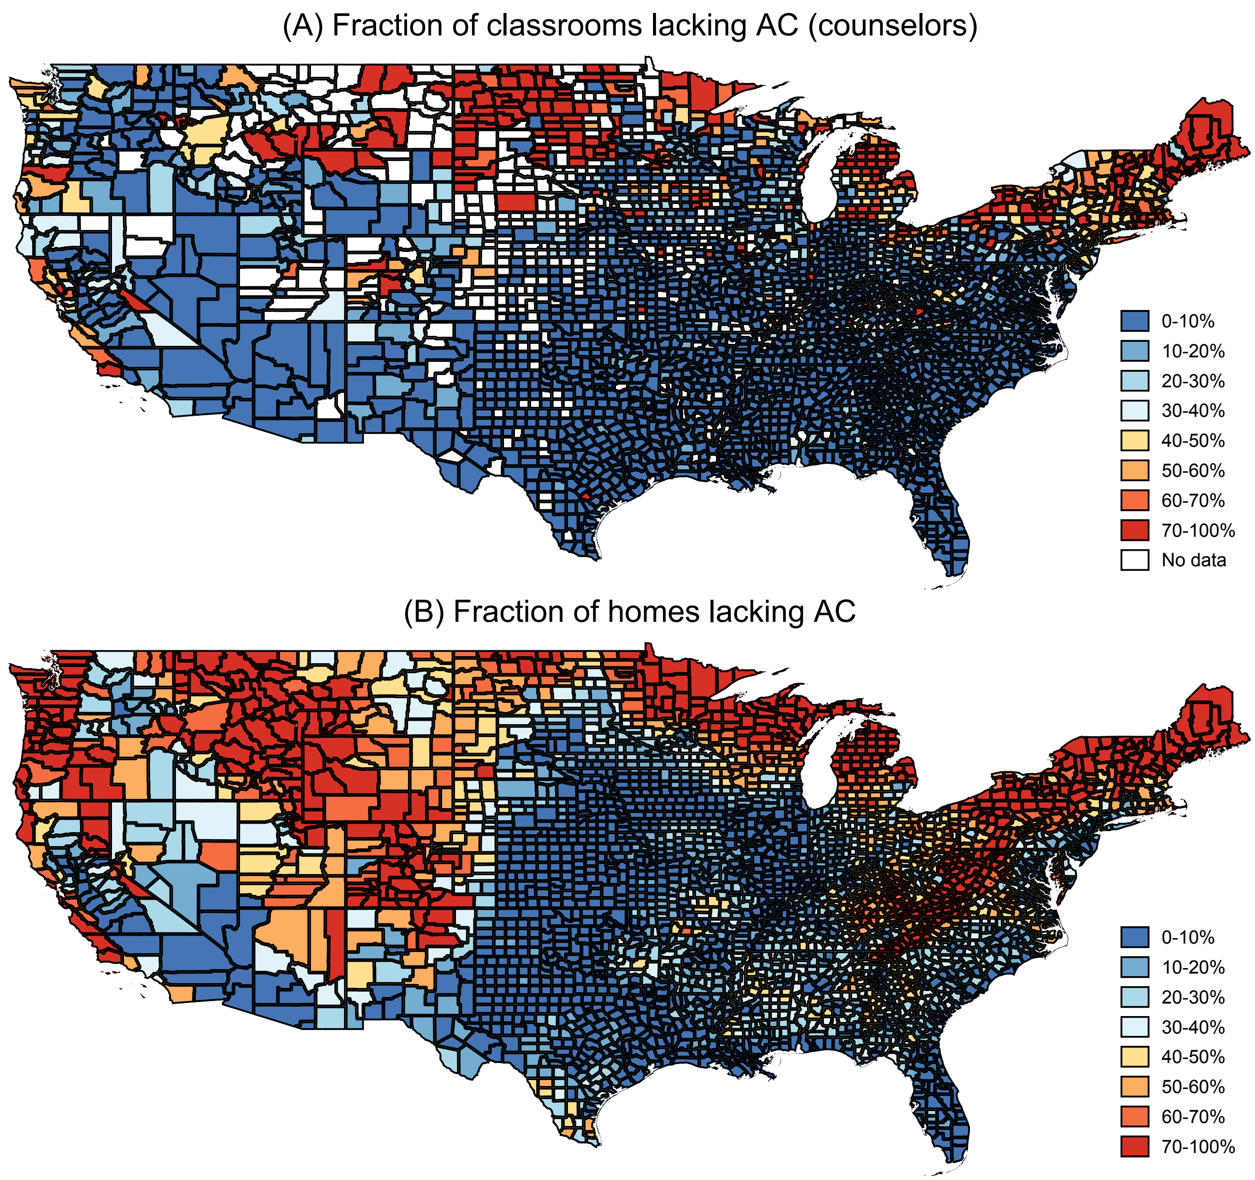

Notes: The above figure shows by county the mean fraction of classrooms (panel A) and homes (panel B) lacking air conditioning. Classroom measures are derived from guidance counselor responses to a survey administered by the College Board. Home measures are derived from the 1980 Census and the 1993-2015 quadrennial Residential Energy Consumption Surveys. The sample consists of all PSAT-takers from the high school classes of 2001-14, whose PSATs were taken between 1997 and 2012.

