

Hot Temperature, Human Capital and Adaptation to Climate Change

Jisung Park*

September 4, 2017

Abstract

Understanding the link between temperature and educational outcomes is important in assessing the welfare impacts of climate change, especially given the potential for persistent impacts on human capital accumulation and income growth. Using student-level administrative data for New York City public schools, I estimate the impact of hot temperature on high-stakes exams and subsequent educational outcomes. Hot days reduce performance by up to 15% and lead to lasting impacts on educational attainment. These effects persist despite the availability of cooling technologies and strategic teacher responses, suggesting that adaptation to climate change may be especially slow in educational settings.

Keywords: climate change, temperature, human capital, education, adaptation

JEL Codes: I21, O18, Q54, Q56

It is well-known that hotter countries tend to be poorer; a country that is 1°F warmer on average has roughly 4.5% lower GDP per capita. It is less well-known that hotter places tend to have substantially lower educational attainment. As shown in Figures 1a and 1b, there is a strong negative association between standardized academic performance and average temperature: both across countries as well as within the United States, even when controlling for per capita income.¹

*Park: Harvard University Economics Department, jisungpark@fas.harvard.edu. Acknowledgements: The author would like to thank Larry Katz, Andrei Shleifer, Robert Stavins, Joe Aldy, Geoffrey Heal, Raj Chetty, Claudia Goldin, Edward Glaeser, Melissa Dell, Michael Kremer, Josh Goodman, Jonah Rockoff, Jeff Miron, Max Auffhammer, Olivier Deschenes and numerous seminar participants at Harvard, Columbia, UC Berkeley, Oxford, IZA, Seoul National, the NYC Department of Health, the NBER Summer Institute and the Bill and Melinda Gates Foundation for valuable comments and feedback. Thanks also to the NYC Department of Education for data access, and to Nicolas Cerkez and Rodrigo Leal for excellent research assistance. All remaining errors are my own. Funding from the Harvard Environmental Economics Program, the National Science Foundation, and the Harvard Climate Change Solutions Fund is gratefully acknowledged.

¹Figure 1a graphs standardized math, reading, and science scores against average annual temperature for the 73 countries who participated in the Program for International Assessments (PISA). Figure 1b shows a similar relationship within the United States. It graphs average 3-8th grade math and reading performance for the near-universe of public school students from over 3,000 U.S. counties between 2009 to 2013, plotting residual variation after controlling for median income and average unemployment at the county level by percentile of the average climate distribution. Student performance data comes from Reardon et al. (2016).

What role – if any – does climate play in explaining this relationship? The central objective of this paper is to shed light on the ways hot temperature may affect the human capital production process.

Assessing the links between climate and human capital may be important on at least two dimensions. First, it may be important in estimating the magnitude of the carbon externality and the design of associated environmental policies, especially if hot temperature is shown to affect the rate of human capital accumulation. Current social cost of carbon estimates do not reflect potential human capital impacts from climate change (Tol, 2009), and assume climate will impact the level of output, not its rate of growth, despite the fact that a temperature-human capital link may lead to compounding growth effects over time (Pindyck, 2013). Second, to the extent that high-stakes exams often pose hurdles to further schooling (and since poorer people tend to live in hotter places), it is possible that temperature may contribute to longstanding achievement gaps. Differences in environmental conditions during exams may be important factors for policymakers to consider when deciding whether and how to administer high stakes exams, which are still common in many parts of the world.

In light of the empirical challenges associated with studying climate and human capital in the cross-section, I use student-level administrative data from New York City which, in conjunction with unique institutional features and quasi-experimental variation in temperature, allows for identification of the causal impact of hot temperature on educational outcomes. The research design is based on a simple premise: that short-run variations in temperature are not caused by unobserved determinants of educational performance. And whereas previous work on this subject (e.g. Graff Zivin et al. (2017), Cho (2017)) has been largely unable to account for avoidance behaviors, this study provides an analysis of the potential mechanisms by which economic agents may adapt, and sheds empirical insight on potential impediments to efficient adaptation.²

I focus on four empirical research questions. First, does acute heat stress meaningfully affect performance on high-stakes exams? That is, do the findings from survey and experimental contexts – wherein cognitive performance declines with elevated temperatures – extend to field settings where there are real economic consequences? Second, can short-run heat exposures, which presumably do not reduce the stock of human capital *per se*, nevertheless affect longer-run outcomes? Depending on the degree of institutional flexibility, the costs of retaking exams, and the presence of dynamic complementarities in the human capital production process, short instances of heat stress may have lasting economic consequences.³ Third, is it

²This builds on recent work suggesting that the realized impact of climate shocks may depend on the quality of institutions (Kahn, 2005; Dell et al., 2012), the adoption of air conditioning (Barreca et al., 2016; Park, 2016), and the possibility of moral hazard arising from government insurance schemes (Annan and Schlenker, 2015).

³Persistent impacts have been found in the context of air pollution (Ebenstein et al., 2016) and grade manipulation (Dee et al., 2016; Diamond and Persson, 2016). For instance, Ebenstein

possible for cumulative heat exposure to influence the rate of human capital accumulation, as suggested by recent findings (Cho, 2017)? Fourth, how do we expect students and teachers to adapt, and how might these adaptive responses interact with existing institutional environments?

The first main result is that heat exposure during a high-stakes exam exerts a causal and economically meaningful impact on student performance, even when controlling for individual ability. Taking an exam on a hot day leads to -0.22% lower performance per °F above room temperature (70°F); in other words, a 90°F day reduces exam performance by 15 percent of a standard deviation relative to a more optimal 70°F day (for a sense of magnitude, the within-school Black-White achievement gap is approximately 25 percent of a standard deviation). At least 18% of the students in my sample experience a hot exam with outdoor temperatures exceeding 90°F.

Looking at longer-run outcomes, I find that heat exposure during exams subsequently affects a student's chances of graduating from high school. For the median student, taking an exam on a 90°F day leads to a 10.9% lower likelihood of passing a subject. This means that a one standard deviation increase in average exam-time temperature reduces a student's likelihood of graduating on time by roughly 2.5 percentage points, despite the fact that students are able to retake failed exams. Heat exposure also substantially reduces chances of achieving key performance thresholds that are used by local universities in college admissions decisions. These results are consistent with a world in which a hot exam day nudges some students to achieve less schooling overall due to institutional rigidities and opportunity costs of time similar to those documented by Dee et al. (2016) and Ebenstein et al. (2016), or dynamic complementarities in the human capital investment process (Cunha and Heckman, 2007; Diamond and Persson, 2016).

I provide further but more speculative evidence on the relationship between cumulative heat exposure and human capital accumulation. Leveraging year-to-year variation in the number of hot days during the school year, I find that repeated heat exposure may reduce the rate of learning as evidenced by end-of-year exam performance – in addition to and controlling for the short-run impact documented above. A one standard deviation increase in the number of days above 80°F reduces Regents performance by approximately 3% of a standard deviation. Though this estimate is relatively under-powered, it is consistent with emerging findings (Cho, 2017; Garg et al., 2017). The effect is similar in magnitude to eliminating the gains associated with having a teacher with half a standard deviation higher value-added for one grade – an intervention which has been shown to increase cumulative lifetime incomes of NYC students by approximately \$14,800 per student, or \$445,000 per classroom (Chetty et al., 2014).

Two pieces of evidence highlight the important role played by institutions in

et al. (2016) find that air pollution exposure during high-stakes exams leads to lower post-secondary schooling attainment and reduced earnings.

determining the efficacy of adaptive responses to environmental shocks. First, building-level AC installation data suggests that public school air conditioning may be sub-optimally provided. Fewer than 62% of NYC public schools had any form of air conditioning as of 2012, compared to residential and commercial AC penetration rates of 90% or higher for the city on average. Comparing schools that do and do not have some form of AC, I find limited evidence for protective effects of having AC equipment present. One possibility is that, in the presence of informational asymmetries in education production, thermal comfort is perceived by administrators to be a consumption amenity as opposed to a production input and is thus inefficiently utilized from the social planner’s perspective.

Second, and possibly in response to inadequate air conditioning, teachers appear to have selectively manipulated grades upward when students experienced hot exam sittings. Previous studies have used bunching estimators to document grade manipulation by teachers (Dee et al., 2016). Following this approach, I estimate the relationship between grade manipulation and exogenous variation in exam-time temperature using a school, subject, and date-specific bunching estimator at passing cutoffs. I find that on average 6% of pre-reform Regents exams exhibit upward grade manipulation, and that the extent of manipulation varied systematically according to the temperature students experienced during the exam. This seems to have mitigated some of the adverse impacts on long-run educational outcomes. Such “adaptive grading” represents a hitherto undocumented and likely sub-optimal channel of adaptation to hot weather.

This paper is broadly related to a growing literature exploring the causal impact of climate on economic outcomes, including impacts of temperature shocks on human health (Barreca et al., 2016), labor supply (Graff Zivin and Neidell, 2014), violent crime (Kenrick and MacFarlane, 1986; Hsiang et al., 2013) and local economic output (Dell et al., 2012; Heal and Park, 2013; Deryugina and Hsiang, 2014), as well as the literature on climate adaptation (Deschênes and Greenstone, 2011; Burke et al., 2015). In particular, this paper shares with a smaller set of papers – including Graff Zivin et al. (2017), Ebenstein et al. (2016), Garg et al. (2017) and Cho (2017) – an emphasis on understanding the mechanisms through which environmental shocks such as temperature may affect student performance. In comparison to the existing literature, this paper establishes a causal connection between hot temperature and both short- and long-run educational outcomes in high-stakes environments, and does so in a context where adaptive responses can be studied in greater detail.

The findings suggest that classroom temperature may be an important factor for policymakers to consider when allocating public resources, especially in contexts where heat exposure is frequent, high-stakes exams pose hurdles to further schooling, and where market failures in infrastructure investment are likely to be present. This paper also provides empirical support for the view that climate and human capital may interact in a way that contributes to the long-debated relationship between hotter climates and slower growth, though more careful research is needed to verify

whether repeated heat exposure may reduce the rate of human capital accumulation in the long run. This is especially true given competing mechanisms in the existing literature, and the potential for endogenous innovation in adaptation technologies.

The rest of this paper is organized as follows. Section I provides a brief overview of the relevant literature on heat and human welfare and describes the underlying conceptual framework. Section II describes the data and institutional context and presents key summary statistics. Section III presents the main results and various robustness checks. Section IV discusses implications and concludes.

1 The Welfare Economics of Temperature Stress

Three stylized facts from the existing literature are of relevance in thinking about the impact of temperature on human capital production. First: heat stress directly affects physiology in ways that can be detrimental to cognitive performance. Second: most individuals demonstrate a revealed preference for mild temperatures close to room temperature, commonly taken to be between 65°F and 74°F, or 18°C and 23°C. Third: the inverted U-shaped relationship between temperature and performance documented in the lab has been confirmed in a range of field settings including mortality and labor supply, but evidence of impacts on human capital remains thin, particularly in high-stakes school environments.⁴

1.1 Heat Exposure and Human Capital

In seminal work, Graff Zivin et al. (2017) provide evidence of the adverse impacts of temperature stress on cognitive performance in US households. They use NLSY survey data which includes short, voluntary assessments that were administered to a sample students at home, and find evidence of contemporaneous impacts of hot days on math performance but not verbal performance, and little evidence of cumulative or persistent impacts over time.

