Timing matters: Heterogenous effects of climatic events on school attendance in Nigeria

Liliana Andriano, University of Southampton
Julia Behrman, Northwestern University
Vineet Xie-Gupta, Northwestern University

Abstract

This paper contributes to the literature on the population-level consequences of climatic events by identifying two often-overlooked sources of temporal bias in studies linking climatic events to children's schooling: (1) exposure misalignment, where climate windows ignore differences in household vulnerability across the agricultural calendar; and (2) outcome misalignment, when schooling is measured too early or too late relative to when shocks take effect. Using Nigeria's Living Standards Measurement Study – Integrated Surveys on Agriculture (LSMS-ISA) linked to high-resolution SPEI, we show that substantive conclusions about the effect of climatic events on school attendance in Nigeria change depending on (1) whether shocks are measured in the preceding harvest season, preceding planting season, or period prior to the start of the school year; and (2) whether schooling was measured at the start or midpoint of the school year. Our findings have implications for both research methodology and educational policy in climate-vulnerable regions.

Keywords: climatic shocks, education, drought, Nigeria, selection

Introduction

Over the last few decades, extreme climate events have become increasingly severe and frequent. This trend together with the growing availability of geospatial data has fueled research on how shocks like droughts or extreme temperatures increases shape social and demographic outcomes ranging from schooling, to marriage, childbearing, migration, health and wellbeing (Agamile and Lawson 2021; Andriano and Behrman 2020; Baez, Caruso, and Niu 2020; Nubler et al. 2021). The relationship between these climate shocks and school attendance has garnered considerable attention in low- and middle-income countries where access to schooling has expanded substantially over the last 40 years (Agamile and Lawson 2021; Joshi 2019; Randell and Gray 2019; Shah and Steinberg 2017). In many of these countries where livelihoods remain heavily dependent on rainfed agriculture, children's schooling is especially sensitive to extreme climate shocks that disrupt household labor, income, and food security (Hasan et al. 2023; Seidu 2019; Siziya, Muula, and Rudatsikira 2007).

Empirical findings on the effects of climatic shocks vary widely across studies. Some studies report substantial declines in attendance and attainment following climatic shocks (Agamile and Lawson 2021; Dhongde and Spyrou 2023; Joshi 2019; Prenhaca and Melo 2024; Randell and Gray 2019; Shah and Steinberg 2017), consistent with increased opportunity costs of schooling (Zimmermann 2020). Others find no significant effects or even modest improvements (Calero, Bedi, and Sparrow 2009; Nordstrom and Cotton 2020; Randell and Gray 2019), suggesting that schooling may serve as an alternative to diminished labor market opportunities following climatic shocks. Similar heterogeneity characterizes research on related outcomes such as child health and nutrition, where studies report both adverse impacts (Dimitrova 2021; Hoddinott and Kinsey 2001; Kumar, Molitor, and Vollmer 2016) and null results (Dimitrova 2021; Randell, Gray, and Grace 2020). This inconsistency across studies could reflect contextual variation in mechanisms as the social and demographic effects of climatic shocks may depend on social, economic, and institutional conditions (Black et al. 2011; Blake 1968; Muttarak 2021).

We argue that beyond contextual factors, divergent findings may also reflect methodological choices that obscure otherwise consistent underlying patterns. Building on prior showing that the timing of climatic shocks within the agricultural cycle matters for outcomes such as schooling (Randell and Gray 2016), we introduce the broader concept of *temporal selection* which captures variation arising from the misalignment in the timing of both exposure and outcome measurement. Temporal selection operates in two distinct ways: through exposure misalignment and outcome misalignment. First, exposure misalignment occurs when climate measures fail to capture season-specific vulnerabilities. In rural agricultural economies, the consequences of prolonged rainfall deficits or sustained temperature anomalies may be mediated by the agricultural calendar, which shapes seasonal patterns of labor demand, income flows, and food availability. For example, a drought during planting may damage germination, increase replanting costs, and prompt coping strategies such as withdrawing children from school (Amondo, Nshakira-Rukundo, and Mirzabaev 2023; Haile, Gizaw, and Biazin 2022). At harvest, labor demands typically peak, and children may be pulled from school to help in the fields, though a shock at this stage can also reduce liquidity for school expenses (Randell and Gray 2019). These season-specific pathways suggest that the timing of climatic shocks within the agricultural cycle may produce distinct effects on schooling.