Figure 7: School Air Conditioning by Percent Black or Hispanic

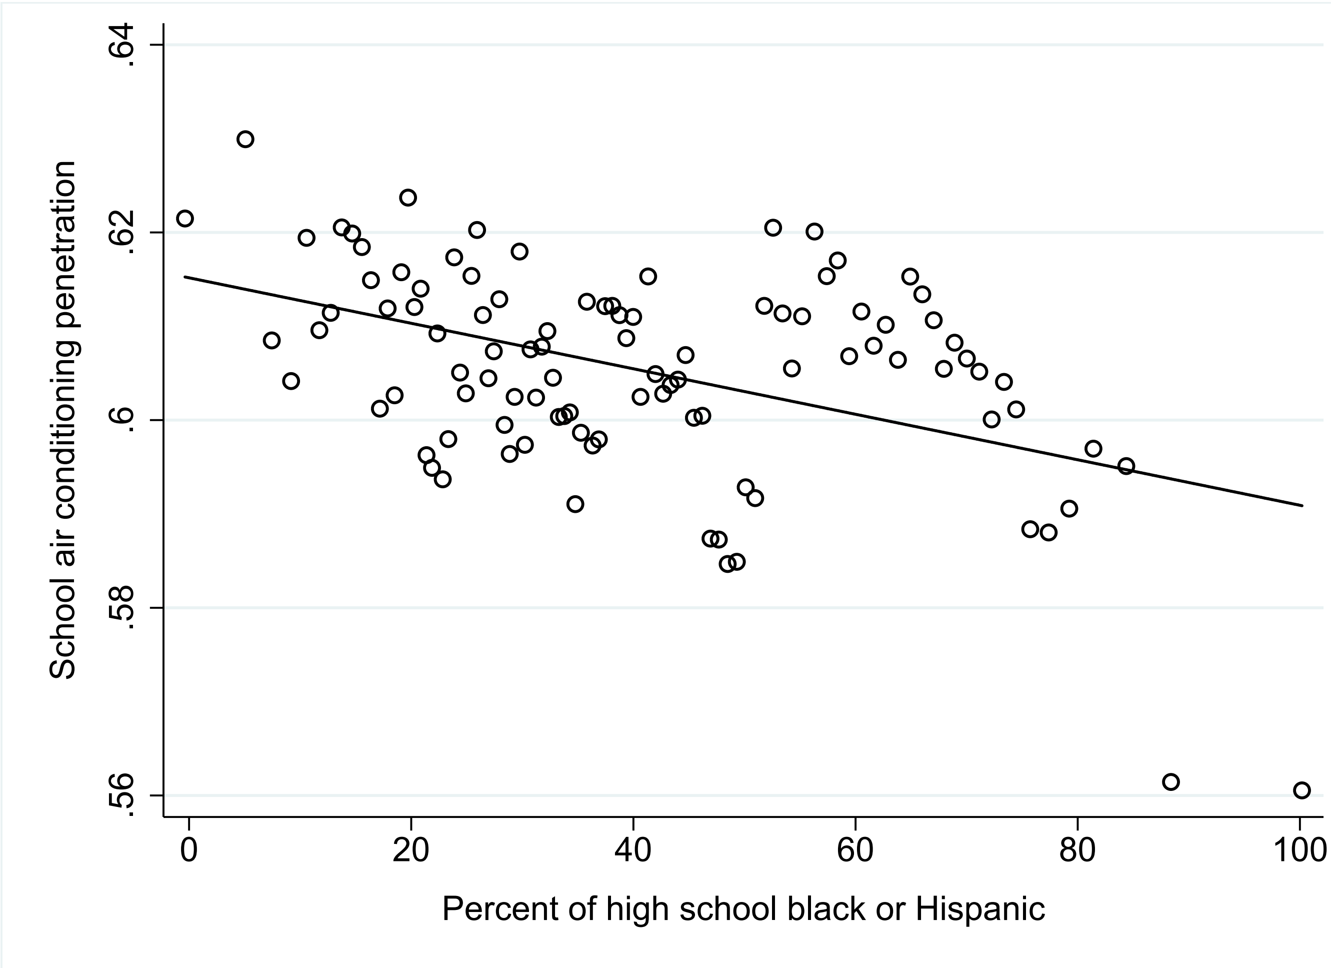

Notes: The above figure is a binned percentile plot of high school air conditioning penetration rates as implied by student survey responses, by percentile of the school-level percent black or Hispanic distribution. It plots residual variation after controlling for average daily maximum temperature by school and average income by zip code between 1997 and 2012.

Table 1: Summary Statistics

		PSAT Retakers				
	All takers (1)	All retakers (2)	Black or Hispanic (3)	White (4)	Lowest income (5)	Highest income (6)
(A) Demographics						
Female	0.53	0.55	0.56	0.54	0.57	0.52
White	0.58	0.58	0.00	1.00	0.29	0.73
Black or Hispanic	0.29	0.28	1.00	0.00	0.61	0.11
Mother has B.A.	0.22	0.33	0.18	0.40	0.14	0.53
ZIP code mean income	63.22	69.58	49.41	78.25	31.59	145.70
(B) PSAT scores						
Retook PSAT	0.36	1.00	1.00	1.00	1.00	1.00
Total takes	1.42	2.15	2.17	2.14	2.16	2.15
First PSAT z-score	-0.00	0.14	-0.49	0.40	-0.40	0.58
(C) Temperature						
Mean temperature (F)	65.1	65.8	68.8	64.2	67.7	65.0
Days above 90 F	11.9	12.2	15.7	10.6	14.8	10.9
(D) Air conditioning						
Classrooms with AC	0.58	0.59	0.60	0.58	0.59	0.61
Homes with AC	0.77	0.80	0.85	0.79	0.79	0.83
N (scores)	38,303,474	21,076,009	6,023,145	12,161,058	4,531,817	4,107,140
N (students)	27,023,119	9,795,654	2,775,607	5,689,371	2,099,070	1,906,942

Notes: Notes: Mean values of key variables are shown. Column 1 includes comprises all students from the high school classes of 2001-14 who took the PSAT at least once. Column 2 includes only those who took the PSAT more than once. Columns 3-6 include subgroups of retakers, with columns 5 and 6 respectively including the lowest and highest quintiles of ZIP code mean income.