Empirical evidence from school settings is limited mostly to short, voluntary cognitive assessments, where it is unclear that the incentives faced by students and teachers are strong enough to induce adaptive responses relevant in more high stakes settings or in the presence of persistently hotter temperatures. Schoer and Shaf-
fran (1973) compare the performance of students in a pair of classrooms set up as a temporary laboratory, with one classroom cooled and one not. They find higher performance in cooled environments relative to hot ones, but the assessments have no bearing on actual schooling outcomes. Garg et al. (2017) study the impact of

⁴See Heal and Park (2016) for a review of the related literature. See Mackworth (1946); Seppanen et al. (2006) on the physiology of heat exposure, Roback (1982); Sinha et al. (2015) for examples of hedonic analyses and the revealed preference for mild temperatures, and Grether (1973); Sudarshan and Tewari (2013); Graff Zivin and Neidell (2014) for impacts of temperature on labor and task productivity.

temperature on Indian students and find that hot days reduce cognitive performance, and that years with more days above 29°C reduce subsequent performance on voluntary cognitive assessments.⁵

Cho (2017) uses variation in summer-time temperature to predict college entrance exam performance by Korean high school students, and finds evidence that cumulative heat exposure reduces subsequent performance. While these exams are high stakes, they occur in November, meaning that one cannot assess adaptive responses or jointly examine the impact of hot temperature on contemporaneous performance and longer-run human capital outcomes. Moreover, the primary identification in Cho (2017) is driven by school fixed effects (student-level identifiers are not present in the data), making the results vulnerable to selection bias arising from changes in student composition as well as to endogenous changes in school AC status, which is not measured.

In related work using weather shocks as instruments for educational impacts, Goodman (2014) shows that snowfall can result in disruptions to learning by increasing absenteeism selectively across different student groups. Peet (2014) uses temperature, precipitation, and wind variation as instruments for pollution exposure in a sample of Indonesian cities and finds evidence of persistent impacts on student performance and labor market outcomes, though it is unclear to what extent temperature exerts a direct impact, and through what channels. A growing number of studies explores the impact of air pollution on student outcomes (Currie et al., 2009; Roth, 2016), and consistently find large impacts on absenteeism and exam performance. In the case of pollution during high-stakes exams in Israel (Ebenstein et al., 2016), there is evidence for persistent and economically meaningful impacts that extend well beyond formal schooling.

This study assesses the impact of temperature on human capital production by using evidence from high-stakes exams in New York City public schools. In contrast to the existing literature, it assesses contemporaneous and cumulative impacts of heat exposure simultaneously, and explores the role of adaptive investments and avoidance behaviors in greater detail. The richness of the administrative dataset permits causal identification using within-student variation in temperature, and allows linking of heat exposure to longer-run outcomes such as high school graduation rates and college eligibility. School-level air conditioning data and estimates of strategic teacher behavior permits an analysis of the potential extent of – and possible constraints to – efficient adaptation responses.

1.2 Conceptual Framework and Empirical Predictions

The basic conceptual framework is a Mincerian human capital model where the marginal value of student (and/or teacher) effort is a function of temperature, and

⁵Noting larger impacts for hot days during the growing season, they interpret the finding as evidence for malnutrition arising from reduced harvests, rather than cumulative impacts of heat exposure.

both exam score and cumulative human capital stock can be affected. The intuition is that students can invest time and effort e in order to accumulate human capital h . At any given point in time t , hot temperature, T_t (expressed as a deviation from optimum or room temperature), may affect the productivity of this investment $a(T)$, and may affect both a given exam score $s_t(T_t)$ as well as the overall amount of learning achieved over a relevant time period $h_T(\sum_t^T T_t)$.

The main empirical predictions from the model, which is described more formally in the appendix, are as follows:

1. We expect acute heat exposure for student i in time t to reduce contemporaneous exam performance, $\frac{\Delta s_{it}}{\Delta T_{it}} < 0$, if any of (a) direct flow utility, (b) marginal cost of effort, or (c) cognitive performance are adversely affected by temperature.
2. Short instances of heat exposure during an exam should not in principle reduce the stock of human capital h_{it} , but might nevertheless reduce overall human capital attainment $\frac{\Delta h_{it+\tau}}{\sum_t^{t+\tau} \Delta T_{it}} \leq 0$, if the schooling environment features high costs of retaking exams, or in the presence of dynamic complementarities due to policies such as tracking.
3. Repeated heat exposure may reduce human capital accumulation and educational attainment over time: $\frac{\Delta h_{it}}{\sum_0^{t-1} \Delta T_{it}} \leq 0$. This may be due to reduced student cognition or teacher performance during study or class time. They may also arise from reduced attendance, heat-related illnesses (e.g. heat stroke, food poisoning), or increased violent behavior at school or at home.
4. Any combination of (a) information asymmetries between students and school administrators or teachers and school administrators, (b) market failures in the provision of local public goods (e.g. collective action problems in providing adequate electrification), or (c) liquidity constraints in the context of air conditioning or other infrastructure investment, can lead to socially sub-optimal adaptive responses to repeated heat exposure in educational settings.

2 Institutional Context, Data, and Summary Statistics

2.1 New York City Public Schools

The New York City public school system (NYCPS) is the largest in the United States, with over 1 million students as of 2017. The median student is relatively low-performing and low-income, though a substantial minority attend high-achieving magnet schools including Stuyvesant Academy and Bronx Science, which consistently rank among the nation's best. Approximately one fifth of NYC students attend private schools which are not included in our sample.

The average 4-year graduation rate, at 68%, is below the national average but comparable to other large urban public school districts (e.g. Chicago). System-wide averages mask considerable discrepancies in achievement across neighborhoods. Schools in the predominantly Black or Hispanic neighborhoods of Brooklyn and the Bronx have four-year graduation rates as low as 35% per year.

2.2 New York State Regents Exams

Each June, students in the state of New York take a series of high-stakes exams called “Regents exams”. These standardized subject assessments are administered by the New York State Education Department (NYSED) and are used to determine high school diploma eligibility as well as college admissions.

Regents exams carry important consequences for the average NYCPS student. Students are required to meet minimal proficiency status – usually a scale score of 65 out of 100 – in five “core” subject areas to graduate from high school.⁶ Many local universities including City University of New York (CUNY) use strict Regents score cutoffs in the admissions process as well: for instance, requiring that students score above 75 on English and Math simply to apply.

The vast majority of students take their Regents exams during a pre-specified two-week window in mid-to-late June each year.⁷ The test dates, times, and locations for each of these Regents exams are determined over a year in advance by the NY State education authority (NYSED), and synchronized across schools in the NYCPS system to prevent cheating. Each exam is approximately 3 hours long, with morning and afternoon sessions each day, and are taken at the student’s home school.⁸

All exams are written by the same state-administered entity and scored on a 0-100 scale, with scaling determined by subject-specific rubrics provided by the NYSED in advance of the exams each year. All scores are therefore comparable across schools and students within years, and the scaling designed in such a way that is not intended to generate a curve based on realized scores.⁹ I use standardized

⁶The core subject areas are English, Mathematics, Science, U.S. History and Government, and Global History and Geography. The passing threshold is the same across all core subjects. Students with disabilities take separate RCT exams, and are evaluated on more lenient criteria.

⁷For any given student, exam takes are spread out across multiple days and years though, in effect, most exams are taken junior and senior year. Apart from the fact that most students take English their junior year, and Living Environment and Global History prior to other “advanced” sciences and US History respectively, there do not seem to be any clear patterns in the timing of subjects throughout students’ high school careers. Some advanced students may take Regents subject exams during middle school or during early January waves, while failing students are required to retake exams in August.

⁸Throughout the study period, students typically took Regents exams at the school in which they were enrolled unless they required special accommodations which were not available at their home school. Students who fail their exams (or are deemed unready by their teachers to progress to the next grade) are required to attend summer school, which occurs in July and August.

⁹In principle, scores are comparable across years as well, as psychometricians in the NYSED

performance at the subject level as the primary measure of exam performance in this study, though the results are robust to using scale scores. While centrally administered, exams were locally graded by teachers in the students' home schools until grading reforms were implemented in 2011 in response to a series of reports suggesting grade manipulation (Dee et al., 2016).

In summary, using scores from Regents exams offers several distinct advantages. First, they are high-stakes exams used to determine diploma eligibility and possibly affecting college enrollment. This means that, in addition to direct welfare relevance, they may also provide information about compensating behavior that is not available in laboratory studies or voluntary cognitive assessments such as those in the NLSY or in Garg et al. (2017). Second, they are offered at a time of year when temperatures fluctuate considerably, allowing for a wide range of quasi-experimental variation within a relatively small geographic locale. Because they occur at the end of the school year, they are also more likely than periodic assessments or college entrance exams to reflect cumulative impacts of hot temperature that may have accrued over the course of the school year. Finally, Regents exams were centrally administered and compulsory for all public school students during the study period, meaning there is relatively little possibility of anticipatory alteration of exam timing based on weather forecasts, or for bias due to selection into taking the exam.

2.3 Student Outcome Data

I obtain student-level information from the New York City Department of Education (NYC DOE). This includes the universe of all public school students who took one or more Regents exams over the period 1999 to 2011.¹⁰ Information on exam dates comes from archived Regents exam schedules from the NYC DOE database. A detailed description of the matching procedure and subsequent sample restrictions can be found in the online appendix.

2.4 Weather Data

Weather data comes from NOAA, which provides daily min, max, and mean temperatures, precipitation and dew point information from a national network of several thousand weather stations over the period 1950-2014. I take daily minimum and maximum temperature as well as daily average precipitation and dewpoint readings

conduct difficulty assessments of each year's subject exams and engage in "equating" procedures prior to their release. The primary identification of short-run impacts include year fixed effects, and thus do not rely on this cross-year comparability.

¹⁰I also use data from standardized math and English language and arts (ELA) exams administered in 3rd through 8th grade from NYC DOE to provide a measure of previous ability. Specifically, I calculate the average combined z-score of each student for whom previous standardized ELA and math exam records are available. For students who are missing these records, I assign imputed average z-scores by decile of the realized Regents score distribution. These exams are generally taken in January and March.

from the 5 official weather stations in the NYC area that were available for the entirety of the sample period (1998-2011). I match schools to the nearest weather station, one for each of the five boroughs: The Bronx, Brooklyn, Manhattan, Queens, Staten Island.¹¹ Given existing evidence on the impact on air quality on student performance, I include controls for pm2.5 and ozone, taken from EPA monitoring data from Manhattan.¹²

2.5 School Air Conditioning Information

Information on building-level air conditioning equipment comes from records originally compiled by New York City School Construction Authority (SCA), which administers detailed, building-level surveys for NYC public schools. While a centralized database was not publicly available, a web-scrape of individual school websites resulted in matchable records for 644 middle and high school buildings in the study sample. The records include information on air conditioning equipment presence and maintenance status as of the year 2012. Unfortunately, the data does not provide AC installation or usage status by year, nor does it provide information regarding where *within* a school AC was available. As such, I take AC installation status as a rough proxy for the true variable of interest, which is effective AC utilization at the classroom level.

2.6 Summary Statistics

The final working dataset consists of 4,509,102 exam records for 999,582 students. It includes data from 91 different exam sessions pertaining to the core Regents subjects over the 13 year period spanning the 1998-1999 to 2010-2011 school years.

Table 3b presents summary statistics for the key outcome variables that form the basis of this analysis. The student body is 40% Latino, 31% African American, 14% Asian and 13% White, and approximately 78% of students qualify for federally subsidized school lunch. The median student scores just around the passing cutoff, with a score of 66 (sd = 17.9), though there is considerable heterogeneity by neighborhood as well as demographic group.

¹¹To account for spatial heterogeneity in outdoor temperature due to urban heat island effects, I also assign spatial correction factors generated by satellite reanalysis data. I impute test-time temperature – for instance, average outdoor temperature between 9:15am to 12:15pm for morning exams – by fitting a fourth-order polynomial in hourly temperature. Further details regarding these corrections are presented in the online appendix. The primary results reported below are not sensitive to either of these corrections. The corrections reduce standard errors but leave implied point estimates relatively unchanged.

¹²The air quality controls used here are admittedly crude, especially for localized pollutants such as ozone. Given the focus of the study, the relatively low levels of particulate matter during the sample period, and the high correlation between ozone and summertime temperature, I run analyses with and without controls for air quality but cannot separately identify the precise effects of pm, ozone, and temperature simultaneously.

Students take on average 7 June Regents exams over the course of their high school careers, and are observed in the Regents data set for roughly 2 years, though some under-achieving students are observed for more than 4 years, as they continue to retake exams upon failing. Fewer than 0.2% of students are marked as having been absent on the day of the exam, corroborating the high-stakes, compulsory nature of these exams.

