The Dynamic Effects of Temperature on Human Capital

Kyle Emerick, Maggie Y. Liu, Yogita Shamdasani & Vis Taraz*

July 21, 2025

Abstract

Little is known about the long-term, dynamic effects of temperature on human capital. We analyze impacts of life cycle heat exposure on school completion using panel data spanning over forty birth cohorts in India, combined with a dynamic framework that incorporates lagged weather shocks and parental investments. Contemporaneous heat exposure systematically reduces grade completion whereas early-life heat exposure is sometimes associated with increased grade completion, likely due to compensatory parental investments. Parental investments also increase with contemporaneous exposure, but these efforts do not fully mitigate heat damages. Our findings underscore the importance of examining temperature impacts over long time frames.

JEL Codes: Q54, O13, J24, I25

*Emerick: Tufts University, Department of Economics, 177 College Avenue, Medford, MA 02155, kyle.emerick@tufts.edu. Liu: U.S. Department of the Treasury, Office of Tax Analysis, 1500 Pennsylvania Avenue NW, Washington, D.C. 20220, yuanyuan.liu@treasury.gov. Shamdasani: National University of Singapore, Department of Economics, 1 Arts Link, Singapore 117570, yogita@nus.edu.sg. Taraz: Smith College, Department of Economics, Wright Hall, 5 Chapin Drive, Northampton, MA 01063-6317, vtaraz@smith.edu.
Archana Dhinakar Bala and Saumya Singla provided excellent research assistance. We are grateful for comments provided by numerous seminar audiences and conference participants. Liu is an employee at the U.S. Department of the Treasury. The findings, interpretations, and conclusions expressed in this paper are entirely those of the author and do not necessarily reflect the views or the official positions of the U.S. Department of the Treasury.

1 Introduction

As climate change continues to accelerate, more frequent extreme weather events disproportionately affect low- and middle-income countries (LMICs) (Diffenbaugh and Burke, 2019). In these settings, adverse weather impacts a host of outcomes, from agricultural productivity (Taraz, 2018; Aragón et al., 2021; Liu et al., 2023) and conflict (Hsiang et al., 2013) to student learning (Park et al., 2021) and academic performance (Garg et al., 2020). Empirical research in this domain has typically linked short-term adverse weather shocks to contemporaneous changes in outcomes, after which these effects are used to infer longer-term implications of climate change. This extrapolation has its limitations, however, given that climate change unfolds over decades and climate-dependent outcomes often form dynamically over time.

Human capital is an example of one such outcome. Its development is a dynamic process in which the returns to educational inputs vary across stages of the life cycle (Todd and Wolpin, 2003; Cunha et al., 2010). Further, the impacts of adverse weather on human capital may also differ by the age of the child (Shah and Steinberg, 2017; Stowell et al., 2022). It is also possible that parental investment in education may respond to lagged weather shocks — for example, parents may spend on after-school tutoring or move their children to schools with better resources. It is thus unclear whether short-term adverse weather shocks translate into persistent declines in human capital or whether parental responses can mitigate these damages. Understanding the effects of climate change on human capital therefore requires analysis that considers weather exposure and educational investment in a dynamic, long-term setting.

In this paper, we analyze how sustained heat exposure influences human capital formation in a large developing economy. Given the pivotal role of human capital in shaping economic growth as well as social outcomes, understanding its vulnerability to climate change is of critical policy importance. We build on prior studies documenting impacts of year-to-year weather shocks on human capital and academic performance by exploring the dynamic impacts of temperature over the life cycle of schooling.[1]

We present a simple model that describes the dynamics of human capital formation, extending the standard human capital production function to incorporate temperature as an input. Temperature can influence human capital through both physiological and income related channels (Dell et al., 2012; Burke et al., 2015; Garg et al., 2020). In each period, households decide how much to invest in a child's education. Exposure to past warming may alter the marginal productivity of these investments: it could reduce returns if it impairs the child's ability to benefit from schooling, or increase returns if additional investment helps the child recover from setbacks. As a result, the cumulative impact of long-term climate change can differ markedly from impacts of short-term, one-off weather shocks. This theoretical ambiguity motivates our empirical analysis of the relationship between heat exposure over the life cycle, educational investment and human capital outcomes.

We build a district-level panel dataset on grade completion by birth cohort across primary, middle, and secondary schooling for individuals born between 1952 and 1996 in India. We link these educational outcomes to district-level temperature and precipitation measures derived from high-resolution reanalysis daily weather data. Our empirical strategy employs a standard panel fixed effects model, exploiting variation in heat exposure experienced by each cohort during their years of schooling. Specifically, we include district and region-by-birth cohort fixed effects to control for time-invariant district characteristics and region-specific trends. Our identification relies on the assumption — in line with the recent climate–economy literature (Dell et al., 2014) — that conditional on these controls, fluctuations in heat exposure are quasi-random and uncorrelated with other determinants of educational attainment.

We find that contemporaneous heat exposure during school years has negative impacts

[1] Adverse impacts have been documented for rainfall and temperature shocks during early life — from the period in utero and early childhood (Maccini and Yang, 2009; Fishman et al., 2019; Hu and Li, 2019; Adhvaryu et al., 2024) — as well as for contemporaneous shocks during school- or college-age years (Shah and Steinberg, 2017; Graff Zivin et al., 2018; Garg et al., 2020; Graff Zivin et al., 2020; Zimmermann, 2020; Park, 2022; Park et al., 2021) in developed and developing country settings.

on educational attainment at all three schooling levels. These impacts are economically meaningful: an additional hundred degree days above 24°C during a school year — which corresponds approximately to a 10% increase in heat exposure — reduces the probability that a cohort completes primary school by 2.3 percent. The corresponding effects for middle and secondary school completion are 3.3 and 3.8 percent, respectively. These results align with prior literature documenting negative effects of adverse weather on human capital, but fail to account for potential dynamic responses that could alter the long-term impacts of heat exposure.

We next examine dynamic effects by estimating a model that considers heat exposure over a child's life cycle, starting from the period in utero up to age 16. We find that the adverse impacts of heat exposure on grade completion are primarily attributable to contemporaneous exposure during school years. In contrast, heat exposure in early life (in utero to age 5) and in the years preceding a particular level of schooling do not have persistent adverse effects. Children who experience greater heat exposure during the in utero period and in their first five years of life are no less likely to finish primary school, and heat exposure during primary school years does not affect middle school completion. Notably, heat exposure in early life is associated with *positive* effects on middle and secondary school completion. Overall, the evidence suggests that while recent heat exposure adversely impacts human capital, the effects of more distant, lagged exposure are smaller, or even positive.

In addition to grade completion, we examine effects of heat exposure over the life cycle on a more intermediate outcome: school enrollment. Greater heat exposure is associated with a decline in school enrollment. These adverse impacts carry over to subsequent grades of schooling, thus highlighting the dynamic nature of human capital accumulation — adverse impacts in early stages of a child's life may compound in later stages, if unmitigated. At first glance, this finding appears to be at odds with our results on grade completion being unaffected by heat exposure in early life and in the years preceding a particular level of schooling. However, these patterns are reconciled once we consider parental responses to heat exposure.

Using nationally representative household survey data on educational expenditures for school–age individuals spanning four decades, we find that early life heat exposure leads to greater educational expenditures, and this pattern holds through to secondary school. Households also increase educational spending in response to contemporaneous heat exposure. Parents are thus responding actively to greater heat exposure by investing more in their children's education, particularly when their children are exposed to heat in early life and during school years. Taken together, our results suggest that educational investment is able to mitigate the adverse impacts of heat exposure, particularly when shocks occur early in a child's life, by allowing parents more time to make compensatory educational investments throughout the life cycle. However, contemporaneous heat exposure continues to reduce educational attainment despite households increasing educational spending, which implies that these efforts are unable to fully mitigate damages from heat.

This paper makes four contributions to the literature on the causal impacts of temperature on education-related outcomes. First, to shed light on the dynamic relationship between climate-driven heat exposure and human capital accumulation, we explicitly model temperature as an input in the production of human capital, allowing for both short-run and long-run effects of temperature on human capital as well as dynamic responses through parental investment. Shah and Steinberg (2017) use a similar approach to examine impacts of favorable rainfall shocks on human capital. Their approach differs from ours in that they assume a direct link between rainfall and wages, and therefore limit the effect of rainfall on human capital accumulation through exogenous changes in income. Our approach allows temperature to affect human capital both directly (e.g. through physiological effects) and indirectly (e.g. through parental income). Further, Shah and Steinberg (2017) examine favorable income shocks and focus on the opportunity cost of schooling by differentiating income and substitution effects. Our empirical setting centers around adverse heat exposure and focuses on intertemporal dynamics — in particular, we explore whether heat-driven setbacks in schooling in the short-run persist as permanent setbacks in the long-run, as well as investigate the role of parental investments in mitigating damages from heat.

Second, our empirical analysis directly links heat exposure, parental responses, and schooling outcomes to offer new insights into how human capital accumulation responds to sustained, climate-driven warming in a long-term, dynamic setting. Prior studies have typically focused on a single, transitory shock — for example, Fishman et al. (2019) look at the impact of higher temperatures in utero on higher education attainment — or have exploited day-to-day or year-to-year variation in temperature exposure (Graff Zivin et al., 2018; Garg et al., 2020; Graff Zivin et al., 2020; Park, 2022). In contrast, our analysis considers heat exposure over the entire life cycle. This approach enables us to capture cumulative effects and identify critical windows of vulnerability. Further, by explicitly examining impacts on school enrollment and parental responses, we shed light on potential compounding or mitigating factors through which long-term heat exposure influences human capital formation.

Third, we focus on a critical indicator of educational attainment: school completion. Much of the existing literature has assessed impacts of temperature on intermediate learning outcomes such as test scores (Cho, 2017; Graff Zivin et al., 2018; Garg et al., 2020; Graff Zivin et al., 2020; Park et al., 2020; Park, 2022). While test scores reflect cognitive achievement, they capture academic performance only at specific points in time, and negative effects may dissipate if students are able to catch up, possibly through physiological adaptation and/or greater compensatory investment by parents. However, if students fall too far behind, they may be forced to repeat a grade or drop out, resulting in a permanent educational setback with potentially far-reaching consequences. By focusing on grade completion as our primary outcome, we capture the cumulative and potentially irreversible impacts of climate-related disruptions. Given the strong links between educational attainment, labor market participation and lifetime earnings, school completion is a critical metric for evaluating the long-term policy implications of climate change on human capital development.

Fourth, we conduct our long-term analysis in the context of a large developing economy. While much of the literature focuses on developed economies, often due to better data availability, impacts in LMICs are likely to differ substantially. In these settings, large segments of the population rely on rain-fed agriculture, making them particularly vulnerable to climate risks. Under this scenario, the agricultural income channel, in addition to the physiological channel, can influence the relationship between heat exposure and human capital. Additionally, limited social protection and restricted access to adaptive mechanisms such as cooling infrastructure exacerbate these vulnerabilities. We address data limitations by leveraging administrative census data covering more than forty birth cohorts who have completed schooling in India, combined with high-resolution, high-frequency weather data. This novel dataset allows us to examine nonlinear effects of heat exposure over a child's life cycle on educational attainment, offering new insights into the long-term consequences of climate change in resource-constrained environments.