Table 2: Prior Year Temperature and PSAT Scores

	(1)	(2)	(3)	(4)	(5)	(6)
(A) Average heat						
Mean temperature (°F)	-0.181***	-0.211***	-0.185***	-0.182***	-0.176***	-0.230***
	(0.028)	(0.036)	(0.029)	(0.028)	(0.027)	(0.042)
(B) Hot days						
Days above 100 °F	-0.067***	-0.077***	-0.075***	-0.065***	-0.064***	-0.098***
	(0.018)	(0.020)	(0.017)	(0.018)	(0.016)	(0.033)
Days in 90s (°F)	-0.053***	-0.061***	-0.059***	-0.053***	-0.053***	-0.064***
	(0.013)	(0.014)	(0.013)	(0.013)	(0.012)	(0.018)
Days in 80s (°F)	-0.035***	-0.037***	-0.039***	-0.035***	-0.029***	-0.046***
	(0.009)	(0.010)	(0.009)	(0.009)	(0.009)	(0.013)
Days in 70s (°F)	-0.024***	-0.024***	-0.026***	-0.024***	-0.023***	-0.020*
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)	(0.012)
Days below 60 °F	0.010	0.013	0.008	0.010	0.010	0.007
	(0.009)	(0.009)	(0.008)	(0.009)	(0.007)	(0.011)
N	21,046,448	21,046,448	21,046,448	21,046,448	21,046,448	5,378,273
Prior year snow, rain	No	Yes	No	No	No	No
Test day weather	No	Yes	No	No	No	No
Pollution	No	No	Yes	No	No	No
Economic conditions	No	No	No	Yes	No	No
State-specific trends	No	No	No	No	Yes	No
Sensor within 5 miles	No	No	No	No	No	Yes

Notes: Heteroskedasticity robust standard errors clustered by weather sensor are in parentheses (* p<.10 ** p<.05 *** p<.01). Coefficients in each column and panel come from a regression of hundredths of a standard deviation in PSAT total (math plus reading) scores on the weather measure(s) shown. Temperature is measured with the daily maximum temperature from school days in the 365 days preceding the PSAT take. All regressions include student fixed effects and fixed effects for each combination of cohort, test date and take number. Column 2 adds controls for prior year rainfall and snowfall, as well as test day temperature, rainfall and snowfall. Column 3 controls for prior year and test day pollution levels (carbon monoxide, ozone, suflur dioxide, nitrogen dioxide and PM10). Column 4 controls for the logarithm of per capita county-level payroll in industries highly exposed to weather. Column 5 adds state-specific linear time trends. Column 6 limits the sample to high schools within 5 miles of the nearest weather sensor. The sample comprises all students from the high school classes of 2001-14 who took the PSAT more than once.

Table 3: Timing of Temperature Shocks

	(1)	(2)	(3)	(4)	(5)
(A) Mean temperature (°F)					
School days, 1 year prior	-0.181***	-0.205***	-0.270***		-0.185***
	(0.028)	(0.030)	(0.043)		(0.041)
Summer days, 1 year prior		0.039		0.047*	
		(0.026)		(0.026)	
Weekend days, 1 year prior			0.114***		
			(0.038)		
School days, post-summer				-0.061**	
				(0.025)	
School days, pre-summer				-0.160***	
				(0.029)	
School days, 2 years prior					-0.010
					(0.052)
School days, 3 years prior					0.012
					(0.043)
(B) Days above 90 °F					
School days, 1 year prior	-0.056***	-0.061***	-0.073***		-0.078***
	(0.012)	(0.011)	(0.016)		(0.014)
Summer days, 1 year prior		0.016		0.018	
		(0.011)		(0.011)	
Weekend days, 1 year prior			0.043		
			(0.028)		
School days, post-summer				-0.074***	
				(0.019)	
School days, pre-summer				-0.074***	
				(0.016)	
School days, 2 years prior					-0.051***
					(0.013)
School days, 3 years prior					-0.048***
					(0.014)
N	21,046,448	21,046,448	21,046,448	21,046,448	21,046,448

Notes: Heteroskedasticity robust standard errors clustered by weather sensor are in parentheses (* $p<.10$ ** $p<.05$ *** $p<.01$). Coefficients in each column and panel come from a regression of hundredths of a standard deviation in PSAT total (math plus reading) scores on the weather measure(s) shown. School day temperature is measured with the daily maximum temperature from school days in the listed 365 day period relative to the PSAT take. Summer temperature is measured across all days in the summer break preceding the PSAT take. Weekend temperature is measured across all weekends and national holidays in the 365 days preceding the PSAT take. All regressions include student fixed effects and fixed effects for each combination of cohort, test date and take number. The sample comprises all students from the high school classes of 2001-14 who took the PSAT more than once.

Table 4: Heterogeneity by Race, Income and Geography

	Black/ Hispanic (1)	White (2)	High minority HS (3)	Low minority HS (4)	Low income ZIP code (5)	High income ZIP code (6)	Coolest areas (7)	Hottest areas (8)
Mean temperature (°F)	-0.320***	-0.093***	-0.379***	-0.071***	-0.296***	-0.134***	-0.215***	-0.170***
	(0.043)	(0.019)	(0.052)	(0.027)	(0.043)	(0.027)	(0.047)	(0.047)
Days above 90 °F	-0.072***	-0.027***	-0.072***	0.003	-0.076***	-0.031**	-0.082***	-0.025*
	(0.015)	(0.008)	(0.017)	(0.014)	(0.016)	(0.013)	(0.029)	(0.013)
N	6,023,145	12,161,058	4,129,153	4,239,916	4,531,817	4,107,140	10,535,013	10,511,435