Second, outcome misalignment arises when the timing of data collection distorts observed impacts. Household surveys typically capture school attendance at one point in time during the academic year. Yet, the impacts of climatic shocks often unfold gradually as income losses, food insecurity, and increased demand for child labor accumulate over months (Hill, Skoufias, and Maher 2019). Coping responses may not be visible at the moment of the shock but emerge later, for example a planting-season rainfall deficit may constraint school attendance only at harvest, when households face lower liquidity. However, if school attendance is measured in the planting season, the impact will be obscured. Conversely, measuring school attendance long after short-term disruptions have resolved can also underestimate climate effects. Moreover, because school attendance is itself seasonal in many rural settings (Humphreys et al. 2015; Kumar and Saqib 2017), misalignment between survey timing and agricultural cycles can confound climate effects with regular seasonal fluctuations.

Both forms of temporal selection have important implications for causal inference and policy design. Analyses that ignore these dynamics may misestimate the magnitude and even the direction of impacts. Conversely, identifying which seasonal conditions most strongly affect school participation can inform the timing of policy interventions – for example, allocating resources during planting seasons when attendance may be most at risk.

Data, Context and Methods

In this paper, we demonstrate how these two forms of temporal selection shape our understating of the impacts of climatic shocks on school attendance by exploiting variation in the timing of the interview in the Nigerian Living Standards Measurement Study – Integrated Surveys on Agriculture (LSMS-ISA). Both Wave 1 (2010–11) and Wave 2 (2012–13) of the LSMS-ISA interview the same households twice within an academic year – once at the start and again at the mid-point – which allows us to capture responses to climatic conditions at multiple points in the school year. We link these individual-level school attendance data to high-resolution measures of climate variability from the Standardized Precipitation Evapotranspiration Index (SPEI) and define exposure windows that correspond to different phases of the agricultural cycle (e.g., planting, harvest, and pre-academic year). This design allows us to examine how climatic shocks at different points in the year influence school attendance at the start and midpoints of the academic year. We focus on Nigeria which is characterized by high climate variability, rapid educational expansion, and a large rural population dependent on rainfed agriculture making it a critical setting for understanding how climatic shocks shape schooling.

To examine the effect of SPEI3 on school attendance, we estimate a set of individual fixed-effects linear probability models with the following equation:

$$\text{Eq. (1) } y_{it}^{v} = \beta_0 + \beta_1 SPEI3_{it}^{s} + \beta_2 AGE_{it}^{v} + \mu_i + \gamma_t + \varepsilon_{it}$$

where y_{it}^{v} is a binary indicator of school attendance for child i in wave t, measured at visit v (either start or mid-point of the academic year). The treatment variable, $SPEI3_{it}^{s}$ captures local climatic conditions in season s (planting, harvest, or pre-academic year), measured as the 3-month standardized precipitation evapotranspiration index leading up to the seasonal cutoff (i.e. June 15, December 31, or August 31, respectively). We also include a time-varying control for the respondent's age measured at visit v (either start or mid-point of the academic year). A key feature of our design is the use of individual fixed effects, μ_i,which controls for time-invariant unobserved individual-level characteristics, such as parental education, family wealth, personal attitudes toward schooling, and unobserved household and community characteristics. Wave fixed effects, γ_t, further account for aggregate changes in school attendance across waves. Standard errors are clustered at the household level.

Results

Descriptive results

Table 1 presents weighted and unweighted descriptive statistics for the key variables in our analysis[1]. On average, 86% of respondents reported planning to attend school at the start of the academic year during visit 1, whereas 84% reported current enrollment at mid-year in visit 2 (paired sample t-test is significant with t-value of 5.4, p-value < 0.001). Specifically, 7.7% of respondents who reported planning to attend school during visit 1 are not enrolled at visit 2, and 29.9% of respondents who reported not planning to attend school during visit 1 are enrolled in school at visit 2. These discrepancies between the start and midpoint of the academic year suggests that school attendance is not stable throughout the school year among children in our sample.