The rest of the paper is organized as follows. Section 2 outlines a model that describes the different channels through which climate change may impact human capital accumulation. Section 3 describes the data sources. Section 4 outlines the empirical strategy. Section 5 describes results and Section 6 concludes.

2 Model

This section presents a model for formalizing different channels through which temperature affects human capital. The model separates the direct effect of current weather from longer-term effects of permanent climate. Following the literature, we model human capital formation as a dynamic process where human capital in any period depends on past human capital, climate events during that period, and parental inputs (Todd and Wolpin, 2003; Cunha et al., 2010). Persistent climate change therefore affects human capital through two channels: a contemporaneous effect of weather and a permanent effect where the amount of human capital that a child enters each period with can depend on climate.

2.1 Setup

We denote A_t as a continuous measure of human capital at the end of period t. We consider three distinct periods of human capital accumulation: $t = 0$ is the period before a child enters school, $t = 1$ is the time of primary school and $t = 2$ is the time of secondary school. Temperature in a given period consists of both permanent climate and an annual disturbance. Following convention, we refer to the permanent part of a temperature realization as climate and the annual disturbance as weather. Specifically, we denote an annual temperature realization as $w_t = \bar{w} + \phi_t$, where $\bar{w}$ is climate and ϕ_t is weather. Households earn income in period t which is temperature dependent, i.e. $y_t = h_t(\bar{w} + \phi_t)$. The function h is indexed by t to account for adaptation where the effect of temperature on income could decline over time.

Parents choose their educational investment x_t at the start of each period after observing all stochastic variables from previous periods, but before observing their realizations in the current period. This amounts to an optimization problem where parents choose investment to maximize lifetime utility. As a result of this maximization, the optimized educational investment can be written as a function of lagged income and human capital at the beginning of the period: $x_t = f_t(y_{t-1}, A_{t-1})$.

2.2 The determinants of human capital

Human capital in period t is a function of human capital at the end of the previous period (A_{t-1}), investment in the current period (x_t), and temperature (w_t); we denote this as $A_t = g_t(A_{t-1}, x_t, w_t)$. Investment could be expenses on after-school tutoring, learning materials, or the time parents spend studying with their children.

In this setup, human capital at the end of *primary school* is written as:

$$A_1 = g_1\left(A_0, f_1(y_0, A_0), w_1\right) = g_1\left(A_0, f_1(h_0(\bar{w} + \phi_0), A_0), w_1\right) \tag{1}$$

Human capital at the end of primary school depends on the amount of human capital that a student enters with, parental investment during primary school, and current temperature. Parental investment depends on income and the amount of human capital. The sign of $\frac{\partial f_t}{\partial A_{t-1}}$ plays a key role in determining how climate affects human capital in the long-run. There are two possibilities. Parents could increase investment in an attempt to overcome deficits in human capital, i.e. $\frac{\partial f_t}{\partial A_{t-1}} < 0$. Parents would choose to do this if current investments are substitutes for the stock of past human capital. Alternatively, parents could invest less in children that have faced adverse weather shocks, which would be the case if current investments are complements to the stock of past human capital. This would imply $\frac{\partial f_t}{\partial A_{t-1}} > 0$. Human capital at the end of *secondary school* is therefore:

$$A_2 = g_2\left(A_1, f_2(y_1, A_1), w_2\right) = g_2\left(A_1, f_2(h_1(\bar{w} + \phi_1), A_1), w_2\right) \tag{2}$$

2.3 Short-run effects of weather on human capital

Many studies focus on the contemporaneous effect of weather on human capital. This effect at the end of secondary school is

$$\frac{\partial A_2}{\partial \phi_2} = \frac{\partial g_2}{\partial \phi_2} = \frac{\partial g_2}{\partial w_2} \tag{3}$$

Notably, this only captures the direct effect of current weather on human capital. For instance, higher temperatures can impair learning — children may not attend school or may encounter learning disruptions while at school (Park et al., 2020, 2021). More directly, extreme heat negatively affects concentration and cognitive function, subsequently lowering test scores (Cho, 2017; Graff Zivin et al., 2018; Garg et al., 2020; Graff Zivin et al., 2020). Higher temperatures may also affect students' time allocation and amount of time spent studying (Alberto et al., 2021). While access to air conditioning or other cooling technologies may offset the adverse impacts of higher temperatures on academic performance (Park et al., 2020), current access to such adaptation technologies in LMICs remains limited (Pavanello et al., 2021).

The literature generally finds that $\frac{\partial A_2}{\partial \phi_2} < 0$. Importantly, this short-run effect does not include the parental investment channel or the long-run effects of climate from the carryover of human capital across periods. Our key contribution is to integrate long-run data on educational attainment and parental investment in order to move beyond this short-run parameter.

2.4 Long-run effects of climate on human capital

Climate produces different long-run effects for two reasons. First, human capital accumulation is dynamic, and climate impacts can accumulate over time if they cause students to enter future grades with lower levels of human capital. Second, parents invest in their children's education over time. The short-term effects of weather could be mitigated (or exacerbated) by changes in parental investment in the long-run.

The long-run effect of climate on human capital at the end of secondary school is

$$\frac{\partial A_2}{\partial \bar{w}} = \frac{\partial g_2}{\partial A_1}\frac{\partial A_1}{\partial \bar{w}} + \frac{\partial g_2}{\partial x_2}\left(\frac{\partial x_2}{\partial y_1}\frac{\partial y_1}{\partial \bar{w}} + \frac{\partial x_2}{\partial A_1}\frac{\partial A_1}{\partial \bar{w}}\right) + \frac{\partial g_2}{\partial w_2} \tag{4}$$

Rearranging terms,

$$\frac{\partial A_2}{\partial \bar{w}} = \frac{\partial g_2}{\partial A_1}\frac{\partial A_1}{\partial \bar{w}} + \frac{\partial g_2}{\partial x_2}\frac{\partial x_2}{\partial y_1}\frac{\partial y_1}{\partial \bar{w}} + \frac{\partial g_2}{\partial x_2}\frac{\partial x_2}{\partial A_1}\frac{\partial A_1}{\partial \bar{w}} + \frac{\partial g_2}{\partial w_2} \tag{5}$$

Three terms drive the difference between long- and short-run effects:

1. $\frac{\partial g_2}{\partial A_1}\frac{\partial A_1}{\partial \bar{w}}$: This is the persistent effect of climate on human capital accumulation over time. Higher temperatures worsen human capital outcomes during earlier periods, and this carries over into later periods because of the dynamic accumulation of human capital. This term is negative i.e., it pushes the long-term effect to be more negative relative to the short-term effect.

2. $\frac{\partial g_2}{\partial x_2}\frac{\partial x_2}{\partial y_1}\frac{\partial y_1}{\partial \bar{w}}$: This is the income effect through parental investment. This term could have either sign. It will be negative if investment is a normal good and higher temperatures under climate change lower income. Adverse weather shocks have been shown to negatively impact agricultural incomes (Schlenker and Roberts, 2009; Taraz, 2018), inhibit labor reallocation to non-agriculture (Liu et al., 2023), and reduce labor productivity (Somanathan et al., 2021). Early life shocks generate larger long-run effects when income is more sensitive to those shocks (Maccini and Yang, 2009; Shah and Steinberg, 2017; Fishman et al., 2019; Hu and Li, 2019; Adhvaryu et al., 2024). Experimental evidence from unconditional cash transfers in the United States shows that recipients spend more time on activities that enrich children's learning (Magnuson et al., 2024), which translates to $\frac{\partial x_2}{\partial y_1} > 0$.

3. $\frac{\partial g_2}{\partial x_2}\frac{\partial x_2}{\partial A_1}\frac{\partial A_1}{\partial \bar{w}}$: This term captures whether current parental investment offsets or exacerbates past climate-related damages to learning. If climate change harms learning during primary school, then $\frac{\partial A_1}{\partial \bar{w}} < 0$. Complementarity between learning during primary school and parental investment during secondary school would mean that $\frac{\partial x_2}{\partial A_1}$ is positive. For example, Cunha et al. (2010) find complementarity between earlier cognitive skill development and later educational investment for low-income children in the United States. If the same were true in the Indian setting, then the long-term adverse effect of climate change would be greater than the short-term effect because it is less productive for parents to invest in children who have fallen behind during primary school. An alternative is that $\frac{\partial x_2}{\partial A_1}$ is negative — this would happen if parental investment serves as a tool to offset past learning losses.

Research has found that households do adjust their investment in response to one-off weather shocks, and these adjustments can either mitigate or intensify the impact of weather shocks on human capital. Wu et al. (2023) find that high rainfall *in the year of birth* increases long-term human capital measures for girls but not for boys in China. The authors find that this is driven by a compensatory mechanism: boys in China are typically breastfed longer than girls, but in high-rainfall years, mothers work fewer agricultural hours and breastfeed female children more than they do in lower rainfall years. More broadly, there is a large literature discussing whether parents reinforce differences in human capital among their children, by investing resources in the child with higher human capital (Akee et al., 2018; Becker and Tomes, 1976; Behrman et al., 1994; Frijters et al., 2009; Rosenzweig and Zhang, 2009) or whether parents instead compensate for differences, by investing more in the child with lower human capital (Almond and Mazumder, 2013; Behrman et al., 1982; Fan and Porter, 2020; Halla and Zweimüller, 2014; Leight, 2017).

2.5 Links with the empirical analysis

This model helps to clarify three mechanisms that we aim to capture in the empirical analysis: the dynamics of learning, the effects of longer-term income, and the role of parental investment in mitigating climate damages. Empirical research on weather and learning tends to focus on the contemporaneous effect of heat on student performance or the longer-term impacts of early life shocks. Neither of these parameters identifies how climate change affects human capital in the long run. The contemporaneous effect ($\frac{\partial g_t}{\partial \phi_t}$) only captures direct cognitive impacts. There is some evidence that individuals can physiologically adapt to higher temperatures in the presence of sustained heat (Sexton et al., 2022), and this heat acclimatization can occur within as little as a week. If physiological adaptation to sustained high temperatures is possible, then we might expect the impact of persistent high temperatures to be smaller than contemporaneous effects. The effect of past shocks ($\frac{\partial g_t}{\partial \phi_{t-k}}$) picks up some of the dynamic channels, but only identifies the effect of a single weather event, not sustained climate change. Our empirical framework instead seeks to capture all three channels by incorporating data on the distribution of weather throughout a child's life cycle.

3 Data

Our empirical analysis combines administrative Census data on educational attainment and school enrollment, nationally representative household surveys on educational investment and hourly reanalysis weather data.

3.1 Educational Attainment

Grade completion data comes from three waves of the Indian Census: 1991, 2001, and 2011. We use publicly available Census tables that report educational attainment by five-year birth cohorts (ages 5-9, 10-14, ..., 80+) at the district level. For each five-year birth cohort bin, the tables report total counts of individuals as well as counts of individuals by their highest level of education. We focus on educational outcomes associated with completing primary, middle and secondary school, and calculate cumulative shares as outcome variables. Specifically, our primary outcome is the proportion of each five-year birth cohort that has completed at least primary, middle, or secondary education.