Notes: Heteroskedasticity robust standard errors clustered by weather sensor are in parentheses (* $p<.10$ ** $p<.05$ *** $p<.01$). Each coefficient comes from a separate regression of hundredths of a standard deviation in PSAT total (math plus reading) scores on the weather measure(s) shown. The first row measures mean temperature using the daily maximum temperature from school days in the 365 days preceding the PSAT take. The second row measures the number of such school days above 90 °F and controls for the number of days in other temperature ranges, so that days in the 60s are the reference category. All regressions include student fixed effects and fixed effects for each combination of cohort, test date and take number. The sample comprises all students from the high school classes of 2001-14 who took the PSAT more than once. Columns 3 and 4 contain students attending high schools in the highest and lowest quintile of fraction of PSAT-takers who are black or Hispanic. Columns 5 and 6 contain students living in the lowest and highest quintiles of ZIP code-level income. Columns 7 and 8 contain students living in areas with below and above median school year average temperatures.

Table 5: Heterogeneity in School Air Conditioning Access

	All schools (1)	Cooler areas (2)	Hotter areas (3)
(A) By high school racial composition			
High minority HS	-0.035***	-0.035***	-0.034**
	(0.010)	(0.013)	(0.014)
Middle minority HS	-0.020***	-0.022**	-0.020*
	(0.008)	(0.009)	(0.012)
(B) By ZIP code income			
Low income	-0.043***	-0.055***	-0.029***
	(0.007)	(0.013)	(0.007)
Middle income	-0.026***	-0.034***	-0.017***
	(0.005)	(0.008)	(0.005)
N	22,347,878	11,176,342	11,171,536

Notes: Heteroskedasticity robust standard errors clustered by weather sensor are in parentheses (* p<.10 ** p<.05 *** p<.01). Each column comes from a regression of fraction of classrooms air conditioned on indicators for categories of high school racial composition or ZIP code income. High (middle) minority high schools are those whose proportion of black or Hispanic students are in the top (middle three) quintile(s) of the distribution. High (middle) income referes to those ZIP codes whose income is in the top (middle three) quintile(s) of the distribution. Each regression controls for a quartic in school-level mean temperatures over the entire time period. Cooler and hotters areas identify schools whose long-term mean temperatures are below or above the median. The sample comprises all students from the high school classes of 2001-14 who took the PSAT.

Table 6: Adaptation through School Air Conditioning

	(1)	(2)	(3)	(4)
Mean temp.	-0.322***	-0.565***	-0.229***	-0.457***
	(0.067)	(0.103)	(0.060)	(0.168)
Mean temp. * School AC penetration	0.253**	0.227***		
	(0.099)	(0.071)		
Mean temp. * School AC penetration change * HS class			0.117***	0.116***
			(0.039)	(0.038)
Mean temp. * HS class			-0.011	0.004
			(0.015)	(0.016)
Mean temp. * School AC penetration change			-0.218	-0.272*
			(0.168)	(0.153)
Mean temp. * Home AC penetration		0.323***		0.213
		(0.110)		(0.185)
N	18,665,967	18,665,967	2,935,907	2,935,907
Interactions with area income, racial composition, temperature	N	Y	N	Y

Notes: Heteroskedasticity robust standard errors clustered by weather sensor are in parentheses (* $p<.10$ ** $p<.05$ *** $p<.01$). Coefficients in each column and panel come from a regression of hundredths of a standard deviation in PSAT total (math plus reading) scores on the heat measure and interaction shown. Mean temperature is measured by the daily maximum temperature from school days in the 365 days preceding the PSAT take. Columns 1 and 2 interact heat with school air conditioning penetration rates as reported by students in 2016. Columns 3 and 4 interact heat with high school class, the change in school air conditioning penetration rates between 2006 and 2016 as reported by guidance counselors, and the interaction of those two variables. Columns 2 and 4 also control for interactions between heat and county-level home air conditioning penetration rates, as well as (not shown) sensor-level mean school year temperature, ZIP code-level income, and the school-level fraction of PSAT-takers who are black or Hispanic. All regressions include student fixed effects and fixed effects for each combination of cohort, test date and take number. The sample comprises all students from the high school classes of 2001-14 who took the PSAT more than once and whose school air conditioning penetration rate (columns 1 and 2) or change in that rate (columns 3 and 4) are non-missing.

