RIDGE

RESEARCH INSTITUTE FOR DEVELOPMENT, GROWTH AND ECONOMICS

SINCE 2014

Climate Risk and Access to Credit: Evidence from Colombian Potato Growers

Camilo Bohorquez-Penuela, Margarita Gáfaro, Karelys Guzmán-Finol, Alex Perez

RIDGE Working Paper 027

This paper was presented at the 2025 RIDGE December Forum

Cooling the Heat on Learning: Temperature Anomalies, Academic Performance, and the Role of Social Protection

Gabriel Cruz
University of Maryland

Marcos Fabian
University of Maryland

September 8, 2025

Abstract

We study the impact of anomalous heat on learning outcomes in Mexico using a nationwide student-level panel that tracks over nine million students annually between 2007 and 2013, following each cohort consecutively for up to six years from grades 3–9. These longitudinal records, which include standardized test scores in math and Spanish and geocoded school locations, are merged with high-resolution daily weather data (0.05°) from 1981–2013 to construct municipality-level temperature and precipitation panels. Heat exposure is measured as deviations from local historical norms, both over the academic year and during the immediate test period. Our estimates show that a one–standard deviation increase in test-period heat anomalies reduces scores by about 0.03 standard deviations overall, with the penalty in primary school math (–0.06 SD) more than triple that in secondary school. Academic-year exposure produces similar magnitudes for both subjects, indicating that prolonged heat harms learning across the curriculum. Exploiting within-classroom variation in Mexico's flagship conditional cash transfer program, we find that beneficiaries are fully protected from the negative effects of test-period heat shocks. Finally, early-grade exposure has persistent effects, with learning losses still detectable three to four years later.

1 Introduction

Climate change's intensifying heat poses a growing threat to human-capital formation by disrupting education. Using student-level data, recent studies have linked rising temperatures to declines in learning and performance (Arceo-Gomez and López-Feldman, 2024; Li and Patel, 2021; Park et al., 2020). However, many questions remain about how ongoing climate change will affect human capital in the long run.

We use a comprehensive Mexican census panel that follows more than nine million students annually from 2007 to 2013. The dataset includes nationwide standardized test scores for students in grades 3–9 and geocoded school locations. We merge these data with daily temperature and precipitation observations (0.05° spatial resolution) from 1981–2013 to construct a municipality-level weather panel. Exposure to anomalous heat is quantified by comparing average temperatures over specific time windows to their historical norms for the same calendar dates. This event-specific deviation approach—widely used to study weather impacts on various outcomes (Arceo-Gomez and López-Feldman, 2024; Albert et al., 2021; Lin et al., 2020; Rocha and Soares, 2015)—is linked to our panel data and incorporates student fixed effects, allowing us to isolate the impact of extreme heat on student performance. By focusing on deviations from local historical averages, we account for regional differences in baseline climate (and any resulting adaptation), thereby capturing the true effect of heat anomalies. As Dell et al. (2014) note, what constitutes an unusual weather shock in one area may be normal elsewhere, where adaptation has already occurred.

With this rich data set, we want to answer the following two questions. First, what is the effect of anomalous heat on student performance, and does this effect vary with grade level? Second, to what extent can conditional cash transfers help protect learning outcomes from the impact of anomalous heat?

For the first question, we estimate our main specification on the full sample and then separately for primary- and secondary-school students.

For the second question, we exploit variation within classrooms in receipt of Mexico's largest conditional cash-transfer program. In other words, we observe side-by-side which students received the transfer and which did not. It is important to note that this program was not designed to mitigate climate-change impacts.

For our estimations, we use two complementary measures of heat deviation. First, test-period exposure, defined as the difference between the average temperature during the exam month and that month's historical average over 1983–2013—capturing the impact of near-exam heat on student performance (Vu, 2022; Li and Patel, 2021; Zivin et al., 2020). Second, academic-year exposure, defined as the difference between the average temperature over the current academic year up to the exam date (instructional days only, excluding weekends and official breaks) and the location's historical average for the same calendar days. This measure captures cumulative heat exposure across the school year and its broader effect on learning.

Additionally, leveraging the information of our data, we explore how early heat shocks can influence later academic performance by exploiting the variation in shocks that students experience across different years.

Our findings show that exposure to anomalous heat significantly reduces student performance, with effects concentrated among younger students and in mathematics. A one–standard deviation increase in test-period heat anomalies lowers scores by about 0.03 standard deviations overall, with the penalty in primary school math (–0.06 SD) more than triple that in secondary school. Using academic-year exposure, we find similar magnitudes for math and Spanish, suggesting that prolonged heat affects both subjects, while shorter-term exposure primarily harms younger students. We also find that Mexico's conditional cash transfer program substantially mitigated these losses, fully offsetting the penalty for beneficiaries during the test period. Finally, early-grade heat shocks have lasting effects, with impacts persisting up to three to four years later, underscoring the long-term consequences of climate-related learning disruptions.

The magnitude of our estimated effects is noteworthy: the impact of extreme heat is about 50% larger than the average effect of educational interventions involving more than 2,000 students (Kraft, 2023). In other words, hot weather can offset the gains from over half of typical large-scale educational interventions.