We arrive at a district-level panel of birth cohorts by constructing our sample as follows. First, district boundaries in India vary across the three Census waves. To account for shifting boundaries, we follow Liu et al. (2023) and create district units that are consistent over our analysis time frame. This gives rise to 441 districts, illustrated in Appendix Figure A.1. Second, the partition of India in 1947 led to high levels of migration and introduced compositional changes in average schooling levels in the population (Bharadwaj et al., 2015). To avoid any discontinuities in cohort education levels set in motion by the partition, we restrict our sample to individuals born after 1950. Third, the same birth cohort is often observed across multiple Census waves as they age. For example, the cohort aged 5 to 9 in 1991 corresponds to the cohort aged 15 to 19 in 2001 and the cohort aged 25 to 29 in 2011. If a cohort is observed too early, its members may have yet to complete a given education level — this is an important consideration in the Indian context where grade repetition is common and the pace of grade progression varies across regions (Ahsan et al., 2018). Conversely, if a cohort is observed too late, compositional changes resulting from mortality or migration may distort "true" grade completion rates of individuals born in a particular locality. To minimize these issues, we use data from the earliest census round by which a five-year birth cohort is likely to have completed a certain level of schooling, allowing for a buffer period. Specifically, we use the earliest Census wave in which the youngest member of a five-year birth cohort is at least five years older than the typical completion age for each schooling level (i.e., age 15 for primary, age 19 for middle, and age 21 for secondary).

Our sample comprises five-year birth cohorts born between 1952 and 1996. Figure 1 presents grade completion rates for these cohorts. We see steady increases in grade completion as we move from the oldest to the youngest cohorts across all three schooling levels, consistent with an upward trend in educational attainment over time in India.

3.2 School Enrollment

School enrollment data comes from three waves of the Indian Census: 1991, 2001, and 2011. We use publicly available Census tables that report district-level counts of individuals attending educational institutions. This data is presented at a single-cohort frequency for individuals between the ages of 5 to 19, and at a five-year frequency thereafter. We focus on primary, middle, or secondary school age children (i.e. between the ages of 7 to 16) at the time of each Census wave, and calculate shares of children enrolled in school as outcome variables.

3.3 Educational Investment

Educational investment data comes from five waves of the Household Social Consumption: Education Module (Schedule 25.2) conducted as part of India's National Sample Survey

(NSS).[2] We rely on questions about education expenditure for individuals who are currently enrolled in primary school or higher. We focus on primary, middle, or secondary school age children (i.e. between the ages of 7 to 16) at the time of each survey wave. For each enrolled child, we compute total education expenditure as the sum spent on course fees, books, stationery and uniforms, transportation, private coaching, and other miscellaneous items. For consistency, we harmonize district of residence in the survey to the same set of district boundaries used in the Census data.

3.4 Weather

Hourly temperature and precipitation data come from the ERA5 Climate Reanalysis dataset (ECMWF, 2023). This dataset extends from 1940 to the present and provides climate variables at a spatial resolution of a 0.25-degree latitude-longitude grid.

We use temperature data that is available at an hourly level and extract four points per day (at six-hour intervals), after which we construct a daily average temperature measure for each grid square. We calculate the area-weighted average of all grid squares within the boundary of a given district to aggregate this measure at the district level. We then use the district-level daily temperature data to construct our primary measure of heat exposure — degree days with a threshold of 24°C, defined as:

$$DD^{24}(T_{dmy}) = \sum (T_{dmy} - 24) \times 1(T_{dmy} > 24),$$

where T_{dmy} is the average temperature on day d, month m, and year y.[3] These daily values of degree days are summed over a 12-month period to arrive at an annual degree day measure that can be paired with birth cohorts. Census data is gathered in February of each census

[2] The five survey rounds and corresponding survey years are as follows: 42nd (1986-1987), 52nd (1995-1996), 64th (2007-2008), 71st (2014) and 75th (2017-2018).

[3] The average diurnal range, or difference between daily maximum temperature and daily minimum temperature in India is 12°C (Rai et al., 2012). Therefore, a day with an average daily temperature of 24°C roughly corresponds to a day with daily maximum temperature of 30°C.

year, so we sum up daily values from March of the previous year through February of the current year.

Earlier literature has found that the relationship between human capital and temperature is nonlinear, with damages being largest when temperatures are especially high (Hu and Li, 2019; Garg et al., 2020; Park et al., 2020, 2021). Daily temperature data in conjunction with a degree day measure allow us to precisely capture this nonlinearity in a parsimonious manner. Our choice of 24°C as a threshold is derived from analysis where we use a flexible temperature bin specification; we describe this in detail in Section 4.1.

We use precipitation data that has been aggregated from hourly time steps to the monthly level. As before, we sum up monthly values of precipitation over a 12-month period (lagged March to current February) to arrive at an annual precipitation measure that can be paired with birth cohorts.

Figure 2 presents long-term climate trends in two key variables: degree days above 24°C (left axis) and precipitation in meters (right axis), with data aggregated into three-year averages between 1950 and 2020. Average exposure to degree days above 24°C has generally increased over time, with a few exceptions. This increase in heat exposure is visibly sharper starting from the early 2000s, consistent with an increased pace of warming in recent decades. In contrast, precipitation demonstrates notable variability, with fluctuations but no clear monotonic trend over the seventy-year horizon.

We construct measures that capture contemporaneous exposure to heat and precipitation during school years, namely ages 6 to 10 for primary school, ages 11 to 14 for middle school, and ages 15 to 16 for secondary school. For example, when evaluating primary school completion for a birth cohort born in 1989, we construct measures of average annual weather from 1995 (age 6) to 1999 (age 10).[4] To further explore dynamic effects, we construct a full set of variables that correspond to distinct stages spanning a child's life cycle — early life (in

[4]To match to five-year birth cohort bins in the educational attainment data, we focus on the middle cohort within each bin and use the middle cohort's school years as the match key. For example, for the 1987–1991 birth cohort, we focus on the middle cohort born in 1989.

utero to age 5), primary school (age 6 to 10), middle school (age 11 to 14), and secondary school (age 15 to 16).

We merge the educational attainment, school enrollment and educational investment data with temperature and precipitation variables to create our three analysis data sets. Appendix Table A.1 presents summary statistics of weather variables and educational attainment for each five-year birth cohort in the analysis sample, which consists of 9 five-year cohorts born between 1952 to 1996 across 425 districts.[5] Appendix Table A.2 reports summary statistics of weather variables and educational investment (panel A) and enrollment (panel B) separately for each age group in the analysis sample.

4 Empirical Strategy

4.1 Educational Attainment

To estimate the impact of heat exposure on educational attainment, we estimate the following panel regression:

$$Y_{cjr} = \beta f(T_{cjr}) + \gamma Precipitation_{cjr} + \alpha_c + \alpha_j + \alpha_{cr} + \epsilon_{cjr}, \tag{6}$$

where Y_{cjr} denotes the log share of individuals in birth cohort c, district j and region r who have completed at least a given schooling level (primary, middle or secondary). Birth cohort c is defined as a five-year bin centered around individuals born in year c.[6]

We first consider contemporaneous exposure to heat and precipitation during school years. Specifically, for primary school completion, $f(T_{cjr})$ represents the average number of degree days above 24°C, scaled by 100, experienced by individuals in cohort c while they are of primary school age (ages 6 to 10).[7] In this specification, β can be interpreted as the percentage

[5] 14 districts from Daman and Diu and Jammu and Kashmir are dropped due to missing Census data and 2 districts from Andaman and Nicobar Islands are dropped due to missing weather data. This leads to a reduction in the number of districts from 441 (as illustrated in Appendix Figure A.1) to 425.

[6] This captures individuals born between years $c - 2$ to $c + 2$.

[7] Age-appropriate exposure windows are 11 to 14 for middle school, and 15 to 16 for secondary school.

change in the share of individuals completing at least primary school associated with an additional hundred degree days above 24°C during school years. The variable $Precipitation_{cjr}$ captures average total precipitation over the same period, standardized using district-specific historical means and standard deviations from 1941 to 1950.[8]

We arrive at a degree day threshold of 24°C using a data-driven approach. We estimate Equation 6 with $f(T_{cjr})$ specified as a vector of temperature bins with 3°C intervals, spanning from below 15°C to above 33°C.[9] Daily average temperature realizations are distributed across these bins. Appendix Figure A.2 presents regression results for each bin on schooling outcomes at primary, middle and secondary school, with 21–24°C as the reference bin. We find that higher temperature exposure significantly reduces grade completion across all three levels of schooling, with effects being particularly pronounced for middle and secondary school completion. Moreover, the relationship is approximately linear above 24°C, motivating our selection of this threshold.

The term α_c represents a vector of birth cohort fixed effects, which control for variation in climate or educational attainment over time. The term α_j represents a vector of district fixed effects, which control for time-invariant district-specific factors that may be correlated with climate or educational outcomes. The term α_{cr} controls for time-varying region-specific effects: either a vector of region-specific linear birth cohort trends or a set of region-birth cohort fixed effects. This term captures unobserved factors that vary across birth cohorts at the regional level, which may be correlated with climate or educational outcomes. Finally, ϵ_{cjr} is an idiosyncratic error term. We cluster our standard errors in two ways. First, we cluster the errors at the district level to allow for arbitrary serial correlation over time within a district. Second, we report Conley standard errors to allow for spatial correlation up to 500 km as well as arbitrary serial correlation (Conley, 1999).

The identifying assumption underlying this panel regression is that, conditional on district

[8] We select this reference period as it pre-dates the first birth cohort in our analysis (1952) and is the earliest data available from ERA5. Results are robust to using a reference period of 30 years, which corresponds to a district's "climate normal."

[9] There are eight bins in total: below 15°C, 15°C–18°C, 18°C–21°C, ..., 30°C–33°C, and over 33°C.

and region-birth cohort fixed effects, residual variation in the heat exposure and precipitation measures is as good as random. As an example, suppose birth cohorts a and b in a given district experience a large difference in heat exposure. Identification compares the difference in outcomes between birth cohorts a and b in that district with a difference in outcomes between birth cohorts a and b in another district in the same region who experience similar heat exposure. The randomness of the residual variation in our heat exposure variable allows us to infer a causal interpretation of β as the effect of heat exposure on educational attainment.

Next, we consider cumulative exposure to heat and precipitation over the educational life cycle of a child, starting from the period in utero. For example, when looking at primary school completion, $f(T_{cjr})$ captures both the average number of degree days above 24 °C during early life years (in which individuals born in year c are in utero to age 5), as well as years in which they are of primary school age (age 6 to 10).[10] As described in Section 2, predictions for the sign of β are theoretically ambiguous — it depends on the dynamics of learning and the carryover of human capital across periods, the effects of long-term income and investment in educational inputs in response to heat exposure over the life cycle.