Figure 2 illustrates the source of identifying variation for short-run temperature impacts, with temperatures weighted by exam observation and school location. Outdoor temperature during exams range from a low of 60°F to a high of 98°F.¹³

Cumulative heat exposure during the school year can vary considerably, as suggested by Figure 3a which presents the incidence of days with maximum temperatures above 80°F by school year and borough. On average, NYC students experience between 19 and 39 days above 80°F per school year, with a mean value of 26.7 and a standard deviation of 5.6.¹⁴ In addition, there are on average 2.5 days during the school year above 90°F. Most of these days occur during the months of September, October, and June. Summer school students, which comprise roughly the bottom third of the student body but are not included in this analysis due to likely selection bias, are subject to an additional 9 days above 90°F on average.

In 2012, 62% of the NYC public school buildings for which I have building assessment data were reported as having some kind of air conditioning equipment on its premises, including window units, which means that 38% of these buildings did not have any form of air conditioning equipment available (Figure 7a). Of the 62% that were reported as having air conditioning, 42% (259 out of 644) were cited as having defective components, according to the third-party auditors conducting the BCAS assessments.

3 Empirical Strategy and Primary Results

To organize the empirical findings, I report the results by research question.

1. Does Exam-Time Heat Exposure Affect Exam Performance?

Figure 4a presents a visual depiction of performance and temperature that motivates the analysis that follows. It shows a binned scatterplot of standardized exam score by percentile of observed exam-day temperature, plotting residual variation after

¹³Day-to-day variation within the June exam period can be considerable, as suggested by Figure 2b, which shows the variation in outdoor temperature by school and exam take across two consecutive test dates within the sample period.

¹⁴Despite documented warming for the US and the world as a whole over the past several decades, temperatures in the NYC area have remained relatively stable over the study period. Tests for stationarity and trend-stationarity do not suggest time trends in these extreme heat day variables.

controlling for school fixed effects and average differences across subjects and years. It suggests that exams taken on hot days exhibit lower scores.

To further isolate the causal impact of short-run temperature fluctuations on student performance, I exploit quasi-random variation in day-to-day temperature across days within student-month-year cells, focusing on the main testing period in June. While it is unlikely that temperature is endogenous to student behavior, nor is it likely for students to select into different temperature treatments given the rigidity of exam schedules, time-varying unobservables may still be correlated with weather realizations. For instance, if certain subjects tend to be scheduled more often in the afternoon when students are relatively fatigued, as in Sievertsen et al. (2016), or toward the end of the exam period (Thursday as opposed to Monday), we may expect mechanical correlation between temperature and test scores that is unrelated to the causal effect of temperature on student cognition or effort. This motivates a baseline specification that includes year, time of day, and day of week fixed effects:

$$Y_{ijsty} = \gamma_{iy} + \eta_s + \beta_1 T_{jsty} + X_{jsty} \beta_2 + \beta_3 Time_{sty} + DOW_{sty} \beta_4 + \epsilon_{ijsty} \quad (1)$$

Here, Y_{ijsty} denotes standardized exam performance for student i taking an exam in subject s in school j on date t in year y . The terms γ_{iy} and η_s denote student-by-year and subject fixed effects respectively. T_{jsty} is the outdoor temperature in the vicinity of school j during the exam (subject s on date t , year y). X_{jsty} is a school- and date-specific vector of weather and air quality controls, which include precipitation, dewpoint, and ozone. $Time_{sty}$ represents a dummy for time of day (morning versus afternoon, $Time=1$ denotes an afternoon exam), and DOW_{sty} represents a vector of fixed effects for each day of the week in which exams were taken.

Student-by-year fixed effects ensure that we are comparing the performance of the same student across different exam sittings within the same testing window, some of which may be taken on hot days, others not, leveraging the fact that the average student takes 7 June Regents exams over the course of their high school career (on average between 3 and 4 per year). Subject fixed effects control for persistent differences in average difficulty across subjects. Year fixed effects control for possible spurious correlation between secular performance improvements and likelihood of hotter exam days due, for instance, to climate change.¹⁵

¹⁵To the extent that temperature variation within student-month-year cells are uncorrelated with unobserved factors influencing test performance, one would expect the coefficient β_1 to provide an unbiased estimate of the causal impact of temperature on exam performance, subject to attenuation bias due to measurement error in weather variables as well as downward bias from positive grade manipulation. It is worth noting that it is possible for measurement error to be non-classical in a way that biases the estimates upward. For instance, if the average classroom has more students in lower performing schools, experienced classroom temperature scales non-linearly with outdoor temperature, and students in lower performing schools are more susceptible to heat stress, then β_1 may actually be biased upwards.

Table 4b presents the results from running variations of equation (4) for the subset of students who take at least 2 exams in any given year.¹⁶ As suggested by the first row of columns (1)-(4), exam-time heat stress exerts a significant causal impact on student performance. The estimates are robust to allowing for arbitrary autocorrelation of error terms within boroughs and test dates, which is the level of exogenous temperature shock recorded in the data, as well as to using the full sample of scores: i.e. retaining all students with 1 or more exam records and replacing student fixed effects with school fixed effects and a vector of demographic controls.

Taking an exam on a hot day reduces performance by approximately -0.008 standard deviations (se=0.002) per °F. This amounts to -5.2 percent of a standard deviation in performance per standard deviation increase in temperature, or -14.7 percent of a standard deviation if a student takes an exam on a 90°F day as opposed to a more optimal 70°F day.¹⁷

The effect of a 90°F day is thus comparable in magnitude to roughly 1/4 of the Black-White score gap, or 3/4 of the within-school Black-White score gap. This effect is comparable in raw magnitude to the impacts on mathematical reasoning found by Graff Zivin et al. (2017), who find a 90°F day to reduce NLSY math scores by approximately -0.12 standard deviations.

A series of robustness checks, including models that replace student-by-year fixed effects with student- or school-by-year fixed effects, are presented in the online appendix. The point estimates using the school-by-year fixed effects specification are slightly larger (more negative) on average, and remain statistically significant. Also presented in the appendix are heterogeneity analyses by gender and ethnicity. I find relatively little evidence of heterogeneity in short-run impacts, though it is possible that adaptive responses by teachers are offsetting impacts disproportionately for certain subgroups (as discussed in greater detail below).

These results provide strong evidence that heat stress affects student performance, either by reducing raw cognitive ability or by increasing the disutility of effort which in turn affects students' willingness or ability to maintain focus during a three-hour exam. They suggest that temperature in the learning environment plays an important role in determining student outcomes, and that whatever com-

¹⁶Results using un-standardized scale scores as the dependent variable are presented in the appendix. In models where student fixed effects are replaced by school fixed effects, a vector of demographic control variables are included to control for possible selection into subjects by student type.

¹⁷Precipitation has a slightly positive effect, and ozone has a negative but insignificant effect, with a 1 standard deviation increase in ozone corresponding to a point estimate roughly 1/5th the size of a 1 standard deviation temperature effect. Despite previous literature documenting adverse impacts of pm2.5 in Israel (Ebenstein et al., 2016), I find little evidence for that here, perhaps because average concentrations of pm2.5 are much lower in NYC than in Israel, as well as the fact that the performance impacts documented by Ebenstein et al. (2016) are highly non-linear, driven mostly by heavily polluted days with pm2.5 above 100 micrograms per cubic meter. The maximum recorded value of pm2.5 in my data is 38.8 micrograms per cubic meter.

pensatory effort is exerted by students due to the high stakes nature of some exams may not be enough to offset the physiological impacts of temperature documented in the lab. However, these results are not in themselves directly informative of welfare impacts or effects on the stock of human capital, which is explored in the next section.

2. Does Short-run Heat Exposure Affect Longer-run Educational Outcomes?

Heat exposure during an exam, while reducing cognitive ability or concentration temporarily, presumably does not reduce the stock of knowledge or human capital *per se* – at least not immediately through the physiological impact of heat stress itself.

In a stylized, friction-less world with fully flexible educational institutions, unlucky students who fail a subject due to a hot exam sitting would immediately retake the exam until she believes her “true ability” has been reflected in the exam score: $s_{it} = h_{it}$. In this world, random heat exposure during exams should not affect the final amount of schooling achieved.

However, in the presence of institutional rigidities that limit the effective number of possible retakes or impose time and effort costs to retaking an exam (by, for instance, requiring students to attend remedial courses), it is possible that even short-run heat exposure can lead to ripple effects on long-run educational attainment. In addition, employers may treat students who graduated from high school in five or six years differently from those who graduated “on-time” for a variety of reasons, reducing the pecuniary return to education once a subject has been failed the first time. Similarly, exam scores may serve as important signals within the education system – to the student herself, to her peers, or to her parents and teachers – leading to dynamic complementarities in human capital investment (Cunha and Heckman, 2007; Diamond and Persson, 2016).

Recent evidence suggests persistent effects of temporary score shocks in the context of teacher manipulation (Dee et al., 2016; Diamond and Persson, 2016) and air pollution (Ebenstein et al., 2016), with as yet inconclusive evidence regarding the specific mechanisms by which they occur.¹⁸

Pass Rates and College Proficiency

¹⁸Dee et al. (2016) find substantial impacts of upward score manipulations on graduation status, especially for students who scored in the manipulable zone. Using administrative records from Swedish middle schools, Diamond and Persson (2016) also find substantial effects of upward score manipulations on subsequent performance, graduation likelihood, and later life income. Ebenstein et al. (2016) find that Israeli high school students who receive lower scores on their Bagrut (high school exit) exams due to air pollution are less likely to receive Bagrut certificates (comparable to high school diplomas) and receive lower wages later in life.

If heat exposure during an exam pushes some students below important (cardinal) score thresholds that affect access to further educational opportunities, one might expect even short “doses” of heat exposure to give rise to lasting impacts on educational attainment.

In NYC, students must score a 65 or above to pass a given subject exam and thus have it count toward receiving a high school diploma. This cutoff does not change based on the realized distribution of performance in any given year.¹⁹ Students are also assigned “proficient” (i.e. college ready) status on each of the subjects in which they receive a grade of 75 or higher and “mastery” status for scores of 85 or higher. Beyond any personal motivational or within-school signalling value, these designations carry real weight externally in the sense that many local colleges and universities such as City University of New York (CUNY) use strict score cutoffs in their admissions decisions.²⁰

To estimate the impact of contemporaneous heat stress on the likelihood that a student scores at or above the passing and proficiency thresholds, I run variations of the following models:

$$p_{ijsty} = \gamma_{iy} + \eta_s + \beta_1 T_{jsty} + X_{jsty} \beta_2 + \beta_3 Time_{sty} + DOW_{sty} \beta_4 + \epsilon_{ijsty} \quad (2)$$

$$c_{ijsty} = \gamma_{iy} + \eta_s + \beta_1 T_{jsty} + X_{jsty} \beta_2 + \beta_3 Time_{sty} + DOW_{sty} \beta_4 + \epsilon_{ijsty} \quad (3)$$

where p_{ijsty} is a dummy variable indicating whether student i passed – that is, scored a 65 or above on – subject s on date t , year y , and c_{ijsty} is a dummy variable indicating college proficiency status: i.e., a dummy for scores at or above 75 points.

The results suggest that acute heat exposure can have significant short term impacts on student performance, with potentially lasting consequences. Exam-time heat stress reduces the likelihood of passing by 0.31 (se=0.12) percentage points per °F, or -0.54% per °F from a mean likelihood of 0.57.²¹ Impacts on the likelihood of achieving proficiency status are slightly larger in aggregate, with a magnitude of -0.96% per °F hotter exam-time temperatures relative to a mean likelihood of 0.32. Unless higher-ability students are more sensitive to heat stress, this discrepancy seems to be driven in part by grade manipulation around the passing threshold.