This paper contributes new evidence to the education–climate literature by providing several key findings. First, we show that heat shocks—measured using an approach that accounts for local adaptation—have a significant and negative impact on student learning. Using a rich dataset, we exploit student fixed effects to control for innate ability. We also show that these impacts occur throughout the academic year, not just in the month or days immediately preceding the exam. Second, we find that temperature shocks have substantially larger adverse effects on academic performance in primary school than in secondary school, suggesting that younger students are more vulnerable to heat-related disruptions. Third, leveraging student-level variation in program participation, we show that a conditional cash transfer program mitigated these negative impacts, effectively eliminating the adverse effect of heat shocks on learning. Finally, we explore whether these effects are transitory and find that they are not: exposure to higher temperatures can be cumulative, with early-grade heat shocks reducing students' test scores even three-four years later, indicating that early learning losses due to heat may persist over time.

The paper is structured as follows. Section 2 provides institutional background related to temperature variation and Conditional Cash Transfer program in Mexico. Section 3 describes the empirical approach. Section 4 presents the main results. Section 5 concludes.

2 Data and Background

2.1 Temperature Variation in Mexico

The main source for temperature data is the Climate Hazards Group InfraRed Precipitation with Stations (CHIRPS) and its companion dataset CHIRTS, which provide quasi-global daily estimates at a 0.05° spatial resolution from 1983 to 2016. CHIRPS combines satellite imagery with in-situ station data to generate bias-corrected precipitation fields, while CHIRTS applies similar methods to produce consistent gridded temperature estimates.[1]

We aggregate daily temperature grids to the municipality level by averaging across all raster cells falling within each municipal boundary. This produces a panel of daily

[1]The dataset was developed through a collaboration between the United States Geological Survey and the Climate Hazards Center to support drought monitoring and early warning systems. Details in `https://chc.ucsb.edu/data`.

municipality-level temperature values (°C) from 1983 to 2014.

To quantify exposure to anomalous heat, we construct standardized temperature deviation measures that compare average temperature over a given period to historical norms for the same calendar window. Similar measures are used in the literature to study the effects of weather anomalies on human outcomes (Arceo-Gomez and López-Feldman, 2024; Lin et al., 2020; Rocha and Soares, 2015). The key measure is a standardized temperature deviation for municipality i in year y, computed as:

$$\Delta T_{iy}^{(h)} = \frac{\overline{T}_{iy}^{(h)} - \mu_i^{(h)}}{\sigma_i^{(h)}} \tag{1}$$

where $\bar{T}_{iy}$ is the average daily temperature over a defined period of interest, and μ_i and σ_i are the historical mean and standard deviation for the same time window in municipality i, calculated over the 1983–2014 period.

We construct two temperature-deviation measures. Test-period exposure deviation is defined as the difference between the average temperature over the exam month and the preceding month[2] and the historical mean for the same two-month window over 1983–2014, standardized by the corresponding historical standard deviation. Academic-year exposure deviation is defined as the difference between the average temperature on instructional days from the start of the school year up to the exam date (excluding weekends and official breaks) and the location's historical average for those same calendar days, again standardized by the historical standard deviation. This second measure captures cumulative heat exposure across the school year and its broader effects on learning.

To construct the historical reference, we apply the same instructional calendar (excluding weekends and official breaks) to each year from 1983 to 2014. The distributions of the short- and long-term deviations are shown in Figure 1, panels (a) and (b), respectively.

[2] For academic years 2007–2008, 2008–2009, 2009–2010, and 2013–2014 we use March–April; for 2010–2011, 2011–2012, and 2012–2013 we use April–May.

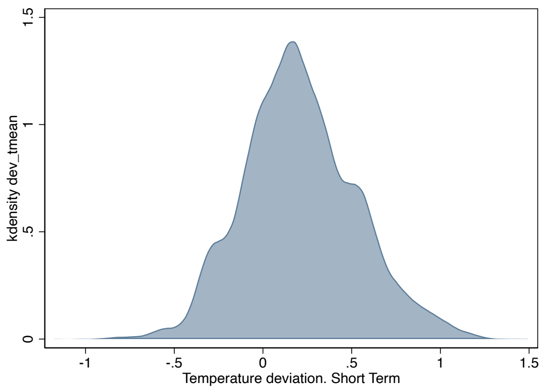

Panel a) Test-Period Exposure

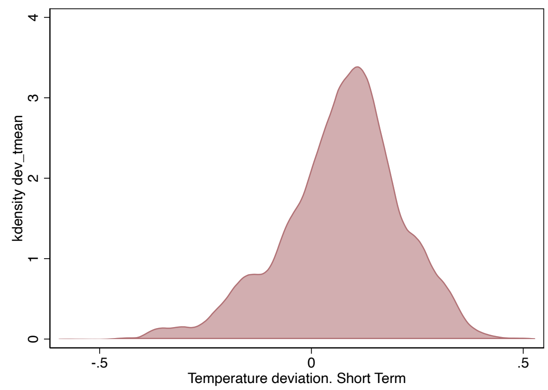

Panel b) Academic Year Exposure

Figure 1: Distribution of Standardized Temperature Deviations

2.2 ENLACE standardized test implementation

Our main data source is the *Evaluación Nacional del Logro Académico en Centros Escolares* (ENLACE), a nationwide standardized assessment implemented annually by the Secretaría de Educación Pública (SEP) between 2006 and 2014. ENLACE was designed to measure student achievement and provide diagnostic information for teachers, schools, and policymakers. Coverage was census for the targeted grades: all students in grades 3 to 6 of primary school and grades 1 to 3 of lower secondary school, with Spanish and mathematics assessed every year and a rotating third subject. We focus exclusively on Spanish and Mathematics, the only subjects tested consistently across the entire period.