Turning to SPEI conditions, Figure 1 shows the distribution of seasonal SPEI3 values for the pooled sample and by wave. Panel A displays conditions during the harvest season (October to December of the year preceding the wave), Panel B covers the planting season (March to June of the year the wave began), and Panel C reflects the pre-academic season (June to August of the year the wave began). Overall, the distributions of the SPEI3 are comparable in aggregate across wave 1 and wave 2, suggesting no major shifts in the frequency of climatic conditions between waves. The only notable difference is that wet conditions were proportionately much more common in wave 1 during the harvest season; in contrast, wet conditions were more common in wave 2 during the pre-academic year season.

Exposure misalignment: Does the effect of SPEI3 on attendance vary by the timing of treatment?

Figure 2 explores whether the effect of SPEI3 on school attendance differs depending on when during the agricultural cycle SPEI3 is measured. Figure 2 displays regression estimates and corresponding 95% confidence intervals for the effect of each of the 2 different measures of SPEI3 on school attendance (continuous SPEI3 and SPEI3 ≤ -1.0). The left column of Figure 2 shows results for the harvest season (Oct 1 to Dec 31), the middle column for the planting season (Mar 15 to Jun 15), and the right column for the pre-academic year season (Mar 1 to Aug 31). Each row corresponds to the survey visit: visit 1 at the start of the academic year (top) and visit 2 at the midpoint of the academic year (bottom).

We start by examining heterogeneity in results across agricultural seasons using the continuous measure of SPEI3 (shown in green). When SPEI3 is measured during the harvest season (left panel) or just before the academic year (right panel), the coefficients are close to zero and not statistically significant at $p<0.05$ in both the start (top row) and the midpoint of the academic year (bottom row). In contrast, when measured in the planting season (center panel), a one-unit increase in SPEI3 corresponds to a 1.6 percentage-point increase in the probability of start-of-the-year school attendance ($p = 0.173$) and a 4.0 percentage-point increase in the probability of mid-year school attendance ($p<0.001$), although the former result is not statistically significant at $p<0.05$.

Next, we explore the same question using a binary measure of moderate-to-extreme dry conditions (SPEI3 ≤ -1.0) rather than the continuous measure of SPEI3 (in red). The top row of Figure 2 shows results for school attendance at the start of the school year. When dry conditions occur in the harvest season (left panel), school attendance *increases* by 1.9 percentage points, although the effect is not statistically significant ($p = 0.165$). In contrast, moderate-to-extreme dry conditions in the planting season (center panel) lead to a 5.0 percentage-point decline in school attendance ($p=0.001$). When dry conditions occur just before the school year (right panel), the effect is null. Results are somewhat similar when school attendance is instead measured at the mid-point

[1] The weights used for the summary statistics are the wave-specific weights provided in the dataset, which adjust the wave sample to be representative at the population level.

of the academic year. When SPEI3 is measured in the harvest season (left panel) (–1.40 pp, p = 0.339) or immediately before the school year (right panel) (–1.66 pp, p = 0.086), coefficients are close to zero (though slightly negative) and not statistically significant at $p<0.05$. However, dry conditions during the planting season correspond to a 4.2 percentage-point decline in mid-year school attendance (p=0.004).

Taken together, results of these analyses suggest that there is notable heterogeneity in the effects of SPEI3 on school attendance depending on the agricultural season in which SPEI3 is constructed (i.e. harvest, planting or pre-academic year). These results suggest that exposure misalignment, or failure to account for when households may be most vulnerable during the year, can mask important differences in how dry conditions affects school attendance.

Outcome misalignment: Does the effect of SPEI3 on attendance vary by the timing of outcome?

Next, we focus on whether the effect of SPEI3 on school attendance differs depending on whether school attendance is measured at the start or middle of the school year. Figure 3 presents the same regression estimates as in Figure 2, reorganized to highlight variation by the timing of the survey visit during the academic year. The left column reports estimates for visit 1, conducted at the beginning of the school year (Sept–Oct), and the right column shows estimates from visit 2, conducted during the midpoint of the school year (Feb–Mar). Each row corresponds to the season when SPEI3 is measured: harvest (Panel A), planting (Panel B), and pre-academic (Panel C).