¹⁹Until 2005, low-performing students were allowed the option of applying to receive a “local diploma” which required scores of 55 and above for exams to count toward the diploma. In the following regressions, I use the more stringent and universally accepted standard of “Regents Diploma” as the definition of passing score, as do Dee et al. (2016). Results of running the regression analyses below using the “Local Diploma” cutoff feature similar (slightly more negative) point estimates.

²⁰The scale score needed to be considered “college ready” differs by subject. For instance, for admission to CUNY schools, a student can demonstrate the necessary skill levels in reading and writing by meeting any of the following criteria: SAT Critical Reading score of 480 or higher; ACT English score of 20 or higher; N.Y. State English Regents score of 75 or higher.

²¹These results are presented in tabular form in the appendix.

Taken together, these estimates suggest that experiencing hot ambient temperatures during a high stakes exam can have non-trivial consequences for student performance, with a 90° day leading to approximately 9.7% lower chance of passing a given exam, and a 17.4% lower probability of achieving proficiency status for the average student.

Graduation Status

Perhaps more importantly, I find evidence that these short-run heat exposures reduce final educational attainment by affecting the likelihood that students graduate from high school. Figure 5a presents a binned scatterplot of 4-year graduation status on average exam-time temperature by student during June Regents exams up through the student’s senior year. It plots residual variation in a dummy variable denoting graduation status, controlling for school-level averages, student-level demographic characteristics, and the number of June Regents exams taken by student. It suggests that students who experience greater exam-time heat stress are less likely to graduate on time.

Whereas short-run impacts of heat stress could be identified *within* student cells, long-run impacts on graduation status cannot because the outcome variable is no longer date-specific. This poses additional challenges to causal identification. Computing a measure of average heat exposure across multiple exam sittings by student results in mechanical correlation between average experienced temperature and the number of exams such that students who take more exams are more likely to be assigned average temperature values closer to the climatic mean in that month, and students who take fewer exams are more likely to be assigned extreme values.²²

The comparison of interest is the difference in graduation likelihood between students who, conditional on the number of draws from the climate distribution, experience different amounts of heat stress. One way to accomplish this is to compare within exam-count and year-count cells, controlling for observable factors. Collapsing the data to the student level, I estimate variations of the following model:

$$g_{ijcn} = \alpha_0 + \alpha_1 \overline{T_{ij}} + X_{ij} \alpha_2 + \chi_j + \theta_c + Z_i \alpha_3 + exams_n \alpha_4 + \epsilon_{ijc} \quad (4)$$

Here, g_{ijcn} is a dummy denoting whether student i in school j and entering cohort c who takes n June Regents exams over the course of her high school career has graduated after 4 years in high school. $\overline{T_{ij}}$ denotes the average temperature experienced by student i while taking June Regents exams in school j , up through her senior year. X_{ij} is a vector of weather controls averaged at the student-by-school level. χ_j denotes school fixed effects; θ_c denotes cohort fixed

²²That is, assuming that the average June climate in New York City can be represented by a distribution of daily temperature realizations, the average temperature across multiple days will exhibit a form of mean-reversion as one increases the number of draws from the underlying climate distribution (see online appendix for a visual representation).

effects; Z_i is a vector of student-level controls including race, gender, federally subsidized school lunch eligibility, and previous ability (combined ELA and math z-scores); and $exams_n$ denotes a vector of fixed effects for the number of June exam takes.

The parameter of interest is α_1 , which captures the impact of an additional degree of heat exposure during exams on the likelihood of graduating on time.²³ School fixed effects account for potential omitted variable bias due to unobserved determinants of graduation rates being correlated with average temperature in the cross-section (e.g. if urban heat island effects are stronger in poorer neighborhoods). Cohort fixed effects in graduation rates allow for the possibility that heat exposure and graduation rates are correlated due to secular trends in both variables – though warming trends and average improvements in NYC schools would suggest this effect to lead to downward rather than upward bias in the estimate of α_1 .

Table 5a presents the results from running variations of equation 4 with and without school and cohort fixed effects, as well as flexible controls for the number of exams. Standard errors are clustered at the borough by date and time level, based on the intuition that this conservatively approximates the level of quasi-random temperature variation, though the results are once again robust to alternative levels of clustering.

Columns (1)-(3) suggest that a 1 degree F increase in average exam-time temperatures is associated with a 0.71 (se=0.17) to 0.76 (se=0.22) percentage point decline in the likelihood of graduating on time. A one standard deviation in average exam-time temperature (+4.4°F) leads to a 3.12 to 3.34 percentage point decline in the likelihood of on-time graduation, or 4.59% to 4.91% decline relative to a mean on-time graduation rate of 68 percent.

These effects are large. Even without correcting for adaptive grading by teachers, I estimate that, over the period 1998 to 2011, upwards of 510,000 exams that otherwise would have passed received failing grades, affecting the on-time graduation prospects of at least 90,000 students.²⁴

²³The intuition is that variation in experienced temperature among students in the same school and cohort will be plausibly uncorrelated with residual variation in graduation status within school and cohort cells. Suppose there are two students, Jill and Karen, who entered high school in 2000. In 2001, because of differences in the sequence of subjects that Jill and Karen took, Jill takes Regents exams on Monday, Wednesday, and Thursday, and Karen takes Regents exams on Monday, Tuesday, and Friday. Suppose a similar phenomenon occurs during their sophomore, junior, and senior years, such that they take the same overall number of June exams. The variation in overall experienced temperature between Karen and Jill in 2001 will likely be exogenous to any unobserved differences in Jill and Karen’s likelihood of graduating from high school.

²⁴See online appendix for a description of the methodology used in calculating the number of students and exams affected. These figures do not explicitly account for upward grade manipulation, and therefore likely represent a lower bound.

3. Does Cumulative Heat Exposure Reduce the Rate of Learning?

Given the effects documented above, and the vastly uneven distribution of hot days across countries, one might be interested in understanding the extent to which repeated heat exposure over time may reduce the effectiveness of learning and thus also affect the rate of human capital accumulation over time. Emerging evidence (Cho, 2017; Garg et al., 2017) suggests that cumulative heat exposure can hinder skill formation – or at least exam preparation – though the exact mechanism(s) through which this may occur is as yet unclear.

Figure 6a presents a binned scatterplot of Regents score on the number of days with max temperatures between 80°F and 90°F, controlling for exam-day temperature and precipitation, as well as school-, subject- and time of day fixed effects. The figure suggests that hot days may be reducing learning attainment, at least as reflected in end-of-year exam performance.

Because Regents exams are subject-specific and usually administered at the end of the school year during which that subject was taken, they provide a suitable opportunity for uncovering potential cumulative learning impacts of heat exposure during the school year. On the other hand, because each subject exam is usually only taken once per year and observed over the course of 13 years in my data set, and because cross-sectional variation in heat exposure within New York City is relatively limited, the analysis is likely to exhibit limited precision compared to the estimates of short-run exam-day impacts.

To identify the impact of cumulative heat exposure on learning, I collapse the data to the school by subject and month (year) level.²⁵ I retain subject-level variation in order to estimate the impact of cumulative heat exposure while controlling for the short-run impacts of contemporaneous heat stress documented above. I estimate variations of the following model:

$$y_{jsty} = \beta_0 + \beta_1 T_{jsty} + X_{jsty} \beta_2 + \sum_d^D \beta_3^d DD_{jT}^d + \chi_j + \eta_s + Z_{jsty} \beta_4 + \beta_5 Time_{sty} + DOW_{sty} \beta_6 + f(Year_y) + \epsilon_{jsty} \quad (5)$$

where y_{jsty} denotes the average Regents z-score (standardized once again by subject over the study period) for students in school j taking subject s on date and time t , during year y ; T_{jsty} denotes exam-time outdoor temperature at school j for subject s on date and time t , during year y ; χ_j denotes school fixed effects; η_s denotes subject fixed effects; and Z_{jsty} represents a vector of demographic controls averaged at the school by take (subject-month-year) level. $Time_{sty}$ represents a dummy for time of day, and DOW_{sty} is a vector of fixed effects for day of week.

²⁵This is for a couple of reasons. Because I do not observe student-specific measures of cumulative heat exposure over the preceding school year, keeping student-level exam observations will likely introduce additional measurement error, since cumulative heat exposure during preceding school years is measured at the school level and some students may have been present for more days than others or live in neighborhoods that are more prone to heat stress than others.

$f(Year_y)$ denotes a non-linear (cubic) time trend in scores. Time trends are included in lieu of year fixed effects to account for possible secular changes in performance over time that may be spuriously correlated with shifts in climate over the study period. The results presented below are robust to including school-specific trends as well as subject-specific trends.

The variable “ DD_{jy} ” denotes a vector of day counts in a series of degree-day bins during the preceding school-year (y), beginning with the first day of the fall semester up to the first day of the testing period the following June.²⁶ I use a number of bin classifications for hot days, motivated by the existing literature (e.g. Barreca et al. (2016)) as well as the analyses presented in the previous section, which find negative impacts of heat stress beginning around 72°F. The preferred analysis flexibly divides temperature days into 10 degree bins, beginning with 10°F to 20°F up to 90°F and above, omitting the “optimal” bin, which the data and previous work using similar approaches – e.g. Deschênes and Greenstone (2011) – suggests to be around 60°-70°F. The coefficients of interest correspond to the “hot” degree day bins, around or above 80°F, and represent the correlation between the number of hot days in a school year and end-of-year exam scores. X_{jsty} denotes a vector of contemporaneous and cumulative weather controls, including precipitation and dewpoint on exam day as well as during the preceding school year, and annual snowfall, which is taken from weather station readings in Central Park and assigned uniformly across all schools in the city.

I run variations of equation 5 that allow for a flexible characterization of the reference category against which we can interpret the impact of hot days. Table 6b presents the results from these analyses, with columns (1), (2), (3), and (4) corresponding to specifications that control for (1) hot (70-80°F and above) days only, (2) hot and cold (30-40°F and below) days only, (3) hot days and average daily maximum temperature over the school year, and (4) all degree day bins from 0°F to 90°F omitting the 60-70°F bin respectively.

The results are suggestive of cumulative learning impacts due to heat exposure during the school year. First, note that the short-run impacts persist in all specifications, with relatively stable point estimates of similar magnitude from the results presented above. Focusing on column (3), which controls for average daily max temperatures during the school year as well as for the contemporaneous effect of exam-time temperature, days between 80° and 90°F have a negative impact of between -0.011 (se=0.0031) and -0.012 standard deviations (se=0.0033), or approximately 1% of a standard deviation per hot day. Estimates for days above 90°F are noisier given the relatively limited number of such days during term. Results in columns (1), (2) and (4) suggest a similar pattern of hot days during the preceding

²⁶School years are defined such that the test year corresponds to the year in which the spring semester of the academic year occurs. For instance, $y=2000$ corresponds to the 1999-2000 school year; $y=2001$ to the 2000-2001 school year, etc.

school year reducing exam performance.²⁷

These estimates imply that a one standard deviation (3.91 day) increase in the number of days with maximum temperatures above 80°F can reduce learning by approximately 0.04 standard deviations, as measured by end-of-year exam performance. These impacts are on par with the learning impacts of a 0.4 standard deviation reduction in average teacher value-added (Chetty et al., 2014), or 1/2 of the impact of reducing class size from 31 to 25 (Angrist and Lavy, 1999). Though data limitations do not permit the analysis of the impacts on later-life outcomes such as wages or health directly, these results should be interpreted in light of recent analyses (Chetty et al., 2014) which examine the same population of NYC students and find significant impacts of improved learning on later-life outcomes.

4. How do agents adaptively respond?

If heat stress affects student performance in economically meaningful ways, standard micro theory would predict that students, parents, and teachers would respond to mitigate this impact, presumably along the most cost-effective margins. However, efficient adaptation to environmental shocks may be contingent on a number of factors, including institutional quality or the effective costs of (and potential barriers to the adoption of) cooling equipment.