Each subject test consisted of approximately 50 multiple-choice items, with raw scores reported on a scale from 200 to 800 points and classified into four achievement levels by SEP. Because the test forms and scaling procedures differed across grades and years, we standardize each subject's score within grade–year cells to have mean zero and standard deviation one. This transformation allows all results to be interpreted in standard deviation units and removes mechanical differences in score distributions due to test form variation.

The ENLACE test was typically administered once per academic year, generally in the spring (April or May), with exact dates varying slightly across years and between primary and secondary levels. In most years, primary school testing preceded secondary school testing by one to two weeks. The testing period was standardized nationally within each level to ensure comparability of conditions across states.[3]

Following De Hoyos et al. (2018), it is possible to construct longitudinal panels of students based on the unique identifiers available in the data. These identifiers allow tracking of academic histories for specific cohorts, although not all students in the panel begin in grade 3, and the length of follow-up varies across cohorts, ranging from four to six consecutive years. Table A.1 documents the cohorts that can be followed over time, showing entry grade, subsequent grade progression, and the academic years covered.

2.3 Progresa Beneficiaries

Progresa, later renamed Oportunidades and Prospera, was Mexico's flagship conditional cash transfer (CCT) program, implemented nationwide during the years covered by our analysis. Launched in 1997 in poor rural localities, the program aimed to break the intergenerational transmission of poverty by fostering investments in children's education, health, and nutrition . Coverage expanded steadily in the early 2000s to include all rural areas and, subsequently, urban communities (Parker and Todd, 2017). By the period of our study, it had reached national scale and was fully integrated into Mexico's federal social policy framework.

Importantly, family enrollment did not occur through voluntary application. Instead, eligibility was determined through a combination of locality-level poverty maps and detailed household socioeconomic censuses. Potentially eligible households were identified administratively and contacted directly by program personnel. Families who accepted the program's rules were incorporated into the beneficiary registry and subject to periodic re-certification of eligibility, but there was no open sign-up or competitive selection process among the poor (Parker and Todd, 2017).

A central feature of the program was its education-related conditionalities. To

[3] Specifically, the ENLACE tests were administered on the following dates: in 2007–2008, primary April 14–16 and secondary April 21–23; in 2008–2009, primary April 20–22 and secondary April 27–29; in 2009–2010, primary April 20–22 and secondary April 27–29; in 2010–2011, primary April 5–7 and secondary April 12–14; in 2011–2012, primary June 19–21 and secondary June 26–28; and in 2012–2013, primary May 14–16 and secondary May 21–23.

continue receiving the education grant, school-age children were required to maintain at least 85 percent attendance. Transfers were paid to the female head of household, and amounts increased with the child's grade level. In lower secondary school, grants for girls were set higher than for boys, reflecting concerns over higher dropout rates among adolescent girls. This design created progressively stronger incentives for families to keep children enrolled through the end of compulsory schooling, especially the transition from primary to lower secondary.[4]

Compliance with education conditions was monitored through attendance records sent by schools to program coordinators. Failure to comply could lead to temporary suspension of the relevant benefit, but families could regain benefits upon meeting requirements in subsequent monitoring periods. These mechanisms were a core part of the program's operational model throughout its expansion.

In the ENLACE microdata, students from grade 6 onward can be linked to Prospera through administrative identifiers, allowing us to distinguish beneficiaries from non-beneficiaries within the same schools and classrooms (De Hoyos et al., 2018). Following Behrman et al. (2025), we assign beneficiary status based on grade-6 participation and carry this indicator backward to grades 3–5 when defining baseline status. Because Prospera enrollment is highly persistent once a family qualifies, and the education grants are available for children enrolled in grades 3 through 9, grade-6 status serves as a reliable proxy for earlier participation.

3 Empirical Approach

We measure the impact of high temperatures on learning by exploiting nationwide standardized tests administered annually to all Mexican students in grades 3–9 in mathematics and Spanish from 2008 to 2013. These repeated assessments allow us to track individual students over time. Building on this, we exploit the quasi-random variation in temperature across years and at the municipality level to identify how changes in temperature affect learning outcomes at the individual level across different grade

[4]Alongside the education requirements, the program imposed health conditionalities, including regular medical check-ups and participation in health talks for all household members.

levels. Specifically, we implement an estimation strategy with student fixed effects, identifying the effect as follows:

$$test^a_{ismgy} = \beta\Delta T_{my} + \alpha precipitation_{my} + \theta CCT_{iy} + \eta_i + \gamma_y + \phi_g + \epsilon_{ismgy} \qquad (2)$$

Here, $test^a_{ismgy}$ is the standardized score in subject $a \in \{\text{Math, Spanish}\}$ for student i in school s, municipality m, grade g, tested in year y. ΔT_{my} is the standardized temperature deviation in municipality m and year y; P_{my} is precipitation in the same municipality–year; CCT_{iy} indicates whether student i received the conditional cash transfer in year y. η_i are student fixed effects, γ_y year fixed effects, and ϕ_g grade fixed effects. As an alternative, we replace ΔT_{my} with the mean temperature $\bar{T}_{my}$ to capture the effect of average (non-shock) heat exposure, allowing for adaptation.