4.2 Educational Investment

To estimate the impact of heat exposure on educational investment, we estimate a regression of the form:

$$Y_{icjrt} = \beta f(T_{cjrt}) + \gamma Precipitation_{cjrt} + \alpha_c + \alpha_j + \alpha_{rt} + \epsilon_{icjrt}, \tag{7}$$

where Y_{icjrt} denotes the inverse hyperbolic sine (asinh) transformation of educational investment for child i of age c in district j, region r and NSS survey wave t. We consider cumulative exposure to heat and precipitation over the child's educational life cycle, starting from the period in utero. For example, for children of primary school age (ages 6 to 10 at the time

[10]We consider cumulative exposure between the period in utero to age 5, ages 6 to 10 and ages 11 to 14 when we look at middle school completion, and further include exposure between ages 15 to 16 when we look at secondary school completion.

of survey), $f(T_{cjrt})$ consists of two terms: the average annual number of degree days above 24°C experienced from in utero through age 5, and during primary school years.[11]

The terms α_c and α_j represent a vector of age and district fixed effects respectively, controlling for time-invariant age-specific or district-specific factors that may be correlated with climate or educational investment. The term α_{rt} represents a vector of region-NSS year of survey fixed effects, capturing unobserved factors that vary across survey waves at the regional level and may be correlated with climate or educational investment. Finally, ϵ_{icjrt} is an idiosyncratic error term. As before, we present both clustered and Conley standard errors.

4.3 School Enrollment

To estimate the impact of heat exposure on school enrollment, we estimate a similar regression of the form:

$$Y_{cjrt} = \beta f(T_{cjrt}) + \gamma Precipitation_{cjrt} + \alpha_c + \alpha_j + \alpha_{rt} + \epsilon_{cjrt}, \tag{8}$$

where Y_{cjrt} denotes the log share of children of age c in district j and region r in census year t who are enrolled in school. As in Section 4.2, we consider cumulative exposure to heat and precipitation over the life cycle starting from the period in utero.[12]

5 Results

We begin by ascertaining whether there remains sufficient variation in our heat exposure measure after accounting for the full set of fixed effects included in the specifications described in Section 4. Table 1 describes residual variation in heat exposure — that is, the average

[11] Similarly, when looking at children of middle school age (i.e. ages 11 to 14 at the time of survey), $f(T_{cjrt})$ consists of three terms that capture the average annual number of degree days above 24 °C during early life years (in utero to age 5), primary school years (age 6 to 10) and during the years in which they are in middle school.

[12] We truncate our contemporaneous heat exposure measure to end the year before the student is observed in the Census. For example, heat exposure during primary school years for eight-year-old children spans ages six and seven.

annual number of degree days above 24°C experienced during primary (panel A), middle (panel B) and secondary school (panel C), after taking out district (row 1) and cohort fixed effects (row 2), region-cohort trends (row 3) and region-cohort fixed effects (row 4). Column 1 reports standard deviations of the residuals, and columns 2 to 5 report the proportion of observations that have residuals with absolute values exceeding 20, 40, 60 or 80 degree days respectively. Evidence presented in Table 1 suggests that there remains substantial variation in heat exposure after accounting for the full set of fixed effects. Across all three panels, at least 50% of observations have residuals with absolute values exceeding 20 degree days, which corresponds approximately to a 2% increase in heat exposure during a school year.

We now examine implications for educational attainment. Specifically, we estimate the relationship between contemporaneous heat exposure and grade completion using the specification in Equation 6, where $f(T_{cjr})$ captures heat exposure during the years in which individuals are of school age. In terms of our model, this estimate informs us about the short-run effect of weather, i.e. $\frac{\partial A_k}{\partial \phi_k}$. We report results in Table 2. For each outcome, the first column controls for region-cohort trends, while the second column controls for region-by-cohort fixed effects.

We find that greater heat exposure during school years significantly reduces grade completion across all levels of schooling. An increase in contemporaneous heat exposure during school years — as measured by a hundred degree day increase — results in a 2.3% decline in primary school completion (column 2), a 3.3% decline in middle school completion (column 4) and a 3.8% decline in secondary school completion (column 6). To put these figures into perspective, annual mean exposure to degree days above 24°C is close to 1000 in our sample (as summarized in Appendix Table A.1). This implies that an approximate 10% increase in heat exposure during a school year leads to a 2-4% reduction in grade completion. In contrast, the effects of precipitation on grade completion are generally weaker: we observe a small, marginally significant positive effect on middle school completion, and no statistically significant effect on primary or secondary school completion.

To assess the robustness of these findings, we carry out three additional checks. First, we run our specification with outcome variables in levels rather than logs. As shown in Appendix Table A.3, results remain similar in both magnitude and precision to those reported in Table 2. Second, we explore sensitivity of our estimates to the choice of degree day threshold by varying it by $+/-3$°C from the baseline of 24°C. Appendix Figure A.4 plots coefficient estimates across these threshold values ranging from 21°C to 27°C. The estimated effects are remarkably stable across these thresholds for all three levels of schooling, suggesting that our main findings are not driven by the specific threshold choice of 24°C. Third, demographic changes driven by migration could bias our estimates if selection into the sample is correlated with education levels.[13] To test for selection, we examine how cohort size by education level responds to heat exposure in Appendix Table A.4.[14] We find that individuals with at least middle and secondary schooling are more likely to be observed in a district with greater heat exposure, which suggests that the direction of sample selection is positive. Our reduced-form estimates of the impacts of heat exposure on educational attainment are thus plausible under-estimates of the true effect and can be interpreted as a conservative lower bound.

We next expand our analysis to a dynamic framework that considers heat exposure over a child's life cycle — starting from the period in utero — on grade completion. We use the specification in Equation 6 with $f(T_{cjr})$ capturing heat exposure both during early life years as well as years in which individuals are of school age. These effects are meant to capture

[13] For example, if higher-educated individuals are more likely to out-migrate in response to heat exposure, the remaining observed population in a district would comprise of fewer educated individuals, and the estimated effect of heat exposure on educational attainment would be an overestimate of the true effect. Alternatively, lower-educated individuals may be particularly vulnerable to heat since they are more likely to engage in agriculture and have limited access to adaptation mechanisms e.g. cooling technologies. If lower-educated individuals are more likely to out-migrate in response to heat exposure, the remaining observed population in a district would comprise more educated individuals and the estimated effect of heat exposure on educational attainment would an underestimate of the true effect.

[14] For each five-year birth cohort born between 1952 and 1971, we observe population size by education level across the three Census waves: 1991, 2001, and 2011. In order to test whether cohort size changes in response to heat exposure vary by education levels, we construct our dependent variable to be the share of a cohort in a census year that has completed a certain grade. We regress this measure on heat exposure and precipitation in the preceding ten years with cohort, district, and region-by-year fixed effects. In this setup, the estimated effects capture whether a segment of the cohort – for example, those who have at least completed middle school – experiences population change at a differential rate relative to the entire cohort.

a shift in the distribution of weather in the model, i.e. $\frac{\partial A_2}{\partial \bar{w}}$. Table 3 reports coefficient estimates, which are also illustrated in Figure 3.

Our results indicate that the adverse effects of heat exposure on grade completion are primarily attributable to contemporaneous heat exposure during school years. Heat exposure in early life appears to have no significant impact on primary school completion (coefficient 0.007, p-value = 0.573). The magnitude of the coefficient for heat exposure during primary school years in Table 3 column 2 is nearly identical to that in Table 2 column 2, and an F-test of equality of coefficients across the two models yields a p-value of 0.318. Similarly, heat exposure during primary school years does not significantly affect middle school completion (coefficient -0.003, p-value = 0.786), and the coefficients for heat exposure during middle school years are statistically indistinguishable across the two models (p-value=0.373). For secondary school completion, heat exposure during primary school years is insignificant (coefficient 0.015, p-value = 0.318), whereas heat exposure during middle school years has a significant negative impact (coefficient -0.025, p-value= 0.030).

Interestingly, we find that heat exposure in early life is *positively* associated with grade completion, particularly for middle (coefficient 0.026, p-value = 0.017) and secondary school (coefficient 0.055, p-value = 0.000). A possible explanation is that parents make compensatory educational investments in response to early-life heat shocks (Yi et al., 2015). Our model in Section 2 outlines how long-run effects of heat exposure on educational attainment depend crucially on how parental investment responds to climate shocks.

We empirically investigate parental responses to heat exposure over the life cycle using nationally representative household survey data from India's National Sample Survey. More specifically, we explore how households' expenditure on education for each child responds to heat exposure both during early life as well as during years in which individuals are of school age. We conduct this analysis with all children of primary to secondary school age (i.e. between the ages of 7 to 16) who are enrolled in school, and present results in Table 4.

Two patterns emerge. First, we find that parental investment in education increases in response to greater heat exposure during school years. A hundred degree day increase in average heat exposure during primary school years leads to a 10.8% increase in total education expenditure for students enrolled in primary school, with the effect significant at the 1% level. Heat exposure during middle and secondary school years also leads to increases in total education expenditure for students enrolled in middle and secondary school, although effect sizes are smaller at 3.2% and 2.8% respectively. Parents thus respond to greater contemporaneous heat exposure by increasing educational investment.

Second, we find that educational investment responds substantially to greater heat exposure in early life. A hundred degree day increase in average heat exposure during early life increases total education expenditure by 11.6%, 11.8%, and 7.3% during primary, middle, and secondary school years respectively. Interestingly, this response is even larger than responses to heat exposure during contemporaneous school years, in some cases. The magnitude of parental responses is significantly larger for early life heat exposure relative to contemporaneous heat exposure for middle and secondary schooling, as reflected in p-values from an F-test of equality of coefficients (reported at the bottom of Table 4). This persistence in increased parental investment in response to early life heat exposure can potentially explain our earlier finding that heat exposure in early life is positively associated with middle and secondary school completion (Table 3).

Besides parental responses, another factor that may contribute to the dynamic effects of heat exposure on educational attainment in Figure 3 is the impact on school enrollment, which can in turn delay or disrupt subsequent grade completion. We explore how school enrollment responds to heat exposure over the life cycle using Census data. We conduct this analysis with all children who are of primary to secondary school age (i.e. between the ages of 7 to 16) at the time of the Census, and present results in Table 5.

We find that contemporaneous heat exposure results in a decline in school enrollment across all three levels of schooling: a hundred degree day increase in average heat exposure during school years reduces the share of individuals enrolled in school by 3 to 5%. Further, these effects persist in subsequent educational stages. For example, increased heat exposure during primary school years not only lowers enrollment at the primary level but also negatively impacts middle school enrollment. A hundred degree day increase in average heat exposure during primary school corresponds to a 4.6% decline in primary school enrollment and a 3.3% reduction in middle school enrollment, as shown in Columns 1 and 2. Likewise, a hundred degree day increase in average heat exposure during middle school results in a 3.2% decline in middle school enrollment and a 6.2% reduction in secondary school enrollment (Columns 2 and 3). This persistence may arise due to the dynamic nature of human capital accumulation described in Section 2: individuals who fall behind or drop out at an earlier grade are less likely to re-enroll or progress to higher grades, thereby compounding the effects of prior heat exposure. These findings highlight a second dynamic channel – beyond contemporaneous effects – through which heat exposure impacts educational attainment over time.

Taken together, our results demonstrate that parental investment in education only partially mitigates the adverse effects of heat exposure on grade completion. These compensatory investments are most effective when shocks occur early in a child's life, thus providing time for parents to offset potential learning losses. However, increased parental investments during school years do not fully mitigate the detrimental effects of contemporaneous heat exposure. Parental responses are thus insufficient in fully shielding children from the cumulative detrimental effects of heat exposure experienced throughout the life cycle.