For instance, the quality of existing institutions has been found to affect the severity of mortality responses to environmental disasters such as earthquakes (Kahn, 2005), and the adoption of heat-resistant crops by farmers in the US (Annan and Schlenker, 2015). In the short run, flexible institutions and smart policies may help to minimize the realized welfare impacts of a string of hot days or a particularly hot year. In the long run, one might expect institutions to mediate the responsiveness of school systems to a changing climate, with some adapting more quickly than others. Rigid institutions or bureaucratic red-tape may hinder such responses.

In the context of energy infrastructure, market failures arising from imperfect information, liquidity constraints, or principal agent problems may drive a wedge between realized and optimal adaptive investments. Gertler et al. (2016) find evidence for liquidity constraints in the adoption of energy-intensive appliances, and a long literature has noted the potential for principal agent problems in the context of renter-occupied homes (Allcott and Greenstone, 2012). It is possible that such market failures are more prevalent in school environments, where the link between educational inputs and market outcomes is imperfectly observable and often involves long lags. The fact that many US public school districts and teachers unions have clashed over installation of air conditioning suggests possible information problems associated with the efficacy of various adaptation options. For instance, during a

²⁷Cold days appear to have a negative impact on end-of-year performance as well, particularly in the case of days with maximum temperatures between 30 and 40 °F. This is consistent with previous work by Goodman (2014) who finds that the number of snow days and the amount of local snowfall adversely predict end-of-year performance in Massachusetts public schools.

major teacher union strike in Chicago in 2012, “Timetable for air conditioning” was listed as one of four major contract demands. An agreement to provide universal air conditioning in Chicago public schools was not reached until 2016.

This section explores two types of adaptive responses as a window into the ways in which institutional factors may or may not constrain the realization of first-best adaptation investments: air conditioning and teacher responses.

Air Conditioning

Air conditioning is a potential adaptation strategy which, at least in the context of health and labor productivity, has proven to be effective at mitigating the realized health and productivity impacts of hot temperature (Barreca et al., 2016; Park, 2016).

Average air conditioning penetration for commercial and residential buildings in New York City is around 90%, suggesting that historical rates of heat exposure may be frequent enough to merit investment in AC in the region. As noted above, however, fewer than 62% of NYC public schools reported having some form of AC equipment as of 2012. Among those that do, a substantial proportion are reported as having defective system components. This AC penetration gap – between public schools and residential and commercial buildings – in the same region is consistent with though by no means necessarily indicative of the presence of institutional or other barriers to efficient adaptation. It is possible that other unobserved factors (e.g. building age, occupancy rates) give rise to this discrepancy.²⁸

To compare the impact of heat exposure on high-stakes exams across schools with and without AC units, I estimate equation 1 separately for sub-groups of students who took exams in schools with and without air conditioning as of 2012. The results from these regressions are reported in Table 7b. Column (1) reproduces the main effect on the entire sample. Columns (2), (3) and (4), (5) present results for sub-groups with central AC, with any AC, without central AC, and without any AC respectively. The point estimates are smaller and insignificant for sub-samples with AC units, -0.0053 (se=0.0029) and -0.00517 (se=0.0027), relative to sub-samples without AC: -0.0065 (se=0.0027) and -0.0062 (se=0.0026), for schools with and without central AC or any form of AC respectively. These differences are not statistically different however, suggesting either that existing AC units have an only mildly protective effect (perhaps because they are not always in operation), or that AC status is measured with substantial noise. It is also possible that the relatively old building stock in NYC places additional constraints to the effectiveness of some AC systems. More research is needed to understand whether and to what

²⁸Summer vacation means that occupancy may be lower for schools during hotter months, making this comparison less informative. However, it is worth noting that a substantial fraction (often 25% or more) of public school students are required to attend summer school due to low achievement status.

extent school air conditioning represents an effective adaptation response to climate change, and what if any market failures may prevent efficient uptake.

Teacher Responses

Using a similar dataset from 2003 to 2012, Dee et al. (2016) document systematic grade manipulation by NYC teachers on NY State Regents exams. They find that most of the manipulating behavior occurred at or around the passing margin of 65 and that, while varied in magnitude across schools and student types, such manipulation was a near-universal phenomenon within the NYCPS system.

Upon careful analysis of competing explanations, the authors suggest the most likely explanation to be the goodwill of teachers who seek to offset the impact of “a bad test day”:

In sum, these estimates suggest that manipulation was unrelated to the incentives created by school accountability systems, formal teacher incentive pay programs, or concerns about high school graduation. Instead, it seems that the manipulation of test scores may have simply been a widespread “cultural norm” among New York high schools, in which students were often spared any sanctions involved with failing exams, including retaking the test or being ineligible for a more advanced high school diploma (pg 27).

A hot test day may be viewed as a bad test day, particularly if air conditioning is inadequately provided. In that case, it seems possible for discretionary grade manipulation to have been a response to perceived performance impacts of heat stress. Teachers may be able to observe or at least intuit the disruptive impacts of elevated classroom temperatures on test day, especially since exams are taken in students’ home schools and graded by a committee of teachers from that school. If benevolently motivated (as Dee et al suggest), they may be inclined to engage in *more* grade manipulation precisely for those exams that took place under disruptively or unusually hot conditions. One might call this selective response by teachers “adaptive grading”, a second-best adaptation strategy undertaken in the presence of institutional barriers to more efficient forms of structural adaptation.

Estimating Adaptive Grading: Bunching at Score Thresholds

Figure 8a provides a histogram of Regents scale scores in all core subjects prior to 2011. As is clearly visible in the graph, there is substantial bunching at the passing kinks, especially at scores of 65 and 55.

We would expect any form of grade manipulation for students who initially score just below the passing cutoff, even “indiscriminate” grade manipulation uncorrelated with exam-time temperature, to downward attenuate the estimates of heat-related performance impacts uncovered above.²⁹ Indeed, running equation 1 on the subset of

²⁹The only case in which grading process may affect our interpretation of causality is if teachers grade differentially according to the temperatures *they* experience while grading, and temperature

grades that fall within the manipulable zone as established by Dee et al. (2016) based on the institutional features of NY Regents exams and described in greater detail below, I find that the point estimate for the impact of temperature is substantially reduced and no longer significant: β_T equals -0.0007 (se=0.0024) as opposed to -0.0082 (se=0.0021) in the full sample.

To assess the presence and magnitude of “adaptive grading”, I first estimate a version of Dee et al’s bunching estimator by school, subject, month, and year (in effect, the level of exam-time temperature variation). Starting with the student-exam level data, I calculate the fraction of observations in each 1 point score bin from 0 to 100 by core Regents subject. I then fit a polynomial to these fractions by subject, excluding data near the proficiency cutoffs with a set of indicator variables, using the following regression:

$$F_{ks} = \sum_{i=0}^q \psi_{ismyj} (\text{Score})^i + \sum_{i \in -M_{cs}, +M_{cs}} \lambda_{ismyj} \mathbb{1}[\text{Score} = i] + \epsilon_{ksmyj} \quad (6)$$

where F_{ks} denotes the fraction of observations with score k for subject s (e.g. ELA), q is the order of the polynomial, and $-M_{cs}, +M_{cs}$ represent manipulable ranges below and above the passing thresholds. The subscripts m, y and j denote, month, year, and school respectively.

Following Dee et al, (2016), I define a score as manipulable to the left of each cutoff if it is between 50 - 54 and 60 - 64, and manipulable to the right if it is between 55 - 57 and 65 - 67 as a conservative approximation of their subject-and-year-specific scale score-based rubric. In practice, I use a fourth-order polynomial (q=4) interacted with exam subject s , but constant across years for the same exam subject. Realized bunching estimates are not sensitive to changes in the polynomial order or whether one allows the polynomial to vary by year or subject.³⁰

This generates a set of predicted fractions by score and subject. I verify that the average amount of bunching observed in my data is similar to that documented by Dee et al. (2016), who find that approximately 6% of Regents exams between 2003 and 2011 exhibited grade manipulation. I then calculate observed fractions for each score from 0 to 100 by school, month, year, and subject, and generate a measure of

during the exam is correlated with temperature during grading. If hot temperatures make teachers less productive and causes more errors, this will simply add noise to the score variable. If hot temperature makes teachers irritable and more punitive in grading, then we might expect the beta coefficient to be picking up some of the correlation between test day temp and grading punitiveness, although the most striking feature of the histograms above is that the majority of grade manipulation seems to be positive in direction, making this unlikely in practice.

³⁰As a robustness check, I also estimate a linear approximation of the above estimator by generating predicted fractions using a linear spline between boundary points along the distribution that are known to be outside the manipulable range by subject. I then generate an estimate of the extent of bunching by school-subject-month-year cell, taking the absolute value of the distance between observed and predicted fractions by Regents scale score. The results are similar using this simplified measure of bunching.

bunching that integrates the differences between observed and predicted fractions: that is, summing the excess mass of test results that are located to the right of the cutoff (above the predicted curve) and the gaps between predicted and observed fractions of test results to the left of the cutoff (below the predicted curve). The bunching estimator can be written as:

$$\zeta_{smj} = \frac{1}{2} \sum_{i \in +M_{ck}} (F_{ks} - \hat{F}_{ksmj}) + \frac{1}{2} |\sum_{i \in -M_{ck}} (F_{ks} - \hat{F}_{ksmj})| \quad (7)$$

where ζ_{smj} denotes the degree of bunching at the passing cutoff for subject s , month m , year y , and school j . I then examine the relationship between ζ_{smj} and exam-time temperature in that cell, which corresponds to the temperature experienced by students taking subject s in school j in June of year y , with controls for precipitation and humidity.

To assess the magnitude of this relationship controlling for school-, subject-, and/or year-level differences in the degree of manipulation that are unrelated to temperature, I run a series of regressions with ζ_{smj} as the dependent variable:

$$\zeta_{smj} = \delta_0 + \delta_1 T_{smj} + \delta_2 X_{smj} + \chi_j + \eta_s + f(Year_y) + \epsilon_{smj} \quad (8)$$

where T_{smj} denotes temperature, X_{smj} denotes precipitation and humidity, χ_j , η_s , and θ_y denote school, subject, and year fixed effects respectively, and $f(Year_y)$ denotes a cubic time trend in scores. The parameter of interest is δ_1 , which represents the increase in grade manipulation due to exam-time temperature.

The amount of bunching increases by approximately 0.10 to 0.16 percentage points per degree F, or 1.7% to 2.8% per degree F hotter exam-time temperature. Coefficients are positive and significant in specifications with and without school and year fixed effects. These estimates are presented in tabular form in the appendix.

While these results are highly suggestive of adaptive grading, it is not possible to infer teachers' intentions based on these results alone. It could be the case that teachers have an intuitive sense of whether a particular student scored below his or her "true ability" – regardless of whether or not this was due to temperature or other exam-time conditions – and that they respond by manipulating grades in the case of students on the passing margin.

It is also possible that, due to the distributional properties of most Regents exams, heat-related performance impacts may lead to a mechanical increase in the number of grades that fall within the manipulable zone. This could in principle lead to a correlation between bunching behavior and exam-time temperature. I find however that, even when controlling for the potential increase in manipulable scores on hot exam days, there is evidence for more grade manipulation after hot exam takes.

Figure 8b presents a binned scatterplot of the bunching estimator and exam-time temperature by subject-month-year-school cell, adding school fixed effects to

allow for arbitrary differences in the average amount of grade manipulation across schools, and expressing the bunching estimate as a proportion of scores within the manipulable zone (50-54, 60-64). It suggests a clear positive relationship between the degree of grade manipulation and the ambient temperature during the exam being graded. It is consistent with teachers boosting students' grades more often when students experienced hot testing environments, and may suggest substantial institutional constraints to more efficient adaptation responses such as air conditioning or changes in exam format or timing.

Irrespective of whether teachers' explicit intentions are to compensate for heat-related impacts, the realized effect has been for this behavior to mitigate the adverse welfare impacts associated with exam-time heat exposure.