Figure A.1: School Year Calendars by State

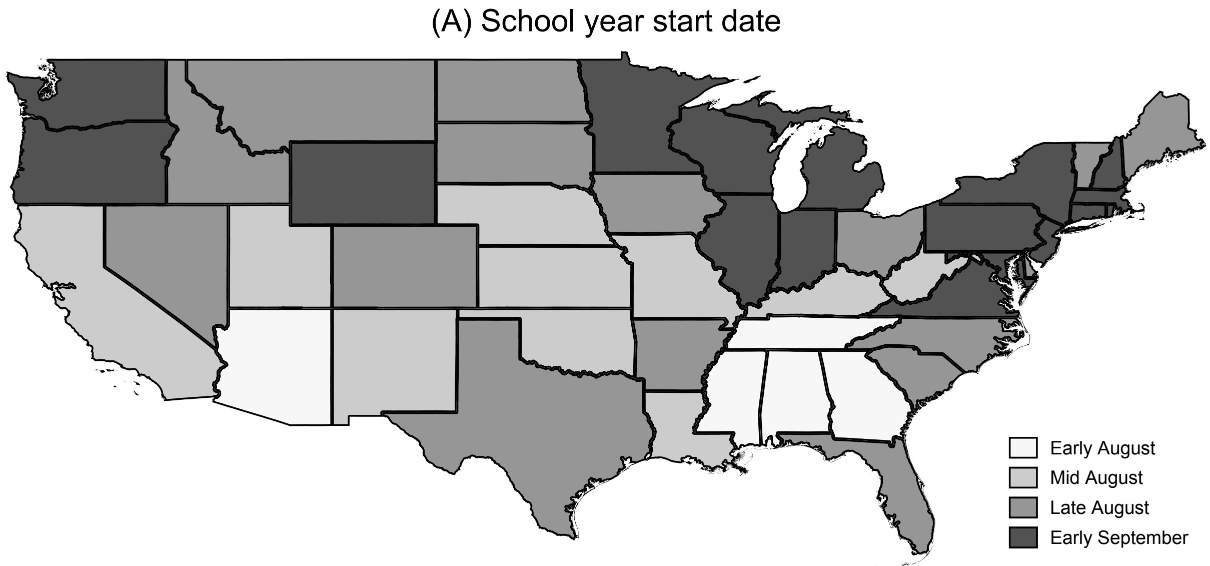

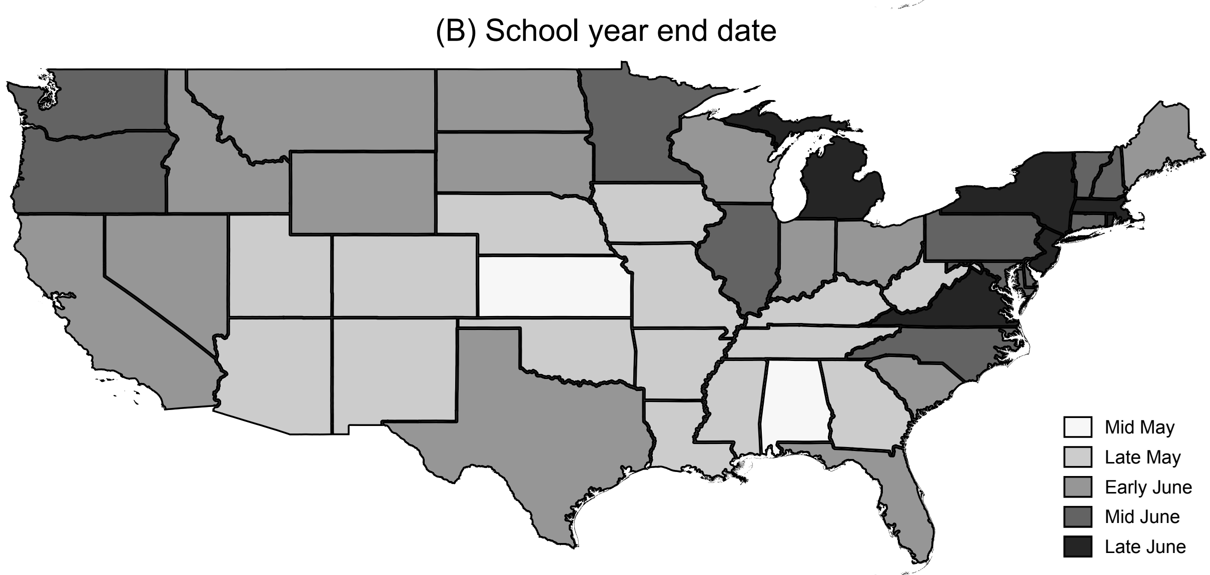

Notes: The above figure shows state's approximate school year start and end dates based on the largest school district in each state and as of 2016.