6 Conclusion

Rising temperatures under climate change have significant adverse impacts on a range of economic outcomes in LMICs. In this paper, we document detrimental long-term impacts of heat exposure on human capital formation in India, focusing on the dynamic interplay between weather shocks, educational attainment, and parental investments across the life cycle. We find strong evidence that contemporaneous heat exposure significantly reduces grade completion across primary, middle and secondary schooling. In contrast, we find limited evidence that heat exposure in early life results in persistent adverse effects on grade completion. In fact, early-life heat exposure is sometimes associated with higher school completion rates, likely due to compensatory parental investments in education.

Our findings carry several important implications. They highlight the importance of studying climate impacts over long time frames, particularly for outcomes that form dynamically over time. Our analysis reveals that parents actively respond to both early life and contemporaneous heat exposure by increasing educational expenditures. These compensatory investments are most effective when shocks occur early in a child's life, as they provide time for parents to offset potential learning losses, while increased parental investments during school years are not able to fully mitigate the detrimental effects of contemporaneous heat exposure. Moreover, we document that heat exposure reduces school enrollment and these effects propagate to subsequent grades, revealing the dynamic and compounding nature of climate impacts on human capital accumulation.

Our findings also highlight the urgent need for policies aimed at protecting vulnerable children from the cumulative damages of heat exposure under climate change. Potential policy responses could range from individual-level interventions such as providing financial support for educational expenses or remedial education programs, to school- and community-level interventions such as large-scale investment in cooling infrastructure. As the frequency and severity of heat events continue to rise, mitigating their long-term consequences is central to promoting equitable and sustainable development across LMICs.

References

Adhvaryu, A., Molina, T., Nyshadham, A., and Tamayo, J. (2024). Helping children catch up: Early life shocks and the progresa experiment. *The Economic Journal*, 134(657):1–22.

Ahsan, M. N., Banerjee, R., and Hari, S. (2018). Social promotion and learning outcomes: Evidence from india. *USC-INET Research Paper*, (18-14).

Akee, R., Copeland, W., Costello, E. J., and Simeonova, E. (2018). How does household income affect child personality traits and behaviors? *American Economic Review*, 108(3):775–827.

Alberto, I. C., Jiao, Y., and Zhang, X. (2021). Too hot or too cold to study? The effect of temperature on student time allocation. *Economics of Education Review*, 84:102152.

Almond, D. and Mazumder, B. (2013). Fetal origins and parental responses. *Annual Review of Economics*, 5(1):37–56.

Aragón, F. M., Oteiza, F., and Rud, J. P. (2021). Climate change and agriculture: Subsistence farmers' response to extreme heat. *American Economic Journal: Economic Policy*, 13(1):1–35.

Becker, G. S. and Tomes, N. (1976). Child endowments and the quantity and quality of children. *Journal of Political Economy*, 84(4, Part 2):S143–S162.

Behrman, J. R., Pollak, R. A., and Taubman, P. (1982). Parental preferences and provision for progeny. *Journal of Political Economy*, 90(1):52–73.

Behrman, J. R., Rosenzweig, M. R., and Taubman, P. (1994). Endowments and the allocation of schooling in the family and in the marriage market: The twins experiment. *Journal of Political Economy*, 102(6):1131–1174.

Bharadwaj, P., Khwaja, A. I., and Mian, A. (2015). Population exchange and its impact on literacy, occupation and gender–evidence from the partition of India. *International Migration*, 53(4):90–106.

Burke, M., Hsiang, S. M., and Miguel, E. (2015). Global non-linear effect of temperature on economic production. *Nature*, 527(7577):235–239.

Cho, H. (2017). The effects of summer heat on academic achievement: A cohort analysis. *Journal of Environmental Economics and Management*, 83:185–196.

Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics*, 92(1):1–45.

Cunha, F., Heckman, J. J., and Schennach, S. M. (2010). Estimating the technology of cognitive and noncognitive skill formation. *Econometrica*, 78(3):883–931.

Dell, M., Jones, B. F., and Olken, B. A. (2012). Temperature shocks and economic growth: Evidence from the last half century. *American Economic Journal: Macroeconomics*, 4(3):66–95.

Dell, M., Jones, B. F., and Olken, B. A. (2014). What do we learn from the weather? The new climate–economy literature. *Journal of Economic Literature*, 52(3):740–798.

Diffenbaugh, N. S. and Burke, M. (2019). Global warming has increased global economic inequality. *Proceedings of the National Academy of Sciences*, 116(20):9808–9813.

ECMWF (2023). ERA5 reanalysis data. `https://www.ecmwf.int/en/forecasts/dataset/ecmwf-reanalysis-v5`. Accessed: 2023-08-19.

Fan, W. and Porter, C. (2020). Reinforcement or compensation? Parental responses to children's revealed human capital levels. *Journal of Population Economics*, 33(1):233–270.

Fishman, R., Carrillo, P., and Russ, J. (2019). Long-term impacts of exposure to high temperatures on human capital and economic productivity. *Journal of Environmental Economics and Management*, 93:221–238.

Frijters, P., Johnston, D. W., Shah, M., and Shields, M. A. (2009). To work or not to work? Child development and maternal labor supply. *American Economic Journal: Applied Economics*, 1(3):97–110.

Garg, T., Jagnani, M., and Taraz, V. (2020). Temperature and human capital in India. *Journal of the Association of Environmental and Resource Economists*, 7(6):1113–1150.

Graff Zivin, J., Hsiang, S. M., and Neidell, M. (2018). Temperature and human capital in the short and long run. *Journal of the Association of Environmental and Resource Economists*, 5(1):77–105.

Graff Zivin, J., Song, Y., Tang, Q., and Zhang, P. (2020). Temperature and high-stakes cognitive performance: Evidence from the national college entrance examination in China. *Journal of Environmental Economics and Management*, 104:102365.

Halla, M. and Zweimüller, M. (2014). Parental response to early human capital shocks: Evidence from the Chernobyl accident.

Hsiang, S. M., Burke, M., and Miguel, E. (2013). Quantifying the influence of climate on human conflict. *Science*, 341(6151):1235367.

Hu, Z. and Li, T. (2019). Too hot to handle: The effects of high temperatures during pregnancy on adult welfare outcomes. *Journal of Environmental Economics and Management*, 94:236–253.

Leight, J. (2017). Sibling rivalry: Endowment and intrahousehold allocation in Gansu Province, China. *Economic Development and Cultural Change*, 65(3):457–493.

Liu, M., Shamdasani, Y., and Taraz, V. (2023). Climate change and labor reallocation: Evidence from six decades of the Indian Census. *American Economic Journal: Economic Policy*, 15(2):395–423.

Maccini, S. and Yang, D. (2009). Under the weather: Health, schooling, and economic consequences of early-life rainfall. *American Economic Review*, 99(3):1006–26.

Magnuson, K., Duncan, G., Yoshikawa, H., Yoo, P., Han, S., Gennetian, L. A., Fox, N., Halpern-Meekin, S., and Noble, K. (2024). Can Cash Transfers Improve Maternal Well-being and Family Processes among Families with Young Children? An Experimental Analysis. *Working paper.*

Park, R. J. (2022). Hot temperature and high-stakes performance. *Journal of Human Resources*, 57(2):400–434.

Park, R. J., Behrer, A. P., and Goodman, J. (2021). Learning is inhibited by heat exposure, both internationally and within the United States. *Nature Human Behaviour*, 5(1):19–27.

Park, R. J., Goodman, J., Hurwitz, M., and Smith, J. (2020). Heat and learning. *American Economic Journal: Economic Policy*, 12(2):306–39.

Pavanello, F., De Cian, E., Davide, M., Mistry, M., Cruz, T., Bezerra, P., Jagu, D., Renner, S., Schaeffer, R., and Lucena, A. F. (2021). Air-conditioning and the adaptation cooling deficit in emerging economies. *Nature Communications*, 12(1):1–11.

Rai, A., Joshi, M. K., and Pandey, A. (2012). Variations in diurnal temperature range over India: Under global warming scenario. *Journal of Geophysical Research: Atmospheres*, 117(D2).

Rosenzweig, M. R. and Zhang, J. (2009). Do population control policies induce more human capital investment? Twins, birth weight and China's one-child policy. *The Review of Economic Studies*, 76(3):1149–1174.

Schlenker, W. and Roberts, M. J. (2009). Nonlinear temperature effects indicate severe damages to US crop yields under climate change. *Proceedings of the National Academy of Sciences*, 106(37):15594–15598.

Sexton, S., Wang, Z., and Mullins, J. T. (2022). Heat adaptation and human performance in a warming climate. *Journal of the Association of Environmental and Resource Economists*, 9(1):141–163.

Shah, M. and Steinberg, B. M. (2017). Drought of opportunities: Contemporaneous and long-term impacts of rainfall shocks on human capital. *Journal of Political Economy*, 125(2):527–561.

Somanathan, E., Somanathan, R., Sudarshan, A., and Tewari, M. (2021). The impact of temperature on productivity and labor supply: Evidence from Indian manufacturing. *Journal of Political Economy*, 129(6):1797–1827.

Stowell, J. D., Sun, Y., Spangler, K. R., Milando, C. W., Bernstein, A., Weinberger, K. R., Sun, S., and Wellenius, G. A. (2022). Warm-season temperatures and emergency department visits among children with health insurance. *Environmental Research: Health*, 1(1):015002.

Taraz, V. (2018). Can farmers adapt to higher temperatures? Evidence from India. *World Development*, 112:205–219.

Todd, P. E. and Wolpin, K. I. (2003). On the specification and estimation of the production function for cognitive achievement. *The Economic Journal*, 113(485):F3–F33.

Wu, J., Lin, J., and Han, X. (2023). Compensation for girls in early childhood and its long-run impact: Family investment strategies under rainfall shocks. *Journal of Population Economics*, 36(3):1225–1268.

Yi, J., Heckman, J. J., Zhang, J., and Conti, G. (2015). Early health shocks, intra-household resource allocation and child outcomes. *The Economic Journal*, 125(588):F347–F371.

Zimmermann, L. (2020). Remember when it rained: Schooling responses to shocks in India. *World Development*, 126:104705.