4 Discussion and Conclusion

This paper explores the impact of hot temperature on the human capital production process. Using administrative data from the largest public school district in the United States, I find that hot temperatures exert a causal and economically meaningful impact on student outcomes by (1) reducing performance on high-stakes exams, (2) possibly reducing the amount of learning achieved over the course of the school year, and (3) ultimately reducing high school graduation rates. The research design exploits quasi-random, within-student temperature variation to identify the causal impact of hot days on performance. The breadth and depth of the data set allows for credible causal estimation heat-related educational impacts as well as an assessment of avoidance behaviors and potential mechanisms through which students and teachers may adapt.

Taking a high stakes exam on a 90°F day results in 15% of a standard deviation reduction in exam performance relative to a more optimal 70°F day, controlling for student ability. These short-run performance impacts can lead to substantial reductions in longer-run educational attainment. A 90°F day results in a 10.9% lower probability of passing a subject, and, for the average New York City student, a 2.5% lower likelihood of graduating on time. I estimate that, over the period 1998 to 2011, upwards of 510,000 exams that otherwise would have passed received failing grades due to hot temperature, affecting at least 90,000 students, possibly many more.

The findings are consistent with emerging evidence suggesting that repeated heat exposure can disrupt learning and reduce the rate of human capital accumulation (Cho, 2017; Garg et al., 2017). Cumulative heat exposure over the course of the preceding school year, measured by the number of days where temperatures exceed 80°F, is associated with reductions in end of year exam performance, controlling for the exam-day effects of heat stress noted above. A year with five additional 80°F+ days is associated with 2.1% of a standard deviation reduction in learning on average, effects that are similar in magnitude with previous findings. While measured with

substantial noise, this effect is comparable in magnitude to eliminating half of the gains from having a one standard deviation higher value-added teacher for a school year, though more careful research is needed to examine whether they result in similar effects on later-life outcomes, given well-documented fade-out in teacher-driven score effects (Cascio and Staiger, 2012) and the possibility that better teachers impart important skills not captured by subject exams (Cunha and Heckman, 2007).

A series of additional analyses provide a window into the important interactions between institutional context and adaptation to climate change. The evidence suggests that, as has been shown in the context of human health (Kahn, 2005), the quality of institutions may be an important factor that determines the realized welfare impacts of any given environmental shock, including hotter temperatures from climate change.

Building level air-conditioning data suggests that school air conditioning may be sub-optimally provided. At less than 62% (as of 2012), the AC penetration rate in New York City public schools seems to be far below residential and commercial sector averages for the region (over 90%). Existing AC does not appear to have a protective effect, though it is possible that this may be due to lack of timely utilization or other factors specific to NYC schools (e.g. old building stock).

Perhaps in response to the lack of first-best adaptation investment, teachers seem to have selectively boosted grades of students who experience hot exam sittings. An analysis using bunching estimators developed in previous work (Dee et al., 2016), suggest a pattern of grade manipulation that is highly correlated with temperature during the test session in which a given subject exam was administered. One interpretation is that teachers may have been trying to offset some of the long-term consequences of short-term heat stress, which presumably affects students' scores but does not reduce human capital *per se*.³¹ The findings underscore the need for more careful research on the impact of school infrastructure on student performance, as well as potential market failures that may drive a wedge between realized adaptation strategies and the efficient adaptation frontier.

These results have several implications. First, they suggest that temperature should be included among the long list of relevant inputs to schooling. For instance, the timing of high-stakes exams – and the characteristics of the built environment in which they are administered – may affect social welfare: either from the standpoint of efficiency or distributional equity. Hot exam days may add noise to the signal-extraction process of high-stakes testing, thus leading to allocative inefficiencies in labor and higher education markets, as documented by Ebenstein et al. (2016). They may also affect distributional equity if individuals taking nationally or internationally standardized exams such as the SAT or ACT in a hotter region or in an

³¹A possible unintended consequence of eliminating teacher discretion in New York City public schools in 2011 may have been to expose more low-performing students to climate-related human capital impacts, eliminating a protection that applied predominantly to low-achieving Black and Hispanic students.

older school building are placed at a disadvantage relative to their peers in cooler regions or climate-controlled buildings. The latter dimension may be of particular importance in developing countries considering the well documented relationship between hotter climates and lower per capita incomes (Acemoglu and Dell, 2010) as well as the strong links between income and air conditioning ownership at the household level (Gertler et al., 2016). It may also be relevant in thinking about the persistence of racial achievement gaps in the United States, given the correlations between race, income, and average climate.

Second, this study raises new questions in an old debate regarding geography and economic prosperity (Acemoglu et al., 2001; Rodrik et al., 2004). How much of the variation in student achievement across countries (Figure 1a) can be explained by the cumulative influence of temperature stress on learning? Important caveats to cross-country comparisons notwithstanding, is it possible that hotter, poorer countries are subject to more challenging baseline learning conditions due to a combination of hot climate, lack of protective capital, and inflexible institutions?

Finally, from the perspective of climate policy, this study suggests that current social cost of carbon estimates omit important elements of the climate damage function: especially those mechanisms, including human capital accumulation, that operate through direct heat-stress of the human body (Burke and Emerick, 2016; Tol, 2009; Heal and Park, 2016). These findings also support the notion that climate change may affect not only the level of economic activity but overall growth rates (Pindyck, 2013), though more research on the impact of cumulative heat exposure on learning is needed. To the extent that future climate change may result in a disproportionate increase in realized heat exposure for the poor within and across countries, these findings lend further support to the notion that climate change may have unequal impacts across the income distribution. It is worth noting that while the average New Yorker is exposed to approximately 14 days above 90°F per year, the average resident of New Delhi experiences over 80 such days annually, with climate forecasts suggesting up to 190 such days per year in New Delhi by 2100.

At the same time, this study also underscores the importance of taking adaptive responses into account when thinking about the realized welfare consequences of climate change, especially when using short-run, weather-based estimates to inform projections about the distant future.

References

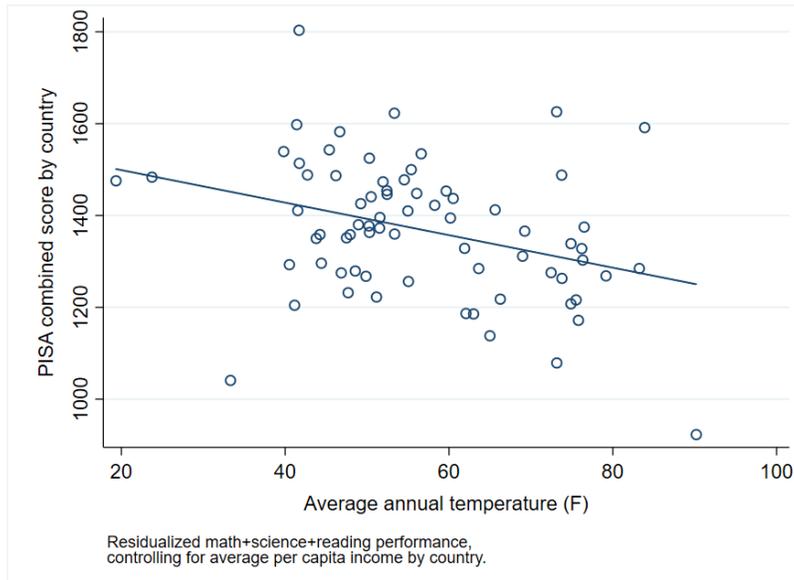
- Acemoglu, Daron and Melissa Dell (2010), “Productivity differences between and within countries.” *American Economic Journal: Macroeconomics*, 2, 169–188.
- Acemoglu, Daron, Simon Johnson, and James A Robinson (2001), “The colonial origins of comparative development: An empirical investigation.” *The American Economic Review*, 91, 1369–1369.

- Allcott, Hunt and Michael Greenstone (2012), "Is there an energy efficiency gap?" *The Journal of Economic Perspectives*, 26, 3–28.
- Angrist, Joshua D and Victor Lavy (1999), "Using maimonides' rule to estimate the effect of class size on scholastic achievement." *The Quarterly Journal of Economics*, 114, 533–575.
- Annan, Francis and Wolfram Schlenker (2015), "Federal crop insurance and the disincentive to adapt to extreme heat." *The American Economic Review*, 105, 262–266.
- 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.
- Burke, Marshall and Kyle Emerick (2016), "Adaptation to climate change: Evidence from us agriculture." *American Economic Journal: Economic Policy*, 8, 106–40.
- Burke, Marshall, Solomon M Hsiang, and Edward Miguel (2015), "Global non-linear effect of temperature on economic production." *Nature*, 527, 235–239.
- Cascio, Elizabeth U and Douglas O Staiger (2012), "Knowledge, tests, and fadeout in educational interventions." *NBER Working Paper*.
- Chetty, Raj, John N Friedman, and Jonah E Rockoff (2014), "Measuring the impacts of teachers ii: Teacher value-added and student outcomes in adulthood." *The American Economic Review*, 104, 2633–2679.
- 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.
- Dee, Thomas S, Will Dobbie, Brian A Jacob, and Jonah Rockoff (2016), "The causes and consequences of test score manipulation: Evidence from the new york regents examinations." *NBER Working Paper*.
- 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.
- Deryugina, Tatyana and Solomon M Hsiang (2014), "Does the environment still matter? daily temperature and income in the united states." *NBER Working Paper*.
- 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.

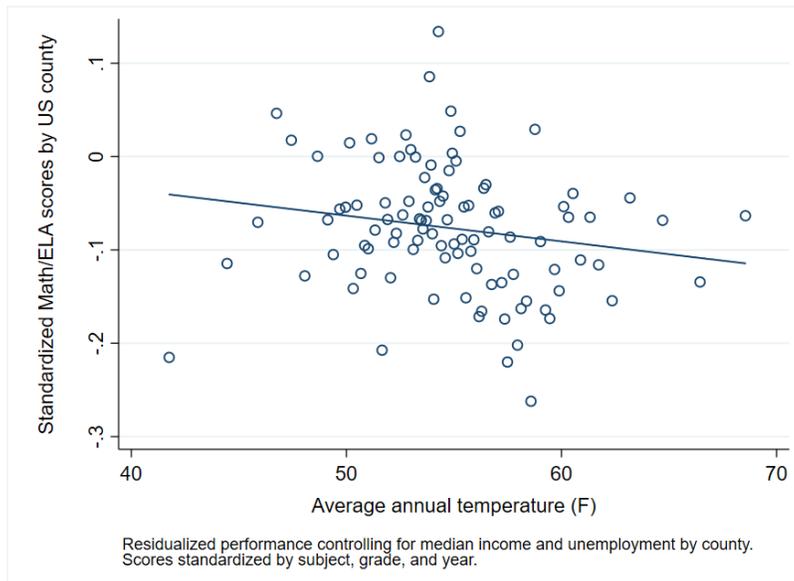
- Diamond, Rebecca and Petra Persson (2016), “The long-term consequences of teacher discretion in grading of high-stakes tests.” *NBER Working Paper*.
- Ebenstein, Avraham, Victor Lavy, and Sefi Roth (2016), “The long-run economic consequences of high-stakes examinations: evidence from transitory variation in pollution.” *American Economic Journal: Applied Economics*, 8, 36–65.
- Garg, Teevrat, Maulik Jagnani, and Vis Taraz (2017), “Human capital costs of climate change: Evidence from test scores in india.”
- Gertler, Paul J, Ori 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 (2014), “Temperature and the allocation of time: Implications for climate change.” *Journal of Labor Economics*, 32, 1–26.
- Grether, WF (1973), “Human performance at elevated environmental temperatures.” *Aerospace Medicine*, 44, 747–755.
- Heal, Geoffrey and Jisung Park (2013), “Feeling the heat: Temperature, physiology & the wealth of nations.” *NBER Working Paper*.
- Heal, Geoffrey and Jisung Park (2016), “Reflectiontemperature stress and the direct impact of climate change: A review of an emerging literature.” *Review of Environmental Economics and Policy*, 10, 347–362.
- Hsiang, Solomon M, Marshall Burke, and Edward Miguel (2013), “Quantifying the influence of climate on human conflict.” *Science*, 341, 1235367.
- Kahn, Matthew E (2005), “The death toll from natural disasters: the role of income, geography, and institutions.” *The Review of Economics and Statistics*, 87, 271–284.
- 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.
- Mackworth, Norman H (1946), “Effects of heat on wireless operators.” *British journal of industrial medicine*, 3, 143.
- Park, Jisung (2016), “Will we adapt? temperature shocks, labor productivity, and adaptation to climate change in the united states (1986-2012).” *Harvard Project on Climate Agreements*, 16–81.