We start by using the continuous measure of SPEI3 (in green). When SPEI3 is measured in the harvest season (top row), there is minimal evidence that SPEI impacts school attendance either at the start or the mid-point of the academic year (0.2 pp, p = 0.778 at start; 1.0 pp, p = 0.187 at mid-point). Turning to the planting season (middle row), SPEI3 has little effect on attendance at the start of the academic year: the coefficient is positive but small and not statistically significant (1.6 pp, p=0.173). By mid-year, however, a one-unit increase in planting season SPEI3 leads to a 4.0-percentage-point higher probability of mid-year school attendance ($p<0.001$). Likewise, a one-unit decrease in SPEI3 leads to a 4.0-percentage point lower probability of mid-year school attendance. Turning to the period just before the school year (bottom row), SPEI3 has negative effect on attendance at the start of the academic year and a positive effect on attendance at the mid-point of the academic year, the coefficients are however small and not statistically significant at $p<0.05$ (–0.8 pp, p = 0.167 at start; 1.0 pp, p = 0.067 at mid-point).

Next, we look at models that use the binary measure of moderate-to-extreme dry conditions (SPEI3 $\leq$ –1.0) (shown in red). The top row (Panel A) shows that dry conditions in the harvest season (SPEI3 $\leq$ –1.0) have a positive impact (1.9 pp) at the start of the academic year, but a negative impact (-1.4 pp) mid-year, albeit neither of these results is statistically significant at $p<0.05$ (p=0.165 and p=0.339, respectively). When SPEI3 is measured during the planting season results are more consistent across both visits: dry planting conditions (SPEI3 $\leq$ –1.0) lead to a 5.0-percentage point ($p<0.001$) decrease in the probability of start of the year attendance and a 4.2 percentage point (p=0.004) decrease in the probability of mid-year attendance. Results are also consistent when SPEI3 is measured in the pre-academic year period: there is no evidence that moderate-to-extreme dry conditions in the period just before the school year impact school attendance at the start or mid-point of the school year.

Taken together, our results show the importance of considering the timing of the outcome variable. In some cases, climatic shocks have no effect at the start of the school year, but have impacts by mid-year, which suggests that impacts of shocks may be realized more slowly.

Figure 1: Distribution of seasonal SPEI3 by wave

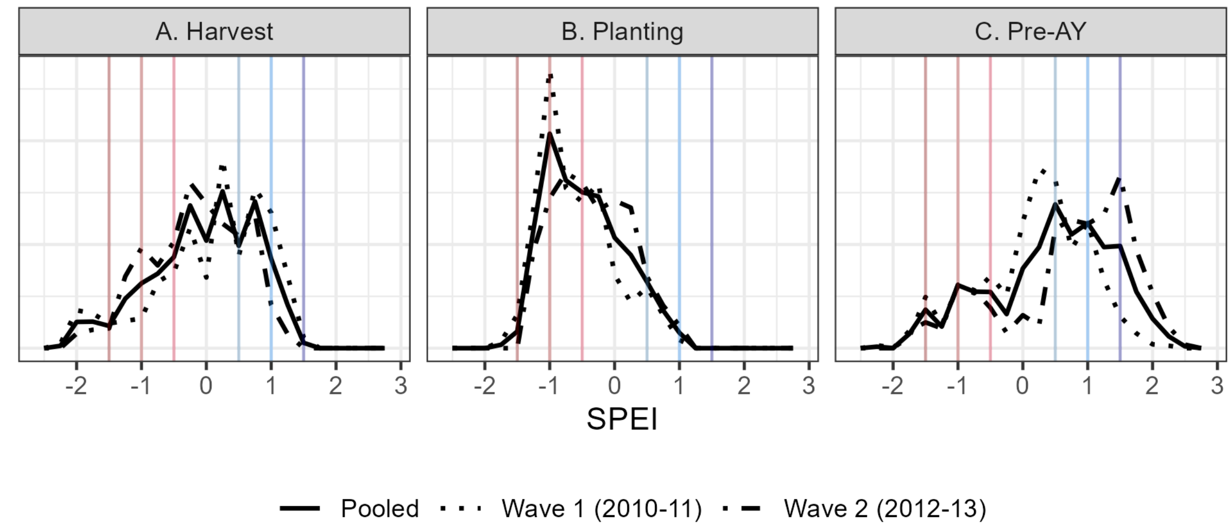

Notes: Graphs show the distribution of SPEI values by observations in the dataset, using the 3-month SPEI. Red and blue lines correspond to thresholds for dry and wet abnormal conditions used in our analyses. Harvest season is Oct 1 to Dec 31, planting season is Mar 15 to Jun 15, and Pre-AY (academic year) season is Jun 1 to Aug 31.