We include student fixed effects η_i, which means identification comes from within-student comparisons of heat exposure and test score variation across multiple test years. This approach also controls for time-invariant unobservable characteristics of the student, such as innate ability or other individual traits. Finally, we include year and grade fixed effects, denoted by γ_y and ϕ_g, respectively. Standard errors are clustered at the municipality level, which is the level of variation in our treatment variable.

We also estimate a specification that replaces the single variable measuring the average impact of heat anomalies with a vector of anomaly levels, allowing for the possibility that the effect of heat exposure is non-linear, following Albert et al. (2021). In this approach, we divide the distribution of ΔT_{my} into 8 bins and compare each bin to the reference category—bin 3—which includes values around $\Delta T_{my} = 0$. Specifically, we estimate the following regression:

$$\begin{aligned} test^a_{ismgy} = {} & \beta_8 Bin^8_{my} + \beta_7 Bin^7_{my} + ... + \beta_2 Bin^2_{my} + \beta_1 Bin^1_{my} \\ & + precipitation_{my} + CCT_{iy} + \eta_i + \gamma_y + \phi_g + \epsilon_{ismgy} \end{aligned} \qquad (3)$$

In this specification, the coefficient β_8, corresponding to the highest-deviation bin, measures the effect of experiencing the largest positive deviation relative to the omitted bin (deviations ≈ 0). This binning approach allows us to examine nonlinear responses to heat deviations.

3.1 Conditional Cash Transfer

Exploiting student-level information on household receipt of the conditional cash transfer (CCT), we test whether the program mitigates the effect of heat shocks on academic performance. Specifically, we estimate the following regression:

$$\begin{aligned} test^{a}_{ismgy} =& \beta\Delta T_{my} + \theta CCT_{iy} + \lambda CCT_{iy} \times \Delta T_{my} + \\ & \alpha precipitation_{my} + \eta_i + \gamma_y + \phi_g + \epsilon_{ismgy} \end{aligned} \tag{4}$$

In this specification, we extend equation 2 by interacting the temperature deviation ΔT_{my} with the student-level CCT indicator CCT_{iy}. The coefficient β measures the effect of a one-standard-deviation increase in ΔT_{my} for non-beneficiaries ($CCT_{iy} = 0$). The coefficient of interest, λ, captures how the effect of temperature deviations differs for students in the CCT program relative to non-beneficiaries.

Although Progresa was not designed to mitigate climate shocks and is plausibly exogenous to temperature, we verify that temperature deviations (ΔT_{my}) do not systematically predict participation in Progresa (CCT_{iy}). Since CCT participation is known to affect educational outcomes, any systematic relationship between ΔT_{my} and CCT_{iy} would indicate that heat shocks influence the likelihood of receiving the transfer. This could raise concerns that the interaction term $CCT_{iy} \times \Delta T_{my}$ might capture not only how the program moderates the effect of temperature on test scores, but also compositional changes in who participates in the program when shocks occur. In this line, we estimate the following regression:

$$CCT_{iy} = \beta\Delta T_{my} + \alpha precipitation_{my} + \eta_i + \gamma_y + \epsilon_{imy} \tag{5}$$

Here, we follow the same notation used in previous estimations. Table A.2 reports the correlation between temperature deviations and Progresa participation for two exposure measures. The results show that academic-year temperature deviations are small, but statistically significantly correlated with participation in the conditional cash transfer program and are sizable in magnitude (0.011), whereas the correlation for test-period deviations is much smaller (0.002) and loses statistical significance.

The academic-year result suggests that in years with more heat shocks, the probability of a student's family receiving a conditional cash transfer increases. This is consistent with Arceo-Gómez et al. (2020), who show that extreme weather in Mexico can reduce rural household earnings and raise poverty rates by pushing vulnerable families below the official income poverty line. Since Progresa targets households under these income thresholds, such climate shocks may have increased the number of families eligible for the program.

These results suggest that families and the government may respond to weather shocks over longer horizons, while exposure during test periods may provide little time for federal agencies to adjust targeting and offer the program. This is worth keeping in mind when analyzing the results. However, the coefficient of interest, λ, still captures how the effect of temperature deviations differs between CCT beneficiaries and non-beneficiaries, conditional on the main effects of ΔT_{my} and CCT_{iy}.

3.2 Earlier Heat Shocks

We take advantage of the fact that we can track students over time, consecutively observing their standardized test scores, to analyze whether early heat shocks influence later academic performance. Specifically, we use one observation per student in grade 9 (or grade 8) and include measures of heat shocks from all previous years. We control for prior academic performance in each of those years and estimate the following regression:

$$Y_{isc}^{9} = \sum_{g=5}^{8} \beta_g \, \Delta T_{sc,g} + \sum_{g=5}^{8} \gamma_g \, P_{sc,g} + \sum_{g=5}^{8} \delta_g \, Y_i^g + \sum_{g=5}^{8} \alpha_g \, CCT_i^g + \lambda_s + \theta_c + \varepsilon_{isc} \qquad (6)$$

where Y_{isc}^{9} is the standardized test score of student i in school s, cohort c, measured in grade 9; $\Delta T_{sc,g}$ is the temperature deviation experienced in school s, cohort c, when the student was in grade g ($g = 5, \ldots, 8$); $P_{sc,g}$ is the precipitation level in school s, cohort c, in grade g; Y_i^g is the student's own standardized test score in grade g; CCT_i^g is an indicator for whether the student's family was a beneficiary of the conditional cash transfer program in grade g; λ_s are school fixed effects (based on the school attended in grade 9); θ_c are cohort fixed effects; and ε_{isc} is the error term. The coefficients β_g capture the effect of a one-unit increase in the temperature deviation in grade g on grade 9 test scores, conditional on precipitation, prior performance, and fixed effects.