Figure A.2: Spatial Distribution of PSAT Takers

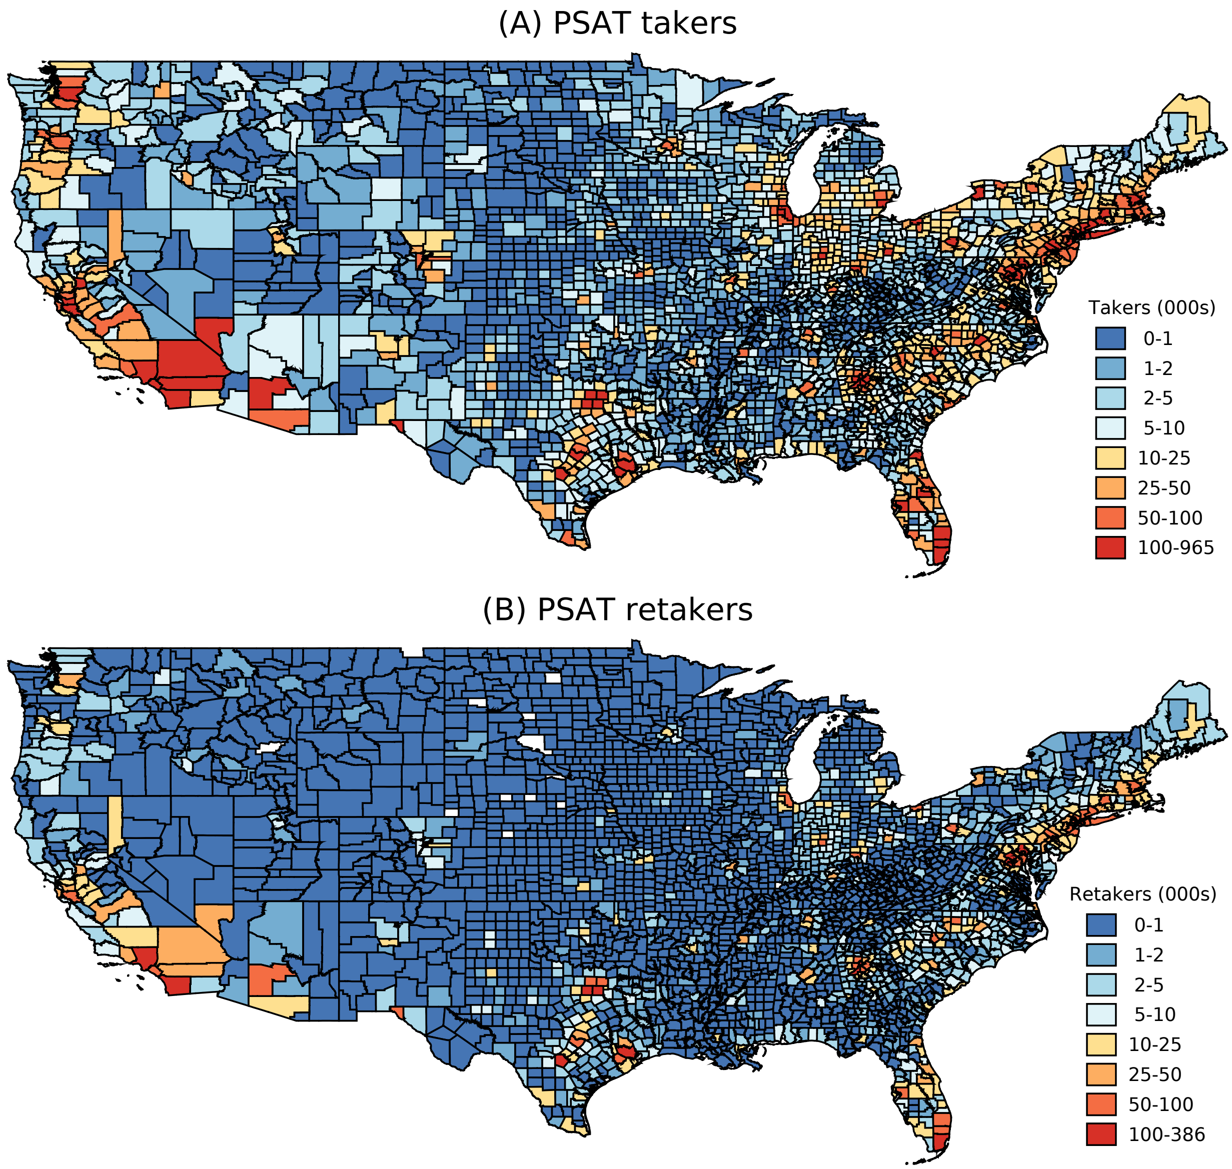

Notes: The above figure shows by county the total number of PSAT takers (panel A) and retakers (panel B) from the high school classes of 2001-14.