- Peet, Evan D (2014), “Environment and human capital: The effects of early-life exposure to pollutants in the philippines.”
- Pindyck, Robert S (2013), “Climate change policy: What do the models tell us?” *Journal of Economic Literature*, 51, 860–872.
- Reardon, Sean F., Demetra Kalogrides, Andrew Ho, Ben Shear, Kenneth Shores, and Fahle Erin. (2016), “Stanford education data archive.”
- Roback, Jennifer (1982), “Wages, rents, and the quality of life.” *Journal of political Economy*, 90, 1257–1278.
- Rodrik, Dani, Arvind Subramanian, and Francesco Trebbi (2004), “Institutions rule: the primacy of institutions over geography and integration in economic development.” *Journal of economic growth*, 9, 131–165.
- Roth, Sefi (2016), “The contemporaneous effect of indoor air pollution on cognitive performance: Evidence from the uk.” *Unpublished Manuscript*.
- Seppanen, Olli, William J Fisk, and QH Lei (2006), “Effect of temperature on task performance in office environment.” *Lawrence Berkeley National Laboratory*.
- Sievertsen, Hans Henrik, Francesca Gino, and Marco Piovesan (2016), “Cognitive fatigue influences students performance on standardized tests.” *Proceedings of the National Academy of Sciences*, 113, 2621–2624.
- Sinha, Paramita, Maureen L Cropper, and Martha Caulkins (2015), “Household location decisions and the value of climate amenities.” *NBER Working Paper*.
- Sudarshan, Anant and Meenu Tewari (2013), “The economic impacts of temperature on industrial productivity: Evidence from indian manufacturing.”
- Tol, Richard SJ (2009), “The economic effects of climate change.” *The Journal of Economic Perspectives*, 23, 29–51.

Figures and Tables

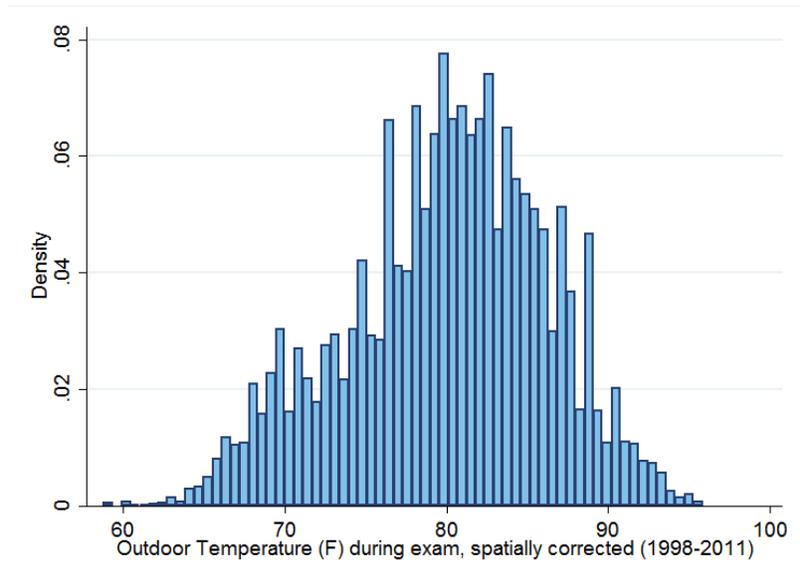


(a) Climate and Standardized PISA Math, Reading, and Science Achievement Across Countries

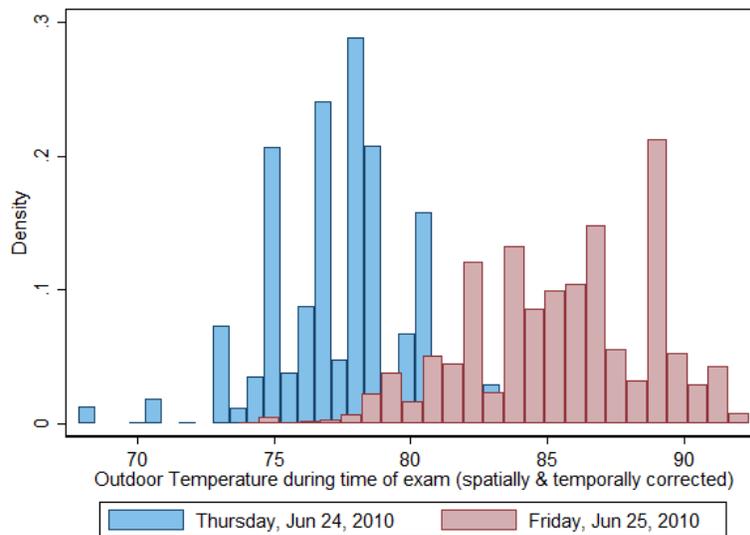


(b) Climate and Standardized Math and Reading scores within the United States

Notes: Panel (a): scatterplot of mean PISA scores (math+reading+science) and average annual temperature by country, plotting residual variation after controlling for mean per capita income in 2012. A standard deviation in PISA scores corresponds to approximately 300 points (100 per subject). Panel (b): binned percentile plot of standardized 3-8th grade math and reading scores (2009-2013) by percentile of the county-level temperature distribution, plotting residualized variation after controlling for median income and average unemployment (scores standardized by subject-grade-year as in Reardon et al. (2016)). Average annual temperatures in both cases are measured over the period 1980-2011.



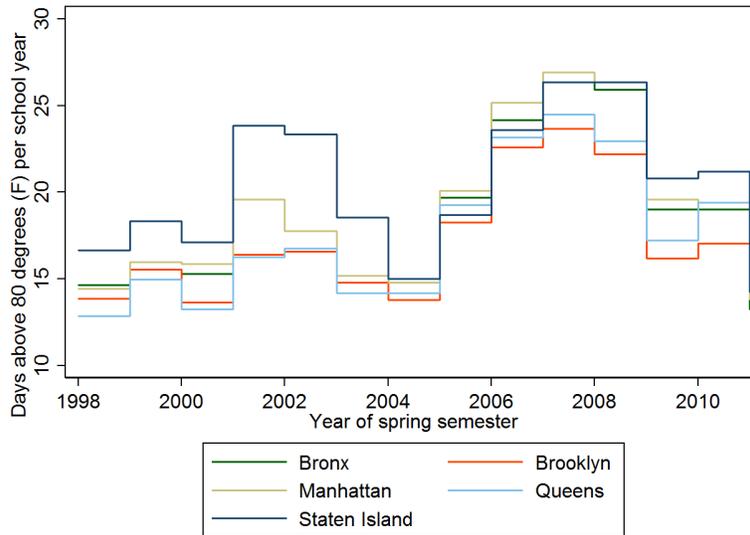
(a) Total variation in exam-time temperature



(b) Temperature on two consecutive test days

Figure 2: Short-Run Identifying Variation in Temperature

Notes: This figure illustrates the source of identifying variation for short-run performance impacts of heat exposure. It presents realized exam-time temperatures for (a) all June Regents exams (1999-2011) and (b) for two subsequent days within a Regents exam period – Thursday, June 24th, 2010, and Friday, June 25th, 2010 – inclusive of spatial and temporal temperature corrections. Temperatures are measured at the school level, weighted by number of exam observations by date and time.

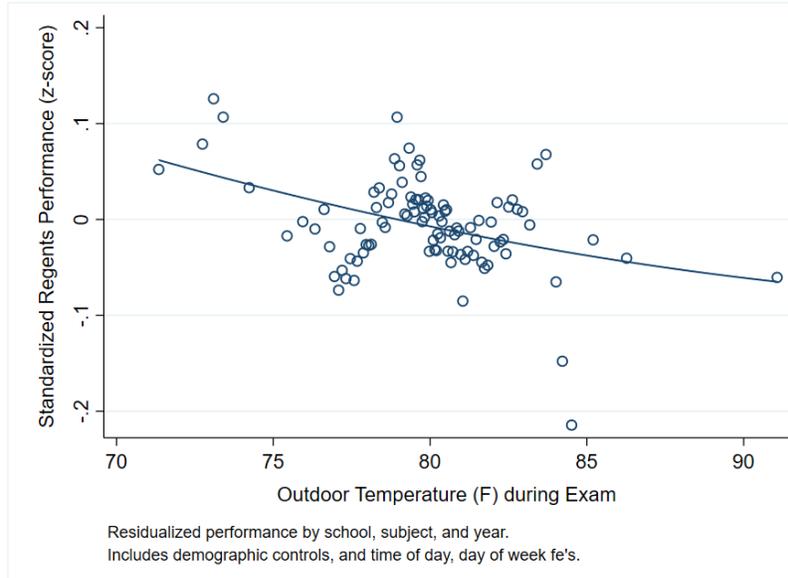


(a) Cumulative Heat Exposure by School Year

Ethnicity	Score	Pass	Proficiency	Previous Ability
Asian	74.73 (16.80)	0.78 (0.41)	0.57 (0.49)	0.98 (1.54)
Black	61.21 (17.05)	0.50 (0.50)	0.23 (0.42)	-0.18 (1.34)
Hispanic	61.49 (17.23)	0.51 (0.50)	0.24 (0.42)	-0.16 (1.32)
Multiracial	69.65 (17.44)	0.69 (0.46)	0.44 (0.50)	0.34 (1.26)
Native American	61.96 (18.08)	0.51 (0.50)	0.26 (0.44)	-0.22 (1.45)
White	72.92 (16.78)	0.75 (0.43)	0.52 (0.50)	1.02 (1.56)
Total	64.86 (17.92)	0.57 (0.49)	0.32 (0.47)	0.16 (1.42)

(b) Summary Statistics by Ethnicity

Notes: Figure 3a illustrates year to year variation in cumulative heat exposure during the school year, measured in terms of the number of days with max temperatures above 80°F per school year. Temperature readings are taken from USGS weather stations, one from each of the five boroughs of NYC. Table 3b: Standard deviations are in parentheses. ‘Pass’ and ‘Proficiency’ denote the fraction of scores above passing and college proficiency thresholds. Previous ability is measured as average z-scores from standardized math and verbal assessments in grades 3 through 8.



(a) Residualized variation in test performance

	(1)	(2)	(3)	(4)
	Z-score	Z-score	Z-score	Z-score
Temperature (F)	-0.00850*** (0.00231)	-0.00736*** (0.00207)	-0.0102*** (0.00233)	-0.0108*** (0.00226)
Afternoon	-0.0297* (0.0130)	-0.0334** (0.0119)	-0.0180 (0.0142)	-0.0156 (0.0127)
Fixed Effects				
Student by Year	X			
Subject	X	X	X	X
Time of Day, Day of week	X	X	X	X
Student		X		
Year		X	X	
School			X	
School by Year				X
N	3581933	3581933	3581933	3581933
r2	0.774	0.717	0.252	0.271

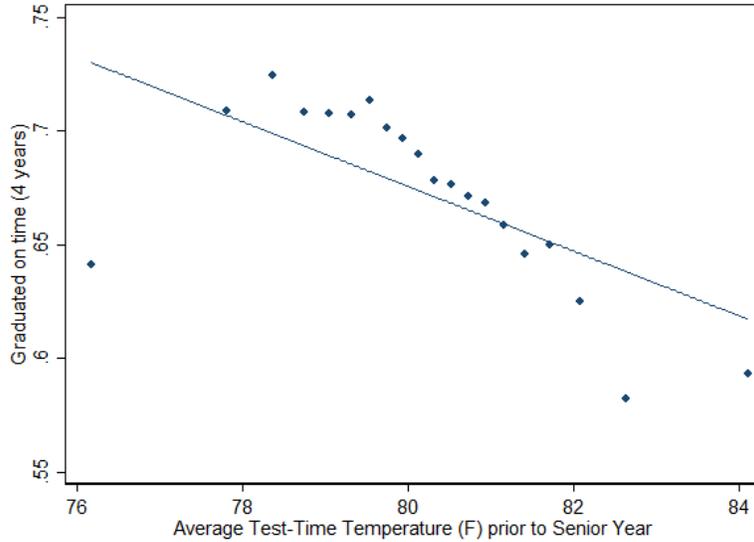
Robust standard errors in parentheses, clustered at the borough (station) by date-time level.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

(b) Dependent variable is standardized performance by subject

Figure 4: Short-Run Impacts of Heat Exposure on Exam Performance

Notes: Panel (a): binned scatterplot of residualized exam performance by percentile of the temperature distribution controlling for school, subject, and year fixed effects. Each dot represents approximately 220,000 exam observations. Panel (b) presents the main regression results. Fixed effects are suppressed in output, and 919,067 singleton observations are dropped. All regressions include controls for daily dewpoint, precip, ozone, and pm2.5.