Specifically, we focus on two sets of cohorts: cohorts 1 and 2 in Table A.1 to measure the effect of earlier shocks from grades 4 to 8, and cohorts 2 and 4 to measure the effect from grades 5 to 9.

In this specification, it is not possible to include student fixed effects to control for unobserved characteristics that do not vary over time. However, we control for previous test scores, which serve as a proxy for underlying ability. This approach provides an estimate of whether early shocks have a lasting effect on later test scores.

4 Results

4.1 Exposure Effects

Panel (A) of Table 1 presents the results of Equation 2 using temperature during the test-period exposure. The first two columns report estimates for all students in primary and secondary grades. The first coefficient indicates that a one-degree Celsius increase in temperature reduces test scores by approximately 0.01 standard deviations in both math and Spanish, and the effect is statistically significant.

However, using the average temperature has limitations: it may reflect students' adaptation to prevailing climate conditions, and a high temperature in one region may be normal in another, thus having different implications (Dell et al., 2014). For this reason, we focus on deviations from the mean temperature (ΔT), as described in the Data section. The results indicate that a one–standard deviation increase in hot anomalies reduces test scores by about 0.03 standard deviations in both math and Spanish. Columns 2 and 3 present the estimates for primary school students, while columns 5 and 6 present those for secondary school students. The effects in math are more than three times larger for primary students than for secondary students: specifically, math scores fall by 0.06 standard deviations—statistically significant—in primary school, compared to a 0.02 standard deviation decline in secondary school, which is not statistically significant. These results suggest that the effects of heat shocks are more pronounced for students in early grades, making them more vulnerable to heat-related impacts on learning.

Table 1: The Effect of Temperature Anomalies on Standardized Test Scores

	(1)	(2)	(3)	(4)	(5)	(6)
	Math	Spanish	Math	Spanish	Math	Spanish
			Primary School		Secondary School	
Panel A. Test-period exposure						
Mean Temperature	-.009	-.012	-.01	-.005	-.007	-.009
	(.004)	(.003)	(.003)	(.002)	(.003)	(.003)
Temperature Deviation (ΔT)	-.031	-.031	-.06	-.03	-.017	-.02
	(.012)	(.01)	(.015)	(.009)	(.01)	(.01)
Observations	34,884,905	34,884,905	15,108,913	15,108,913	17,003,783	17,003,783
Panel B. Academic Year Exposure						
Mean Temperature	-.004	-.009	-.004	-.007	-.01	-.011
	(.004)	(.003)	(.003)	(.003)	(.004)	(.004)
Temperature Deviation (ΔT)	-.04	-.056	-.045	-.043	-.046	-.042
	(.025)	(.02)	(.022)	(.017)	(.026)	(.024)
Observations	34,884,905	34,884,905	15,108,913	15,108,913	17,003,783	17,003,783

Note: Panel A reports the effect of *test-period exposure* on standardized test scores, where the test period is defined as the exam month and the preceding month. Panel B reports the effect of *academic-year exposure*, defined as the period from the start of the school year through the month of the exam. For each panel, three measures are reported: (i) **Mean Temperature**: the effect of exposure to the average temperature over the relevant period; and (ii) **Temperature Deviation** (ΔT): the effect of a temperature anomaly, as defined in Section 2.1. Columns 1–2 present results for the full sample, columns 3–4 for primary school students (grades 3–6), and columns 5–6 for secondary school students (grades 7–9). All regressions include student, year, and grade fixed effects, and control for municipality-year precipitation. Standard errors are clustered at the municipality level, which is the level of variation in the treatment variables.

With these results, we analyze the potential non-linearity of the temperature deviation. Specifically, we estimate Equation 3, dividing the distribution of ΔT into eight bins and comparing each bin to the reference category—bin 3—which includes values around $\Delta T = 0$. Panel (a) of Figure 2 presents the estimates for each bin using test-period exposure for all students in our sample. The results show a decrease of about 0.02 standard deviations in bins 4 to 6, and a larger decline—exceeding 0.04 standard deviations—in bins 7 and 9 relative to bin 3, which represents no deviation.