Figures

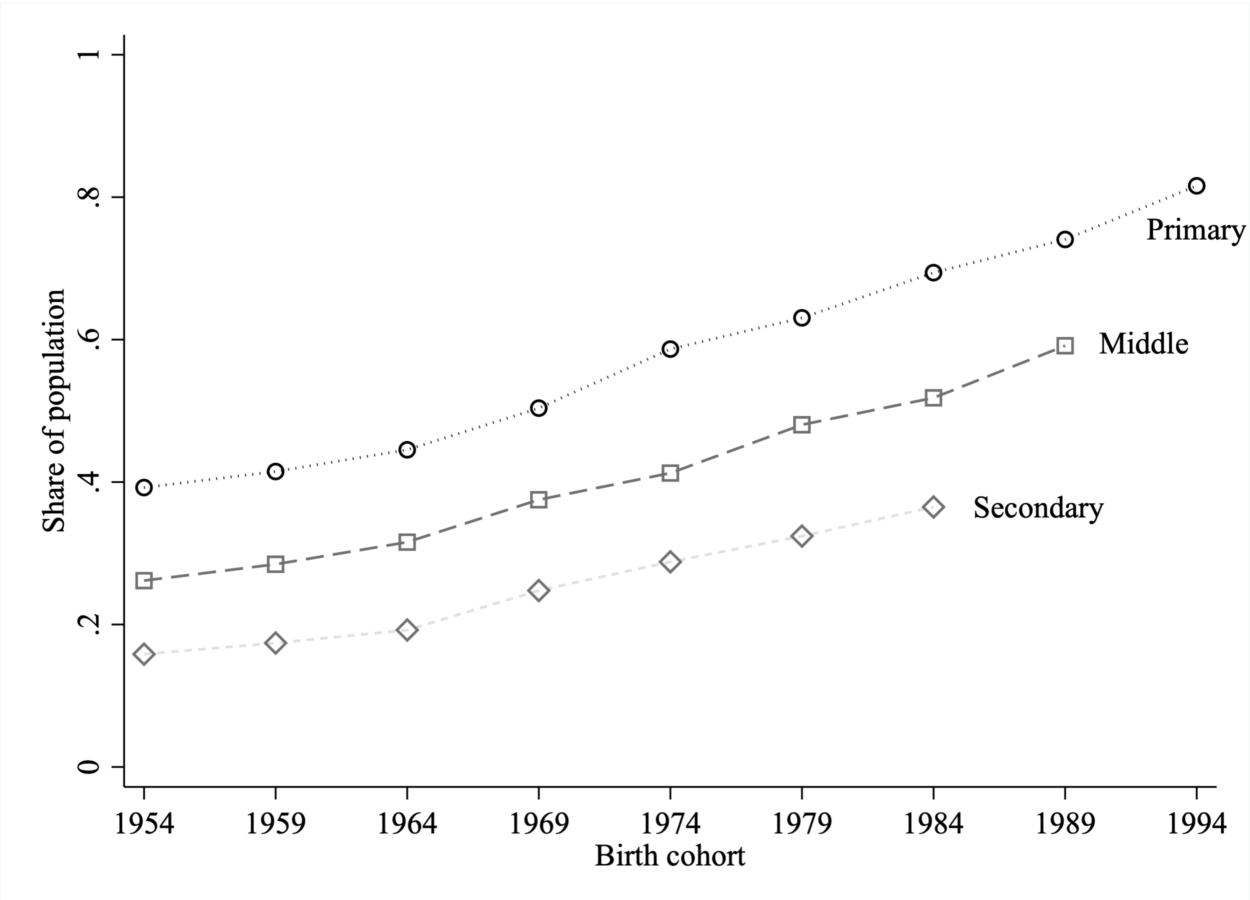

Figure 1: Educational attainment over time

Figure illustrates grade completion at each level of schooling (primary, middle and secondary) for all 5-year birth cohorts in our sample, with the horizontal axis indicating the birth year of the middle age cohort within each bin.

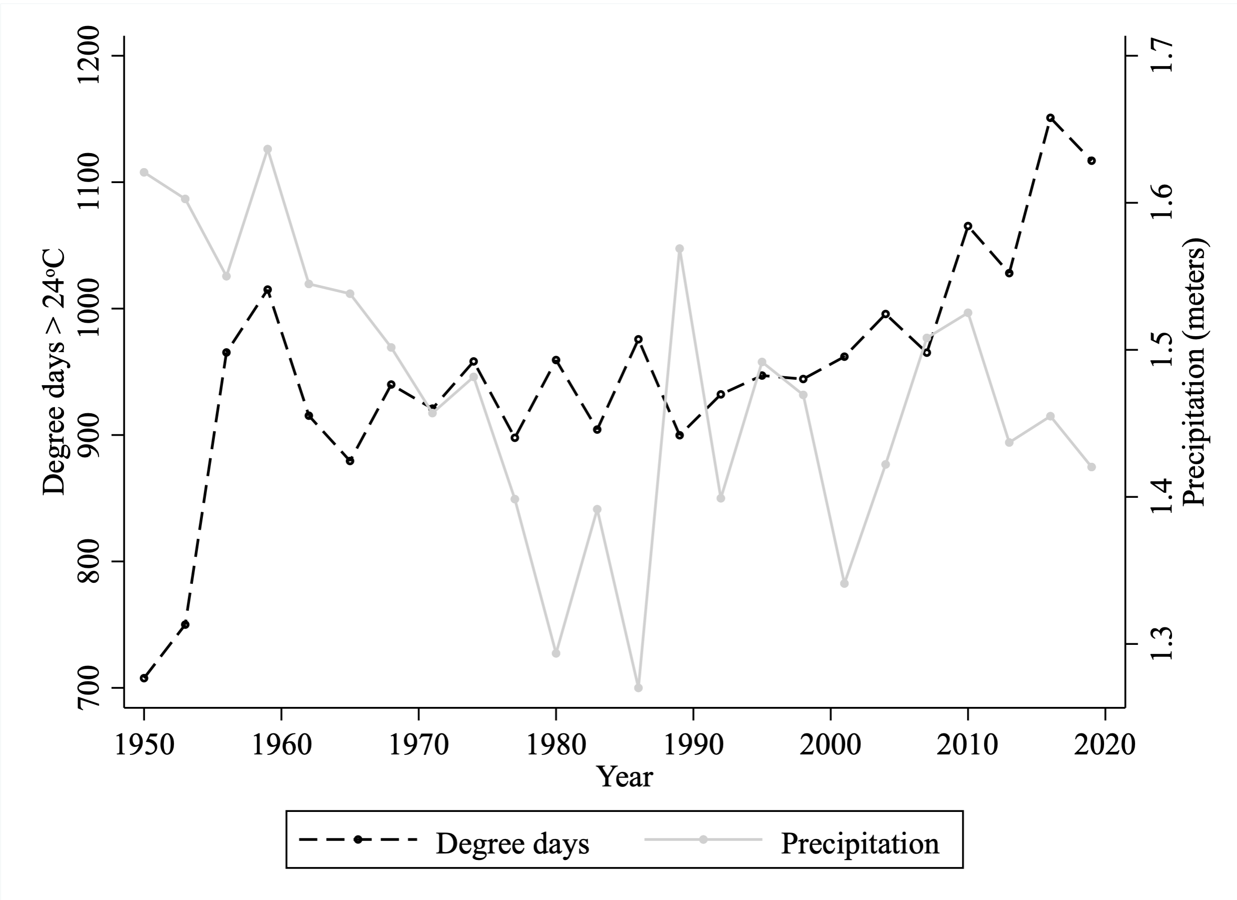

Figure 2: Climate trends over time

Figure illustrates annual degree days above 24°C on the left axis and annual precipitation in meters on the right axis, binned into 3-year averages between 1950 to 2019.

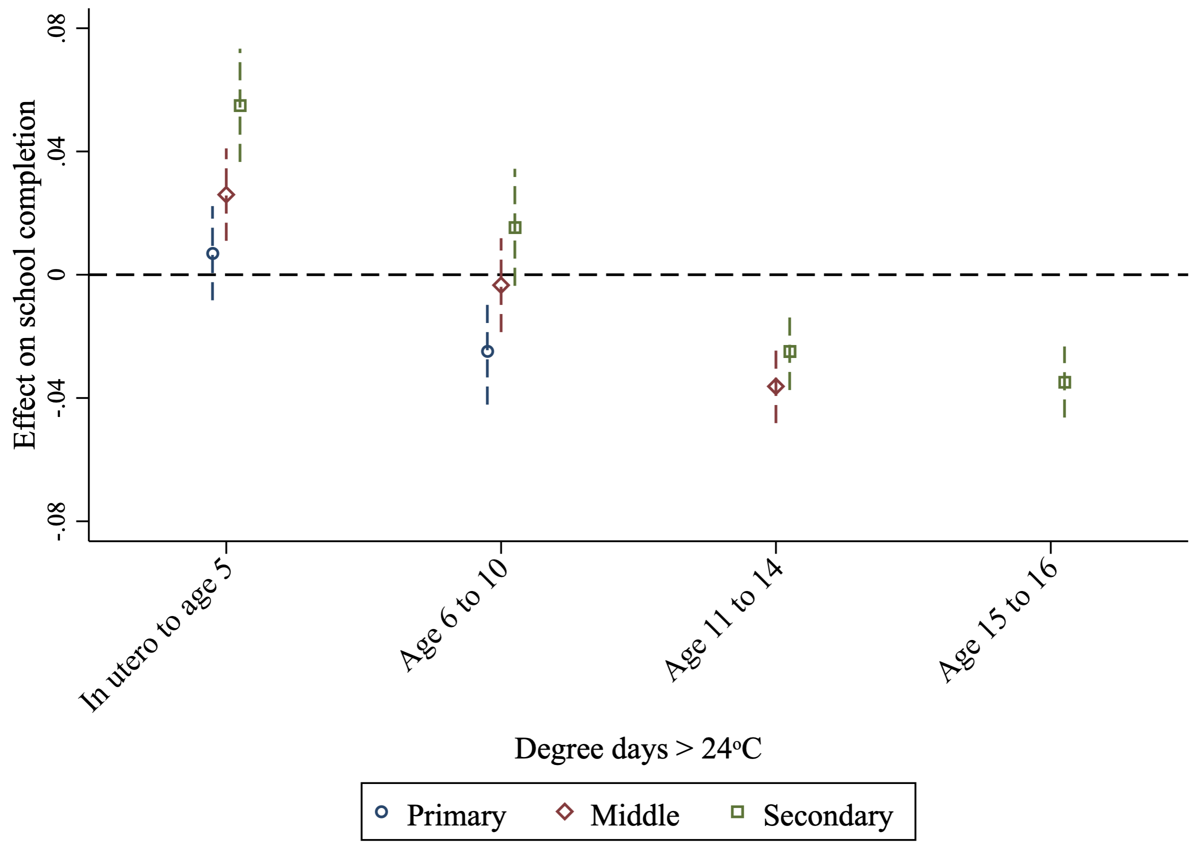

Figure 3: Heat exposure over the life cycle & educational attainment

Figure plots coefficient estimates that capture impacts of heat exposure during the life cycle on grade completion using the panel fixed effects model with region-by-cohort fixed effects (Equation 6). Standard errors are clustered at the district level, and 95% confidence intervals are illustrated with dashed lines.