(a) 4-year graduation status and exam-time temperature

	(1)	(2)	(3)
	Graduated	Graduated	Graduated
Avg Exam-Time Temp ($^{\circ}$ F)	-0.00712*** (0.00173)	-0.00758*** (0.00223)	-0.00733** (0.00231)
Number of June exams			0.193*** (0.00688)
Number of June exams ²			-0.0151*** (0.000809)
Number of June exams ³			0.000312*** (0.0000225)
Fixed Effects			
School	X	X	X
Number of June exams	X	X	
Cohort		X	X
N	515192	515192	515192
r ²	0.232	0.238	0.236

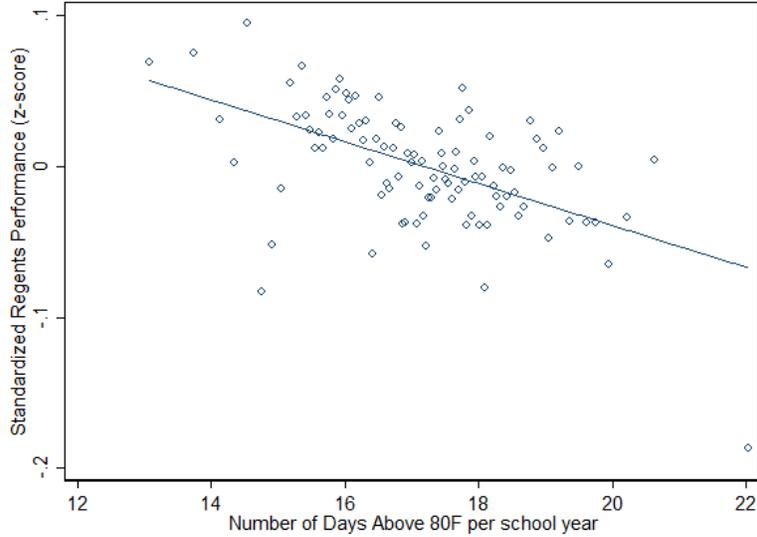
Robust standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

(b) Impacts on graduation status by regression specification

Figure 5: Long-Run Consequences of Short-Run Heat Exposure: Graduation Status

Notes: Panel (a): binned scatterplot of 4-year graduation status by quantile of exam-time temperature distribution. Temperatures are averaged by student for June exam sessions up through senior year. Residual variation after controlling for school and number of exam fixed effects, student-level observables, and weather/air quality controls. In (b), the dependent variable is a dummy for whether or not student graduated in four years. All regressions include controls for daily precipitation, ozone, and dewpoint. Fixed effects are suppressed in output.



(a) Cumulative impacts

	(1)	(2)	(3)	(4)
	Z-score	Z-score	Z-score	Z-score
70°F-80°F days	0.000669 (0.00171)	0.00506 (0.00322)	0.00131 (0.00187)	0.00412 (0.00348)
80°F-90°F days	-0.0108*** (0.00310)	-0.0121*** (0.00329)	-0.0114*** (0.00279)	-0.00337 (0.00668)
90°F+ days	0.0209* (0.00937)	0.00459 (0.0129)	0.00804 (0.00973)	0.0198 (0.0182)
Exam-Time Temp (°F)	-0.00580*** (0.00152)	-0.00830*** (0.00172)	-0.00790*** (0.00165)	-0.00824*** (0.00179)
Cold Day Bins		X	X	
Avg School Year Temp			X	
All Degree Day Bins				X
N	22563	22563	22563	22563
r2	0.587	0.587	0.587	0.587

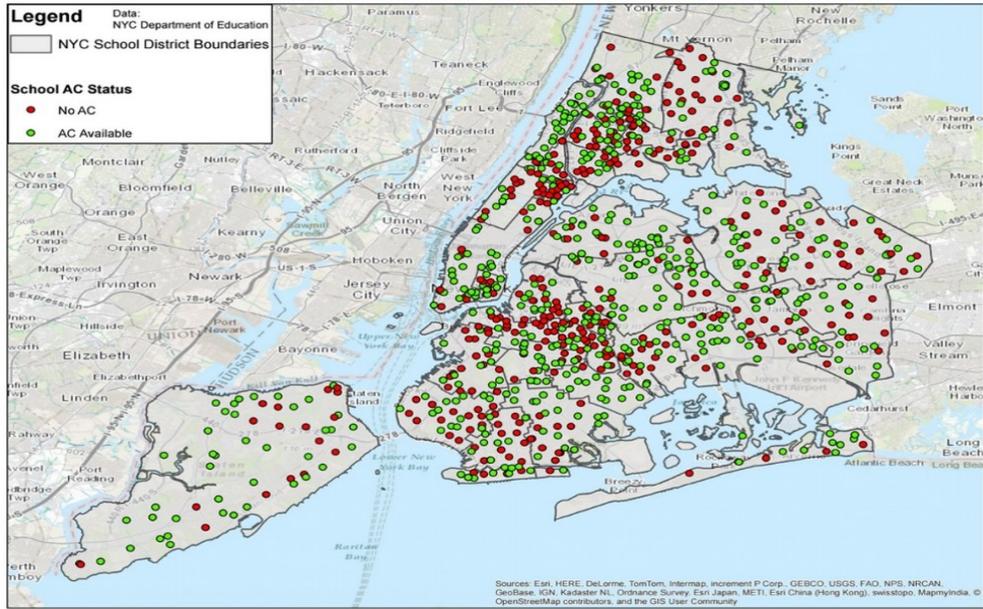
Robust standard errors in parentheses, clustered at the station-by-year level

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

(b) Dependent variable is standardized Regents performance

Figure 6: Cumulative Learning Impacts of Heat Exposure

Notes: Panel a: binned scatterplot of residualized exam performance at the school level on amount of cumulative heat exposure, measured as the number of days per degree bin (days above 80°F on x axis), including controls for exam-day precipitation and school, subject, time of day, day of week, and day of month fixed effects. All regressions in (b) include school, subject, and time of day fixed effects which are suppressed in the output. Cumulative degree day variables are assigned by closest weather station and summed beginning on the first day of the preceding fall semester up through the first day of June Regents exams that year.



(a) AC status by school

	(1)	(2)	(3)	(4)	(5)
	All	Central	Any AC	No Central	No AC
Temp (°F)	-0.00613** (0.00219)	-0.00530 (0.00293)	-0.00517 (0.00273)	-0.00649* (0.00274)	-0.00621* (0.00261)
Precip (mm)	0.00183 (0.00128)	0.00212 (0.00154)	0.00177 (0.00141)	0.00219 (0.00156)	0.00197 (0.00147)
Afternoon	-0.0417** (0.0149)	-0.0439* (0.0181)	-0.0473** (0.0169)	-0.0338* (0.0171)	-0.0369* (0.0163)
N	4347718	906451	1019279	1611336	1724670
r ²	0.710	0.720	0.717	0.710	0.709

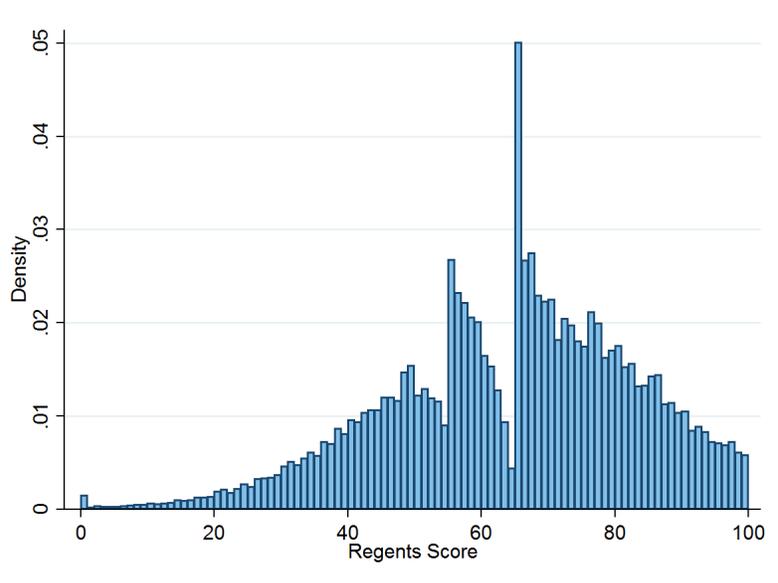
Robust standard errors in parentheses, clustered at the station-by-date-time level

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

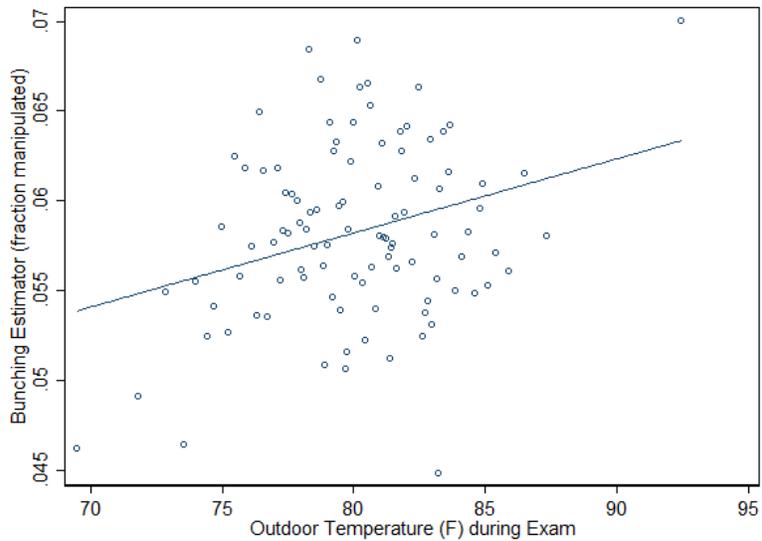
(b) Short-Run Impacts by School Air Conditioning Status

Figure 7: Adaptive Responses: Air Conditioning

Notes: Panel (a) provides a map of New York City public schools, with green dots representing schools that had any air conditioning equipment as of 2012, and red dots representing schools that did not. The dependent variable in (b) is standardized Regents performance, with observations at student, exam, and date-and-time-level. All regressions include student, subject, and year fixed effects, as well as controls for dewpoint, ozone, pm2.5 and precipitation. Fixed effects are suppressed in output, and singleton observations are dropped.



(a) All Regents exams in core subjects prior to NYC grading reforms in 2011-2012.



(b) Grade Manipulation varies with exam-time temperature by subject, school, and take.

Figure 8: Adaptive Responses: Grade Manipulation

Notes: Panel (a) presents a histogram of Regents exam scores from June 1999 to June 2011. A large number of observations bunch at the pass/fail cutoffs, scores of 55 and 65 for local and Regents diploma requirements respectively. Panel (b) presents a binned scatterplot of bunching at the school-subject-date level by quantile of the exam-time temperature distribution, controlling for subject and year fixed effects and daily weather and air quality controls. Included in the analysis are all June Regents exams in core subjects between 1999 and 2011.