Panel B of Table 1 reports the estimates using temperature deviations for academic-

Figure 2: Non-linear Effects of Temperature Anomalies Exposure on Standardized Test Scores

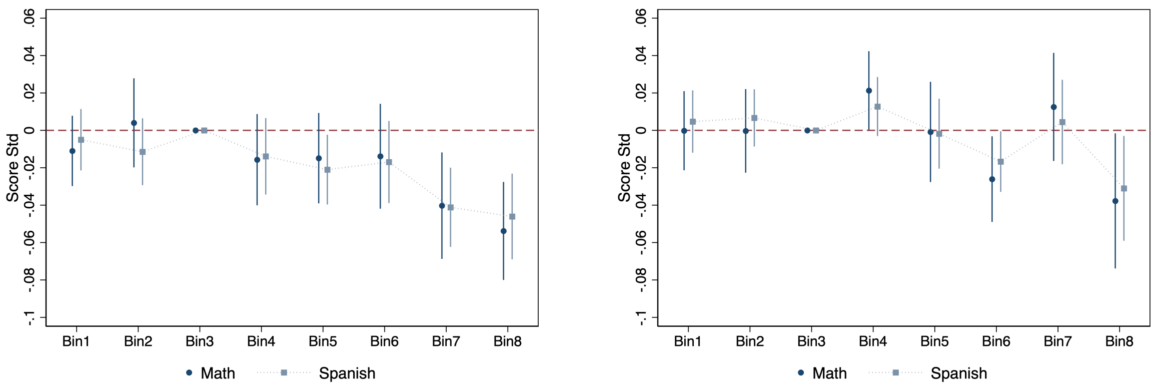

Panel a) Test-period exposure

Panel b) Academic Year Exposure

Notes: Panel A reports the effect of *test-period exposure* on standardized test scores, where the test period is defined as the exam month and the preceding month. Panel B reports the effect of *academic-year exposure*, defined as the period from the start of the school year through the month of the exam. Instead of a continuous ΔT_{my}, exposure distribution is divided into anomaly bins to allow for non-linear effects. Coefficients are relative to the reference category—bin 3—with ΔT_{my} values near zero. All regressions include student, year, and grade fixed effects, and control for municipality-year precipitation. Standard errors are clustered at the municipality level.

year exposure. On average, the effects are larger than those observed for test-period exposure. Specifically, a one–standard deviation increase in hot anomalies reduces math test scores by about 0.040 standard deviations and Spanish test scores by about 0.056 standard deviations. The magnitudes of the coefficients are similar for primary and secondary school students. However, the standard errors are smaller for primary school students, making the coefficients statistically significant, while in secondary school the larger standard errors reduce statistical significance. This pattern may suggest that older students are more heterogeneous and that some have better tools or strategies to cope with heat shocks. Moreover, the absence of a difference in effect size between younger and older students when using academic-year deviations indicates that prolonged heat exposure affects both groups similarly, leading to comparable learning losses. In contrast, when using deviations around the test period, younger students appear more affected, possibly because they have fewer coping mechanisms and are less able to mitigate the immediate impact of heat shocks during the exam period.

In the case of academic-year exposure, Panel (b) of Figure 2 shows a noisier pattern of non-linearity compared to the test-period exposure. While most point estimates are negative, their magnitude fluctuates across bins, and several have wide confidence intervals that include zero. Only in the highest anomaly bin (Bin 8) do we observe a more pronounced decline in test scores, suggesting that extreme heat over the course of the school year may have a stronger adverse effect on learning.

4.2 Conditional Cash Transfer Effect

Table 2 presents estimates of equation 4. Columns 1 and 3 report specifications with only the cash-transfer main effect (Mathematics and Spanish, respectively), while Columns 2 and 4 include the interaction $CCT_{iy} \times \Delta T_{my}$. In Panel A, we interpret the interaction as capturing how the temperature–score relationship differs for beneficiaries relative to non-beneficiaries, reflecting the program's potential mitigation effect.[5]

In Panel A of Table 2, the coefficient on ΔT_{my} is negative across all specifications, consistent with the baseline estimates in Table 1, indicating that higher-than-usual temperatures during the test period are associated with lower test scores. The estimated magnitudes range from about -0.03 to -0.05 standard deviations per one–standard deviation increase in hot anomalies, depending on the subject and specification. These effects are slightly more negative when the interaction term is included. The cash-transfer main effect is positive in all cases, suggesting that, at a given temperature, Progresa beneficiaries tend to score higher than non-beneficiaries.

The interaction term $CCT_{iy} \times \Delta T_{my}$ is positive and statistically significant in both subjects, indicating that the temperature penalty is smaller for beneficiaries. In Mathematics, the interaction is 0.068 and exceeds the magnitude of the temperature coefficient in column 2 (-0.052), so the estimated penalty from a one–standard deviation increase

[5] For the interaction $CCT_{iy} \times \Delta T_{my}$ to be interpreted as Progresa's mitigation effect, temperature variation in the test window must not itself influence program participation, so that differences in exposure are unrelated to take-up decisions during this period. In addition, conditional on fixed effects and controls, the interaction term must be uncorrelated with other unobserved determinants of scores. See Angrist and Pischke (2008) for discussion of interpreting interaction coefficients in this context.