Tables

Table 1: Residual variation in heat exposure

	σ_e	$\lvert e\rvert > 20$	$\lvert e\rvert > 40$	$\lvert e\rvert > 60$	$\lvert e\rvert > 80$
Panel A: Degree days $>$ 24°C, *age 6 to 10*	(1)	(2)	(3)	(4)	(5)
District FE	48.3	0.59	0.34	0.19	0.11
District FE + Cohort FE	42.9	0.60	0.31	0.15	0.07
District FE + Cohort FE + Region-cohort trends	38.4	0.55	0.26	0.12	0.05
District FE + Region-cohort FE	33.3	0.50	0.21	0.07	0.02
Panel B: Degree days $>$ 24°C, *age 11 to 14*					
District FE	49.2	0.58	0.36	0.20	0.11
District FE + Cohort FE	45.5	0.59	0.31	0.17	0.09
District FE + Cohort FE + Region-cohort trends	43.2	0.56	0.30	0.16	0.08
District FE + Region-cohort FE	37.0	0.53	0.25	0.10	0.04
Panel C: Degree days $>$ 24°C, *age 15 to 16*					
District FE	74.6	0.69	0.51	0.36	0.26
District FE + Cohort FE	57.1	0.71	0.48	0.27	0.13
District FE + Cohort FE + Region-cohort trends	54.8	0.71	0.46	0.26	0.12
District FE + Region-cohort FE	43.5	0.56	0.30	0.15	0.08

Note: Table summarizes residual variation in our measure of heat exposure during school years – the annual average number of degree days above 24°C during the years in which individuals are age-appropriate for the corresponding level of schooling – from a regression with fixed effects and/or controls. Each row lists the relevant set of fixed effects and controls included in each regression. Column 1 reports the standard deviation of the residuals (denoted as e). Columns 2 to 5 report the proportion of observations that have a residual e with absolute values exceeding 20, 40, 60 or 80 degree days.

Table 2: Contemporaneous heat exposure & educational attainment

	Grade Completion					
	Primary		Middle		Secondary	
	(1)	(2)	(3)	(4)	(5)	(6)
Degree days > 24°C	-0.018	-0.023	-0.030	-0.033	-0.025	-0.038
	(0.007)**	(0.010)**	(0.006)***	(0.007)***	(0.005)***	(0.006)***
	[0.013]	[0.014]	[0.010]***	[0.011]***	[0.011]**	[0.010]***
Precipitation z-score	-0.002	0.007	-0.006	0.014	-0.010	-0.013
	(0.006)	(0.008)	(0.004)	(0.006)**	(0.005)**	(0.006)**
	[0.010]	[0.012]	[0.007]	[0.008]*	[0.009]	[0.009]
Region-cohort trends	Y	N	Y	N	Y	N
Region-cohort FE	N	Y	N	Y	N	Y
Dep. variable mean (in levels)	0.580	0.580	0.405	0.405	0.250	0.250
Observations	3,825	3,825	3,400	3,400	2,975	2,975

Note: Table presents regression estimates of Equation 6, where $f(T_{cjr})$ captures contemporaneous exposure to heat during school years. More specifically, $f(T_{cjr})$ captures the annual average number of degree days above 24°C during the years in which individuals are age-appropriate for the corresponding level of schooling. The dependent variable is the natural logarithm of the share of individuals in a birth cohort who have completed at least primary school (columns 1 and 2), middle school (columns 3 and 4) and secondary school (columns 5 and 6). All columns include controls for precipitation as well as district and birth cohort fixed effects. We present standard errors clustered by district in parentheses, and Conley standard errors that allow for spatial correlation up to 500 km and arbitrary serial correlation in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3: Heat exposure over the life cycle & educational attainment

	Grade Completion					
	Primary		Middle		Secondary	
	(1)	(2)	(3)	(4)	(5)	(6)
Degree days > 24°C, in utero to age 5	0.005	0.007	0.016	0.026	0.040	0.055
	(0.006)	(0.008)	(0.006)**	(0.008)***	(0.008)***	(0.009)***
	[0.011]	[0.012]	[0.010]	[0.011]**	[0.013]***	[0.013]***
Degree days > 24°C, age 6 to 10	-0.019	-0.025	-0.001	-0.003	0.011	0.015
	(0.007)***	(0.009)***	(0.007)	(0.008)	(0.009)	(0.010)
	[0.013]	[0.014]*	[0.012]	[0.012]	[0.017]	[0.015]
Degree days > 24°C, age 11 to 14			-0.029	-0.036	-0.022	-0.025
			(0.005)***	(0.006)***	(0.006)***	(0.006)***
			[0.010]***	[0.011]***	[0.013]*	[0.011]**
Degree days > 24°C, age 15 to 16					-0.029	-0.035
					(0.005)***	(0.006)***
					[0.010]***	[0.010]***
Region-cohort trends	Y	N	Y	N	Y	N
Region-cohort FE	N	Y	N	Y	N	Y
Dep. variable mean (in levels)	0.580	0.580	0.405	0.405	0.250	0.250
Observations	3,825	3,825	3,400	3,400	2,975	2,975

Note: Table presents regression estimates of Equation 6, where $f(T_{cjr})$ captures cumulative exposure to heat over the life cycle. More specifically, $f(T_{cjr})$ is a vector of variables that capture the annual average number of degree days above 24°C during the years in which individuals are in utero to age 5, ages 6 to 10 etc. The dependent variable is the natural logarithm of the share of individuals in a birth cohort who have completed at least primary school (columns 1 and 2), middle school (columns 3 and 4) and secondary school (columns 5 and 6). All columns include controls for precipitation as well as district and birth cohort fixed effects. We present standard errors clustered by district in parentheses, and Conley standard errors that allow for spatial correlation up to 500 km and arbitrary serial correlation in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Heat exposure over the life cycle & educational investment

	Household Expenditure on Education		
	Primary	Middle	Secondary
	(1)	(2)	(3)
Degree days > 24° C, in utero to age 5	0.116	0.118	0.073
	(0.023)***	(0.022)***	(0.023)***
	[0.034]***	[0.034]***	[0.025]***
Degree days > 24° C, age 6 to 10	0.108	0.042	-0.010
	(0.018)***	(0.018)**	(0.023)
	[0.025]***	[0.020]**	[0.024]
Degree days > 24° C, age 11 to 14		0.032	0.007
		(0.012)***	(0.019)
		[0.017]*	[0.017]
Degree days > 24° C, age 15 to 16			0.028
			(0.011)***
			[0.012]**
Pval: DD(in utero to age 5) = DD(age 6 to 10)	0.795		
Pval: DD(in utero to age 5) = DD(age 11 to 14)		0.000	
Pval: DD(in utero to age 5) = DD(age 15 to 16)			0.058
Observations	148,183	151,116	62,595

Note: Table presents regression estimates of Equation 7, where $f(T_{cjrt})$ captures cumulative exposure to heat over the life cycle. The sample consists of children between the ages of 7 to 10 (i.e. primary school age) in Column 1, children between the ages of 11 to 14 (i.e. middle school age) in Column 2, and children between the ages of 15 to 16 (i.e. secondary school age) in Column 3. The dependent variable is the inverse hyperbolic sine (asinh) transformation of total household expenditure on education in Indian rupees. All columns include controls for precipitation as well as age, district and region x year of survey fixed effects. We present standard errors clustered by district in parentheses, and Conley standard errors that allow for spatial correlation up to 500 km and arbitrary serial correlation in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Heat exposure over the life cycle & school enrollment

	Enrollment		
	Primary	Middle	Secondary
	(1)	(2)	(3)
Degree days > 24° C, in utero to age 5	0.003	-0.009	-0.036
	(0.010)	(0.008)	(0.013)***
	[0.018]	[0.013]	[0.019]*
Degree days > 24° C, age 6 to 10	-0.046	-0.033	0.021
	(0.009)***	(0.008)***	(0.011)*
	[0.015]***	[0.012]***	[0.016]
Degree days > 24° C, age 11 to 14		-0.032	-0.062
		(0.004)***	(0.010)***
		[0.008]***	[0.016]***
Degree days > 24° C, age 15 to 16			-0.031
			(0.005)***
			[0.008]***
Dep. variable mean (in levels)	0.749	0.753	0.578
Observations	4,836	4,836	2,418

Note: Table presents regression estimates of Equation 8, where $f(T_{cjrt})$ captures exposure to heat over the life cycle. The sample consists of children between the ages of 7 to 10 (i.e. primary school age) in Column 1, children between the ages of 11 to 14 (i.e. middle school age) in Column 2, and children between the ages of 15 to 16 (i.e. secondary school age) in Column 3. The dependent variable is the natural logarithm of the share of individuals who are enrolled in school. All columns include controls for precipitation as well as age, district and region x year of census fixed effects. We present standard errors clustered by district in parentheses, and Conley standard errors that allow for spatial correlation up to 500 km and arbitrary serial correlation in brackets. * p < 0.10, ** p < 0.05, *** p < 0.01

Appendix Figures

Figure A.1: Consistent district boundaries

Figure illustrates 441 consistent district boundaries that span the same area between 1991 and 2011. The various splits and boundary changes between 1991 and 2011 can be deduced from the gray boundaries that trace out the 2011 Census districts delineation.

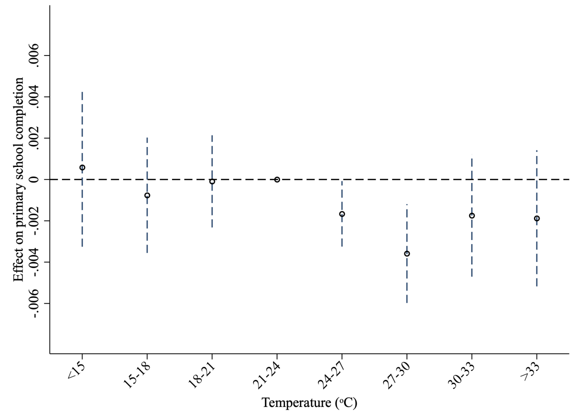

(a) Primary

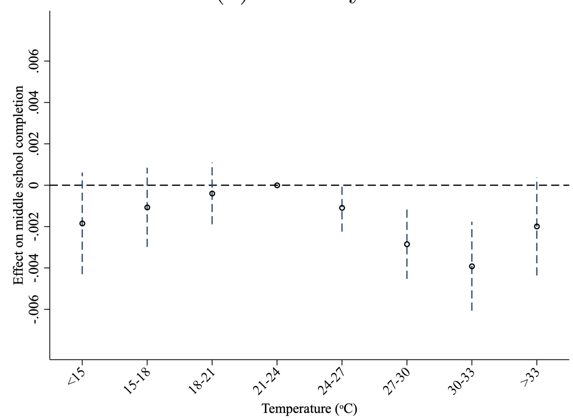

(b) Middle

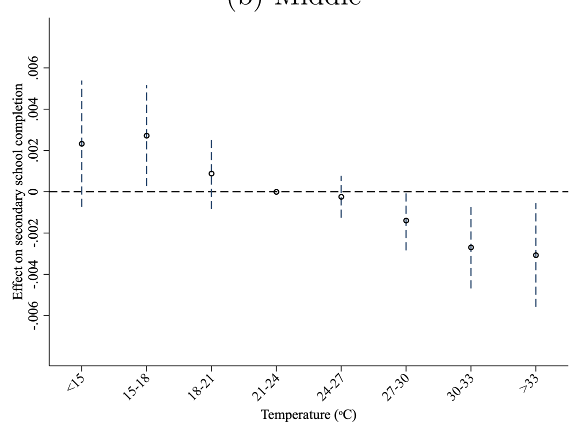

(c) Secondary

Figure A.2: Binned regression illustrating relationship between heat exposure and educational attainment

Figure illustrates coefficient estimates of the effect of contemporaneous heat exposure on grade completion, with 95% confidence intervals. Each regression includes a vector of variables for temperature exposure during school years in 3°C bins from 15°C to 33°C, with 21°C-24°C as the reference category, a linear control for precipitation, region-cohort fixed effects, and Conley standard errors.