Figure A.3: Temporal Variation in Prior Year Temperature

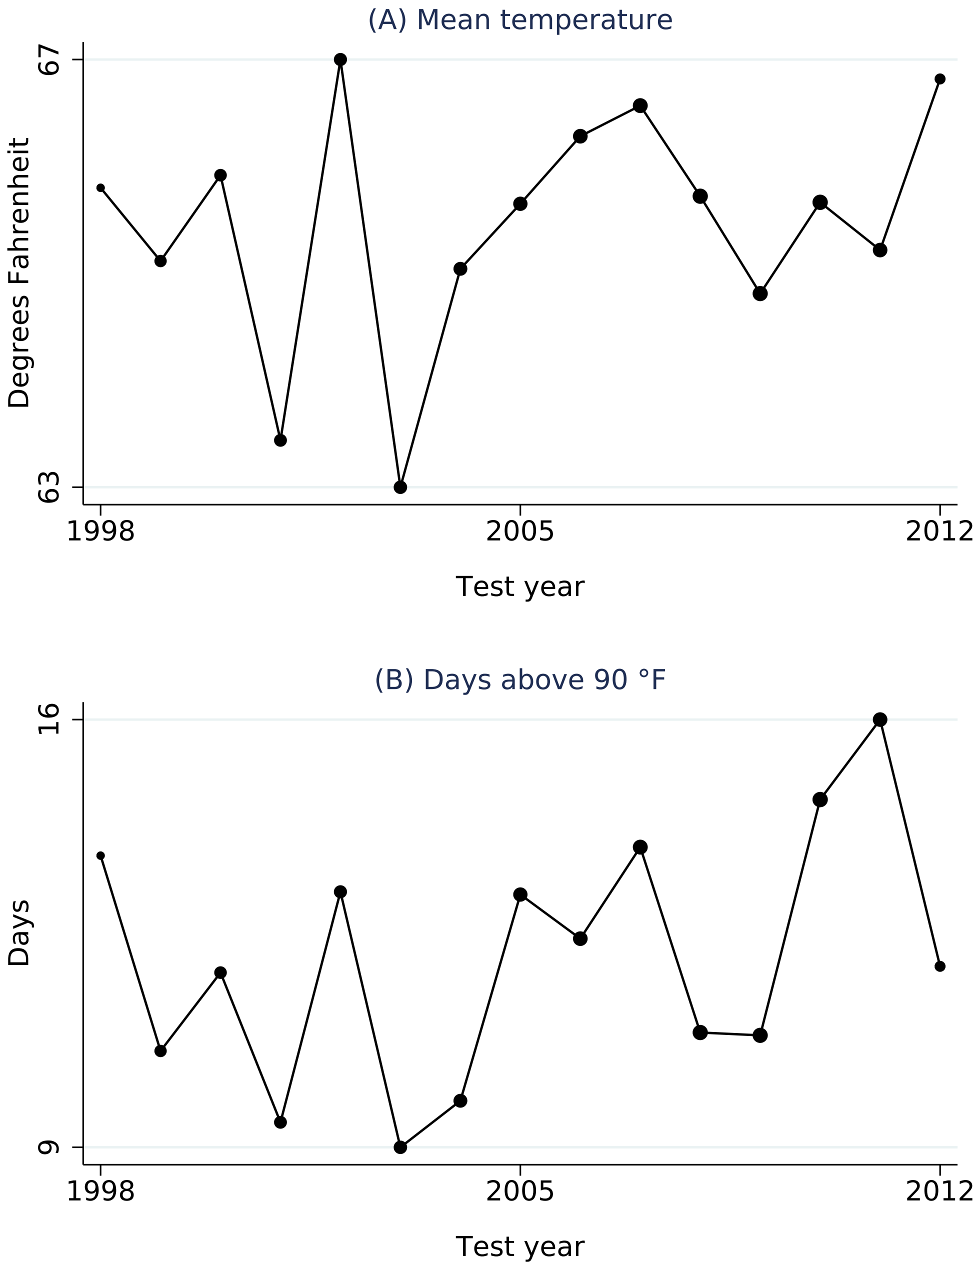

Notes: The above figure shows the mean daily maximum temperature (panel A) and number of days above 90°F (panel B) experienced by students on school days in the 365 days prior to taking the PSAT, by test year. The sample consists of all PSAT-takers from the high school classes of 2001-14, whose PSATs were taken between 1997 and 2012. Dot size is proportional to the number of students in each test year. Test year 1997 is excluded due to the small number of observations.

Figure A.4: Residuals of Prior Year Temperature

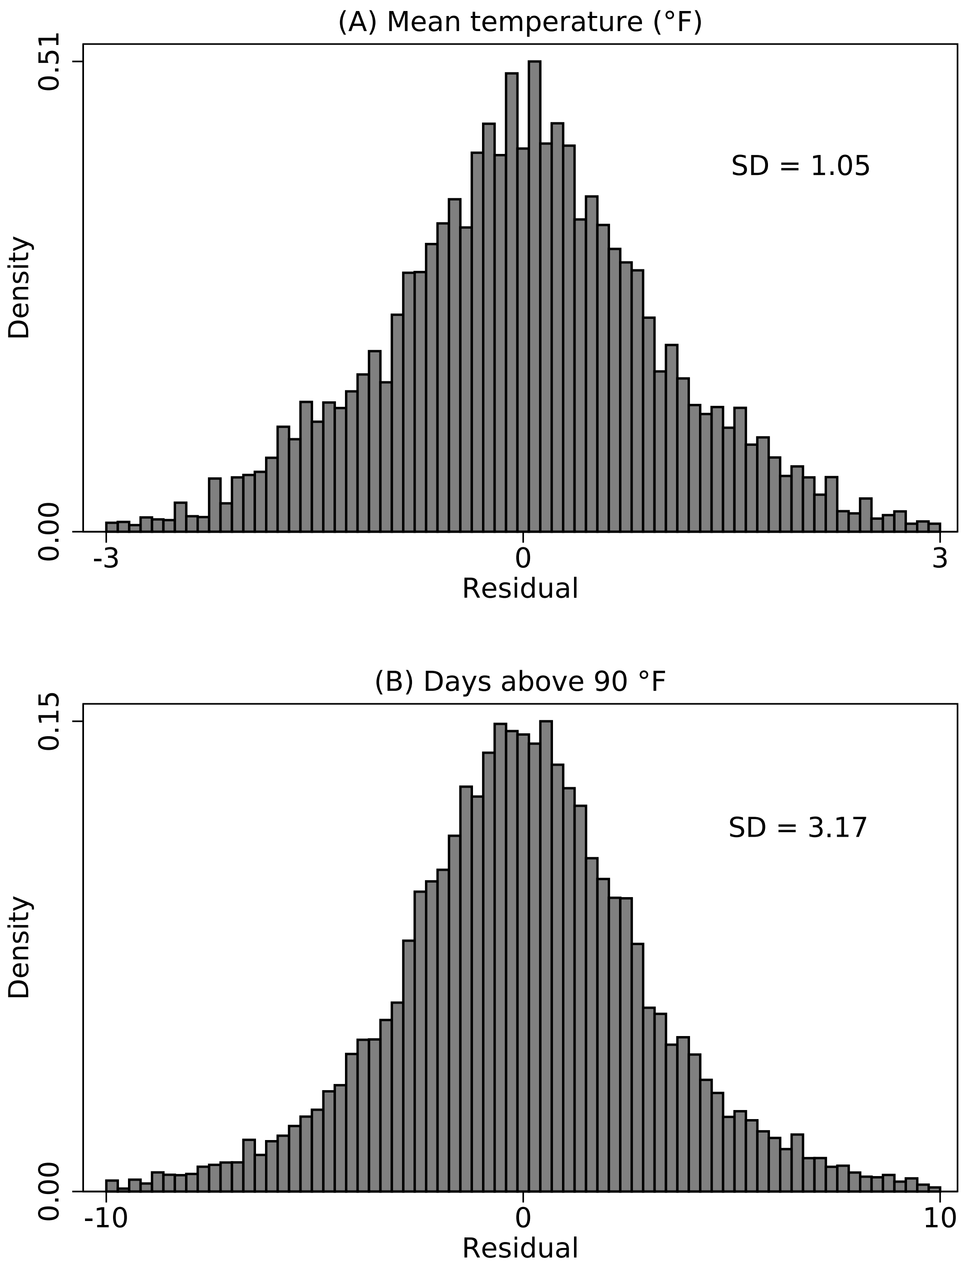

Notes: The above figure shows the distribution of residuals resulting from regressions on student fixed effects of the mean daily maximum temperature (panel A) and number of days above 90°F (panel B) experienced by students on school days in the 365 days prior to taking the PSAT. All regressions include fixed effects for each combination of cohort, test date and take number. The figure excludes residuals with magnitude above three (panel A) and 10 (panel B). The standard deviation of the full set of residuals is shown in each panel.