Table 2: Cash Transfer Mediation Effects

	(1)	(2)	(3)	(4)
	Math		Spanish	
Panel A. Test-period exposure				
Temperature Deviation (ΔT)	-.031	-.052	-.031	-.048
	(.012)	(.013)	(.01)	(.011)
Cash Transfer	.035	.023	.02	.01
	(.005)	(.005)	(.004)	(.004)
Cash Transfer X ΔT		.068		.054
		(.009)		(.007)
Observations	34,884,905	34,884,905	34,884,905	34,884,905
Panel B. Academic Year Exposure				
Temperature Deviation (ΔT)	-.04	-.063	-.056	-.061
	(.025)	(.029)	(.02)	(.024)
Cash Transfer	.037	.033	.022	.021
	(.005)	(.006)	(.004)	(.005)
Cash Transfer X ΔT		.067		.014
		(.031)		(.023)
Observations	34,884,905	34,884,905	34,884,905	34,884,905

Note: Panel A reports the effect of *test-period exposure* on standardized test scores, where the test period is defined as the exam month and the preceding month. Panel B reports the effect of *academic-year exposure*, defined as the period from the start of the school year through the month of the exam. We estimate the effect of exposure to anomalous heat, as described in Section 2.1. Columns 2 and 4 include an interaction between extreme heat and receipt of a cash transfer. All regressions include student fixed effects, year and grade fixed effects, and controls for precipitation at the municipality-year level. Standard errors are clustered at the municipality level, which is the level of variation in the treatment variable.

in test-period temperature is more than fully offset for beneficiaries. In Spanish, the interaction is 0.054 against a temperature coefficient of -0.048 in column 4, implying a slightly smaller but still substantial attenuation.[6]

These results show that the conditional cash transfer program can fully mitigate the effect of hot temperature anomalies. One possible mechanism is that, to continue receiving the transfer, children must maintain regular school attendance. As a result,

[6]Panel B repeats the exercise using academic-year temperature deviations. With the identification caveats discussed above in mind, the patterns are qualitatively similar: the coefficients on ΔT_{my} are negative and of comparable magnitude (roughly -0.04 to -0.06), and the cash-transfer main effect remains positive. The interaction $CCT_{iy} \times \Delta T_{my}$ is again positive but smaller and less precisely estimated than in Panel A (statistically significant in Mathematics and small and not statistically significant in Spanish). Because academic-year deviations may influence program participation, we view these estimates as descriptive.

students from beneficiary families are less likely to miss classes during hot periods, while non-beneficiaries may skip school, disrupting the learning process. This could explain why, when we include the interaction between temperature deviation and CCT, the negative effect for non-beneficiaries increases relative to the average effect, thereby widening the gap in performance between beneficiaries and non-beneficiaries.

4.3 Early Shocks

Figure 3 presents the estimated effects of earlier heat shocks on later standardized test scores from equation 6. Panel (a) reports the effects of shocks from grades 5–8 on 9th-grade scores, while Panel (b) reports the effects of shocks from grades 4–7 on 8th-grade scores.

Figure 3: Effect of Earlier Heat Shocks on Later Test Scores

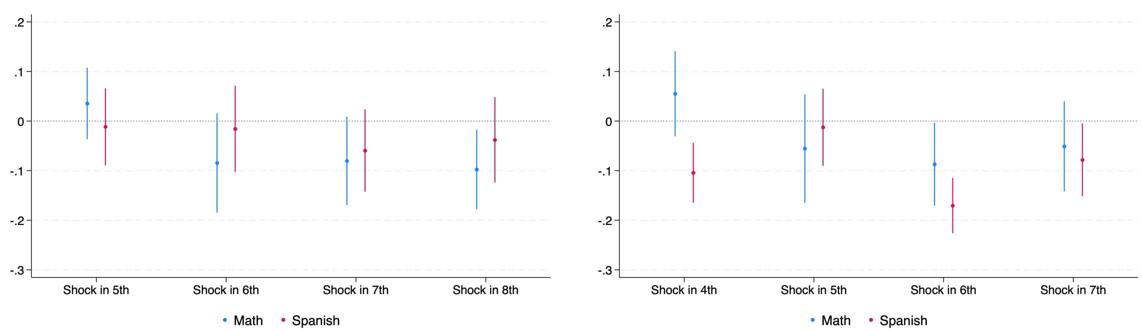

(a) Effect on 9th-Grade Test Scores

(b) Effect on 8th-Grade Test Scores

Notes: Coefficient estimates and 95% confidence intervals are from estimating equation 6. In panel (a), the dependent variable is the standardized test score of students in grade 9; in panel (b), it is the standardized test score of students in grade 8. The coefficients β_g capture the effect of a one-unit increase in the temperature deviation in grade g on grade 9 or grade 8 test scores (depending on the panel), conditional on precipitation, prior performance, participation in the conditional cash transfer program, and school and cohort fixed effects. The number of observations is 2,139,829 in panel (a) and 2,221,768 in panel (b).

For math, the pattern is clearer: in both panels, earlier heat shocks are generally associated with lower later scores, with several coefficients negative and, in some cases, statistically significant. The magnitude of the estimated effects ranges from about –0.05 to –0.10 standard deviations, with the largest declines often linked to shocks occurring closer to the test grade (e.g., grade 8 for 9th-grade outcomes, grade 7 for 8th-grade outcomes). For Spanish, the results are less consistent. Most coefficients are negative but smaller in magnitude, and the wider confidence intervals indicate lower statistical precision. In several cases, the estimates are close to zero, suggesting weaker evidence of a lasting impact of earlier heat shocks on Spanish performance compared to math.

These results suggest that earlier heat shocks have a more robust and precisely estimated negative effect on math achievement than on Spanish, with the largest effects occurring when shocks take place in the grades immediately preceding the test grade.