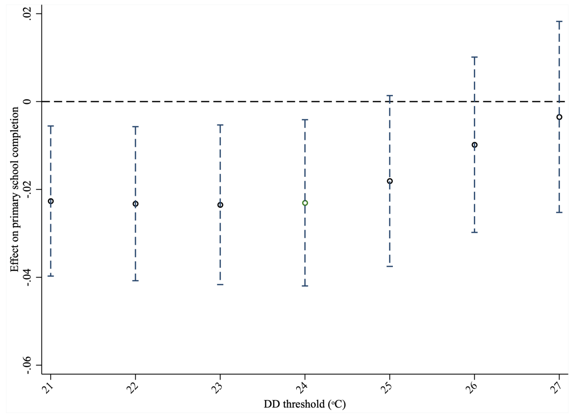

(a) Primary

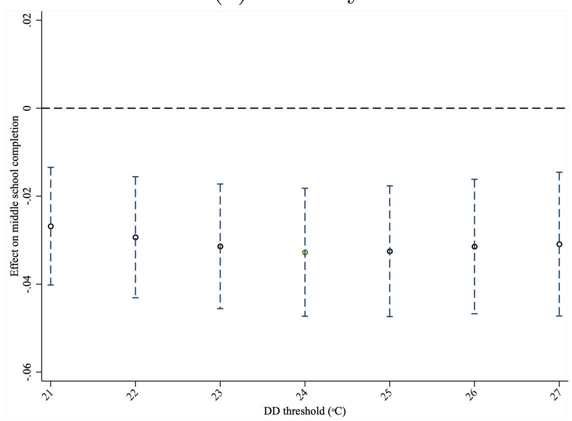

(b) Middle

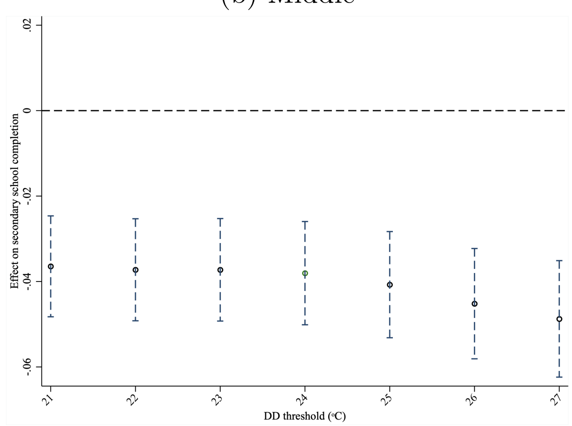

(c) Secondary

Figure A.4: Alternate degree day thresholds

Figure plots coefficient estimates that capture impacts of contemporaneous heat exposure during school years on primary (panel a), middle (panel b) and secondary (panel c) school completion respectively using the panel fixed effects model with region-by-cohort fixed effects (Equation 6) at varying degree day thresholds ($+/-3$ relative to the baseline specification of 24°C). Standard errors are clustered at the district level, and 95% confidence intervals are illustrated with dashed lines.

Appendix Tables

Table A.1: Trends in climate and educational attainment

	1954	1959	1964	1969	1974	1979	1984	1989	1994	Total
Degree days > 24°C, in utero to age 5	880.2	1001.5	917.5	954.8	978.6	954.3	964.8	962.3	951.7	951.7
	(445.3)	(477.8)	(464.2)	(474.7)	(481.3)	(463.1)	(467.0)	(468.0)	(455.5)	(467.1)
Degree days > 24°C, age 6 to 10	990.5	924.6	974.0	952.6	951.7	973.1	939.6	965.5	998.0	963.3
	(478.1)	(471.2)	(482.3)	(468.8)	(457.7)	(471.7)	(458.3)	(459.4)	(482.9)	(470.2)
Degree days > 24°C, age 11 to 14	925.4	941.3	969.2	966.6	942.6	932.4	966.3	980.6	1014.5	959.9
	(476.3)	(474.7)	(478.2)	(468.1)	(456.4)	(455.6)	(462.2)	(475.7)	(478.7)	(469.9)
Degree days > 24°C, age 15 to 16	922.8	1039.3	919.5	922.1	1034.2	954.1	963.9	1032.7	988.5	975.2
	(463.6)	(501.9)	(452.4)	(441.5)	(506.1)	(465.7)	(457.7)	(499.7)	(463.9)	(474.9)
Share completed primary school	0.392	0.415	0.445	0.504	0.587	0.631	0.694	0.741	0.816	0.580
	(0.149)	(0.160)	(0.161)	(0.164)	(0.161)	(0.149)	(0.144)	(0.126)	(0.0993)	(0.205)
Share completed middle school	0.262	0.285	0.316	0.375	0.413	0.480	0.518	0.591		0.405
	(0.110)	(0.122)	(0.127)	(0.139)	(0.139)	(0.149)	(0.154)	(0.151)		(0.176)
Share completed secondary school	0.158	0.174	0.192	0.248	0.288	0.324	0.365			0.250
	(0.0757)	(0.0819)	(0.0829)	(0.0982)	(0.112)	(0.132)	(0.143)			(0.129)

Note: Table presents summary statistics of the weather variables constructed using ERA5 and Census education variables for each 5-year cohort bin in our sample (N=425 districts). The column headers represent the birth year of the middle age cohort within each bin.

Table A.2: Trends in climate, educational investment and enrollment

Panel A: Educational investment	7	8	9	10	11	12	13	14	15	16	Total
Degree-days > 24°C, in utero to age 5	1021.2	1012.6	1001.2	1005.6	1022.8	1001.5	1010.2	983.4	992.2	983.9	1002.7
	(450.3)	(436.6)	(447.9)	(437.7)	(444.6)	(438.9)	(450.5)	(441.8)	(447.8)	(450.3)	(444.3)
Degree-days > 24°C, age 6 to 10	1027.3	1044.4	1030.0	1034.8	1036.5	1014.8	1011.2	1013.6	1014.8	1006.6	1023.1
	(447.5)	(449.8)	(457.2)	(443.2)	(447.8)	(443.5)	(447.6)	(450.0)	(454.4)	(455.0)	(449.4)
Degree-days > 24°C, age 11 to 14					1048.0	1028.3	1043.5	1037.4	1037.9	1029.8	1036.6
					(455.5)	(444.3)	(462.2)	(458.6)	(460.0)	(462.9)	(456.8)
Degree-days > 24°C, age 15 to 16									1037.8	1031.0	1034.4
									(463.9)	(462.2)	(463.1)
Household expenditure on education (Rs.)	2625.7	2528.6	2733.7	2705.4	3118.3	3107.8	3674.0	4208.6	5182.2	6702.3	3572.4
	(5066.3)	(5223.2)	(5222.8)	(5315.7)	(5900.6)	(5517.7)	(6278.1)	(6633.8)	(7853.9)	(9870.7)	(6450.3)
Panel B: Enrollment											
Degree days > 24°C, in utero to age 5	988.0	977.8	972.0	968.9	971.3	972.7	966.5	972.3	977.2	971.8	973.9
	(466.1)	(463.3)	(463.4)	(462.4)	(462.8)	(464.5)	(461.2)	(462.9)	(464.5)	(462.4)	(463.3)
Degree days > 24°C, age 6 to 10	1010.6	1009.1	1001.1	1004.2	994.9	987.5	984.4	978.0	963.1	968.2	990.1
	(472.7)	(472.2)	(468.4)	(469.5)	(466.9)	(464.2)	(465.2)	(464.5)	(461.4)	(462.5)	(467.0)
Degree days > 24°C, age 11 to 14					1019.0	1010.6	1009.1	1001.1	1000.5	993.1	1005.6
					(484.4)	(472.7)	(472.2)	(468.4)	(468.5)	(468.0)	(472.4)
Degree days > 24°C, age 15 to 16									1019.0	1010.6	1014.8
									(484.4)	(472.7)	(478.5)
Share enrolled in school	0.713	0.767	0.817	0.789	0.829	0.772	0.774	0.721	0.638	0.578	0.740
	(0.185)	(0.184)	(0.153)	(0.162)	(0.133)	(0.160)	(0.148)	(0.160)	(0.177)	(0.179)	(0.181)

Note: Table presents summary statistics of the weather variables constructed using ERA5 and NSS educational investment variables in panel A, and Census enrollment variables in panel B, for each age group in our sample. Column headers represent age at the time of survey. Household expenditure on education in panel A is presented in real terms, deflated using the national CPI with 2010 as the base year.

Table A.3: Contemporaneous heat exposure & educational attainment in levels

	School Completion					
	Primary		Middle		Secondary	
	(1)	(2)	(3)	(4)	(5)	(6)
Degree days > 24°C	-0.013	-0.017	-0.012	-0.012	-0.003	-0.005
	(0.003)***	(0.004)***	(0.002)***	(0.002)***	(0.001)**	(0.002)***
	[0.005]**	[0.006]***	[0.003]***	[0.004]***	[0.003]	[0.003]
Precipitation z-score	-0.001	0.001	-0.004	0.002	0.000	-0.000
	(0.002)	(0.003)	(0.001)**	(0.002)	(0.001)	(0.001)
	[0.004]	[0.004]	[0.003]	[0.003]	[0.003]	[0.003]
Region-cohort trends	Y	N	Y	N	Y	N
Region-cohort FE	N	Y	N	Y	N	Y
Dep. variable mean (in levels)	0.580	0.580	0.405	0.405	0.250	0.250
Observations	3,825	3,825	3,400	3,400	2,975	2,975

Note: Table presents regression estimates of Equation 6, where $f(T_{cjr})$ captures contemporaneous exposure to heat during school years. More specifically, $f(T_{cjr})$ captures the annual average number of degree days above 24°C during the years in which individuals are age-appropriate for the corresponding level of schooling. The dependent variable is the share of individuals in a birth cohort who have completed at least primary school (columns 1 and 2), middle school (columns 3 and 4) and secondary school (columns 5 and 6). All columns include controls for precipitation as well as district and birth cohort fixed effects. We present standard errors clustered by district in parentheses, and Conley standard errors that allow for spatial correlation up to 500 km and arbitrary serial correlation in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table A.4: Heat exposure & cohort size

	Cohort Size		
	Primary	Middle	Secondary
	(1)	(2)	(3)
Degree days > 24°C, last ten years	0.014	0.027	0.033
	(0.008)	(0.009)***	(0.010)***
	[0.013]	[0.013]**	[0.013]***
Precipitation z-score, last ten years	0.007	0.017	0.013
	(0.009)	(0.009)*	(0.011)
	[0.011]	[0.010]*	[0.011]
Dep. variable mean (in levels)	0.444	0.307	0.203
Observations	5,124	5,124	5,124

Note: The sample consists of five-year birth cohorts born between 1952 and 1971, as observed across 3 Census waves: 1991, 2001 and 2011. The dependent variable is the natural logarithm of the share of individuals in a cohort who have completed at least primary school (column 1), middle school (column 2), and secondary school (column 3). The key independent variable captures the annual average number of degree days above 24°C in the preceding ten years. All columns include controls for precipitation as well as cohort, district, and region x census year fixed effects. We present standard errors clustered by district in parentheses, and Conley standard errors that allow for spatial correlation up to 500km and arbitrary serial correlation in brackets. * p < 0.10, ** p < 0.05, *** p < 0.01