Table A.1: Temperature and PSAT-Taking

	Takers (1)	Ln(takers) (2)	Female (3)	Black or Hispanic (4)	Mother has B.A. (5)	Father has B.A (6)
(A) Baseline						
Mean temperature (°F)	-0.0107	-0.0021	-0.0001	0.0007***	0.0003	0.0003
	(0.0902)	(0.0017)	(0.0001)	(0.0002)	(0.0003)	(0.0003)
(B) State trends						
Mean temperature (°F)	-0.0924	-0.0013	0.0001	0.0001	-0.0001	-0.0001
	(0.0812)	(0.0016)	(0.0001)	(0.0002)	(0.0002)	(0.0002)
N	686,977	686,977	27,021,552	27,021,552	27,021,552	27,021,552

Notes: Heteroskedasticity robust standard errors clustered by weather sensor are in parentheses (* $p<.10$ ** $p<.05$ *** $p<.01$). Each coefficient comes from a separate regression of the listed characteristic of PSAT-takers on the average daily maximum temperature from school days in the 365 days preceding a student's first PSAT take. All regressions include fixed effects for each high school and for each combination of cohort and test date. Panel B also includes state-specific linear time trends. The sample comprises all students from the high school classes of 2001-14 who took the PSAT at least once.

Table A.2: Temperature and Retaking

	(1)	(2)	(3)	(4)	(5)
Prior year temperature (°F)	0.0005		0.0005	0.0003	-0.0001
	(0.0005)		(0.0005)	(0.0005)	(0.0004)
Following year temperature (°F)		0.0003	0.0003	0.0003	-0.0004
		(0.0004)	(0.0005)	(0.0005)	(0.0004)
N	27,021,551	27,021,551	27,021,551	27,021,551	27,021,551
Test day temperature	No	No	No	Yes	Yes
State-specific time trends	No	No	No	No	Yes

Notes: Heteroskedasticity robust standard errors clustered by weather sensor are in parentheses (* $p<.10$ ** $p<.05$ *** $p<.01$). Each coefficient comes from a separate regression of the probability of retaking the PSAT on the weather measure(s) shown. Yearly temperatures are measured with the daily maximum temperature from school days in the 365 days preceding and following a student's first PSAT take. All regressions include fixed effects for each high school, for each combination of gender, race, income and parental education, and for each combination of cohort and test date. Columns 4 and 5 control for temperature on the day of the first PSAT take. Column 5 includes state-specific linear time trends. The sample comprises all students from the high school classes of 2001-14 who took the PSAT at least once.

Table A.3: Temperature Effects by Test Subject

	Math (1)	Verbal (2)
(A) Average heat		
Mean temperature (°F)	-0.159***	-0.177***
	(0.034)	(0.023)
(B) Hot days		
Days above 90 °F	-0.042***	-0.062***
	(0.014)	(0.010)
N	21,046,448	21,046,448

Notes: Heteroskedasticity robust standard errors clustered by weather sensor are in parentheses (* $p<.10$ ** $p<.05$ *** $p<.01$). Coefficients in each column and panel come from a regression of hundredths of a standard deviation in PSAT math or reading scores on the weather measure shown. Panel A measures temperature with the daily maximum temperature from school days in the 365 days preceding the PSAT take. Panel B measures the number of such school days above 90 °F. All regressions include student fixed effects and fixed effects for each combination of cohort, test date and take number. Panel B also controls for the number of days in other temperature ranges, so that days in the 60s are the reference category. The sample comprises all students from the high school classes of 2001-14 who took the PSAT more than once.

Table A.4: Heterogeneity by Take Number

	(1)	(2)	(3)	(4)
Mean temp. (°F)		-0.152***	-0.200***	-0.269***
		(0.034)	(0.027)	(0.069)
Mean temp. * 1st take	-0.152***		0.048**	0.116
	(0.034)		(0.021)	(0.079)
Mean temp. * 2nd take	-0.200***	-0.048**		0.069
	(0.027)	(0.021)		(0.065)
Mean temp. * 3rd take	-0.269***	-0.116	-0.069	
	(0.069)	(0.079)	(0.065)	
N	21,046,448	21,046,448	21,046,448	21,046,448

Notes: Heteroskedasticity robust standard errors clustered by weather sensor are in parentheses (* p<.10 ** p<.05 *** p<.01). Coefficients in each column and panel come from a regression of hundredths of a standard deviation in PSAT total (math plus reading) scores on interactions between take number and the average daily maximum temperature from school days in the 365 days preceding the PSAT take. All regressions include student fixed effects and fixed effects for each combination of cohort, test date and take number. The sample comprises all students from the high school classes of 2001-14 who took the PSAT more than once.

Table A.5: Future Temperature Shocks

	(1)	(2)	(3)	(4)	(5)
Mean temp., 1 year prior (°F)	-0.181***	-0.229***	-0.182***	-0.178***	-0.228***
	(0.028)	(0.039)	(0.028)	(0.029)	(0.049)
Mean temp., 1 year after (°F)		-0.090*			-0.092
		(0.048)			(0.080)
Mean temp., 2 years after (°F)			0.053		-0.006
			(0.032)		(0.073)
Mean temp., 3 years after (°F)				-0.037	-0.037
				(0.033)	(0.051)
N	21,046,448	21,046,448	21,046,448	21,046,448	21,046,448

Notes: Heteroskedasticity robust standard errors clustered by weather sensor are in parentheses (* $p<.10$ ** $p<.05$ *** $p<.01$). Coefficients in each column and panel come from a regression of hundredths of a standard deviation in PSAT total (math plus reading) scores on the weather measure(s) shown. School day temperature is measured with the daily maximum temperature from school days in the listed 365 day period relative to the PSAT take. All regressions include student fixed effects and fixed effects for each combination of cohort, test date and take number. The sample comprises all students from the high school classes of 2001-14 who took the PSAT more than once.