5 Conclusion

Our results show that anomalous heat meaningfully impairs student learning in Mexico, with the largest and most precisely estimated impacts on primary school mathematics performance. A one–standard deviation increase in test-period heat anomalies reduces scores by about 0.03 standard deviations overall, with early-grade effects persisting up to three to four years later. Importantly, Mexico's conditional cash transfer program fully offsets these losses for beneficiaries, underscoring the potential of social protection to buffer climate-related educational harms.

The magnitude of these impacts is substantial: the penalty from extreme heat is roughly 50% larger than the average effect of large-scale educational interventions involving more than 2,000 students (Kraft, 2023). In other words, a single hot testing period can erase over half the gains from typical education programs, highlighting the urgency of integrating climate resilience into education policy. Similarly, this effect size is comparable to those found in other studies using data similar to ours, such as the effect of being taught by a teacher with one year of teaching experience as opposed to none (Ladd and Sorensen, 2017), the effect of teacher gender on student academic achievement (Paredes, 2014; Dee, 2005), and the effect of being assigned to a teacher with a growth mindset (Claro et al., 2025).

References

Albert, C., Bustos, P., and Ponticelli, J. (2021). The effects of climate change on labor and capital reallocation. Technical report, National Bureau of Economic Research.

Angrist, J. D. and Pischke, J.-S. (2008). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, Princeton.

Arceo-Gomez, E. O. and López-Feldman, A. (2024). Extreme temperatures and school performance of the poor: Evidence from mexico. *Economics Letters*, 238:111700.

Arceo-Gómez, E. O., Hernández-Cortés, D., and López-Feldman, A. (2020). Droughts and rural households' wellbeing: evidence from mexico. *Climatic Change*, 162(3):1197–1212.

Behrman, J. R., Parker, S. W., Todd, P., and Zhang, W. (2025). Prospering through prospera: A dynamic model of cct impacts on educational attainment and achievement in mexico. *Quantitative Economics*, 16(1):133–183.

Claro, S., Paredes, V., Cruz, G., and Cabezas, V. (2025). Do students improve their academic achievement when assigned to a growth mindset teacher? *Educational Researcher*, page 0013189X251333775.

De Hoyos, R., Estrada, R., and Vargas, M. J. (2018). Predicting individual wellbeing through test scores: evidence from a national assessment in mexico. Policy research working paper, The World Bank.

Dee, T. S. (2005). A teacher like me: Does race, ethnicity, or gender matter? *American Economic Review*, 95(2):158–165.

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.

Kraft, M. A. (2023). The effect-size benchmark that matters most: Education interventions often fail. *Educational Researcher*, 52(3):183–187.

Ladd, H. F. and Sorensen, L. C. (2017). Returns to teacher experience: Student achievement and motivation in middle school. *Education Finance and Policy*, 12(2):241–279.

Li, X. and Patel, P. C. (2021). Weather and high-stakes exam performance: Evidence from student-level administrative data in brazil. *Economics letters*, 199:109698.

Lin, Y., Liu, F., and Xu, P. (2020). Effects of drought on infant mortality in china. *Health Economics*, 30(2):248–269.

Paredes, V. (2014). A teacher like me or a student like me? role model versus teacher bias effect. *Economics of Education Review*, 39:38–49.

Park, R. J., Goodman, J., Hurwitz, M., and Smith, J. (2020). Heat and learning. *American Economic Journal: Economic Policy*, 12(2):306–339.

Parker, S. W. and Todd, P. E. (2017). Conditional cash transfers: The case ofprogresa/oportunidades. *Journal of Economic Literature*, 55(3):866–915.

Rocha, R. and Soares, R. R. (2015). Water scarcity and birth outcomes in the brazilian semiarid. *Journal of Development Economics*, 112:72–91.

Vu, T. M. (2022). Effects of heat on mathematics test performance in vietnam. *Asian Economic Journal*, 36(1):72–94.

Zivin, J. G., 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.

Appendix

A Tables

Table A.1: Identifiable student panels by cohort

Academic Year	**3**	**4**	**5**	**6**	**7**	**8**	**9**
2007–2008	C1	C2	C3	C4			
2008–2009		C1	C2	C3	C4		
2009–2010			C1	C2	C3	C4	
2010–2011				C1	C2	C3	C4
2011–2012					C1	C2	C3
2012–2013						C1	C2

Notes: In Mexico, compulsory basic education consists of six years of primary school (grades 1–6) followed by three years of lower secondary school (grades 7–9, also referred to as 1st, 2nd, and 3rd year of secondary). The table summarizes cohorts of students that can be observed across multiple academic years taking the ENLACE test from 2006 to 2013. Labels "C1", "C2", etc. indicate specific cohorts that can be followed longitudinally, with each cohort representing the same group of students tracked as they progress through grades.

Table A.2: Effect of Temperature Anomalies on Progresa Status

	(1) Academic Year	(2) Test Period
Temperature Deviation (ΔT)	.011	.002
	(.003)	(.001)
Observations	34,759,174	34,759,805

Note: Column 1 defines the *test-period exposure* as the exam month and the preceding month. Column 2 *Academic-year exposure* refers to the period from the start of the school year through the exam month. All regressions control for precipitation and include individual fixed effects and test-year fixed effects. Standard errors are clustered at the municipality level.