Figure 2: Estimates of the effect of different measures of SPEI3 on school attendance; columns reflect season when SPEI3 was measured and rows reflect timing during the academic year

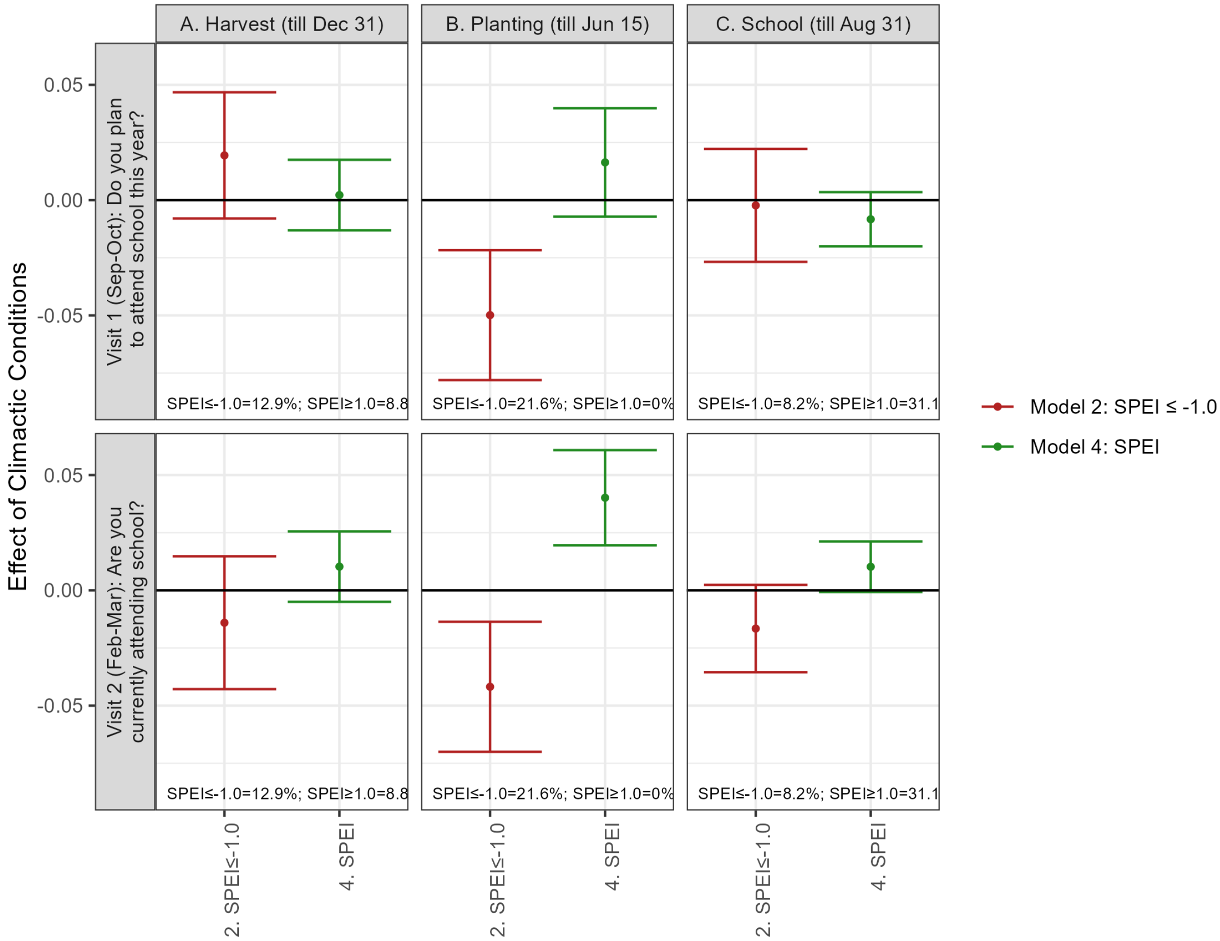

Figure 3: Estimates of the effect of different measures of SPEI3 on school attendance; columns reflect timing during the academic year and rows reflect season when SPEI3 was measured

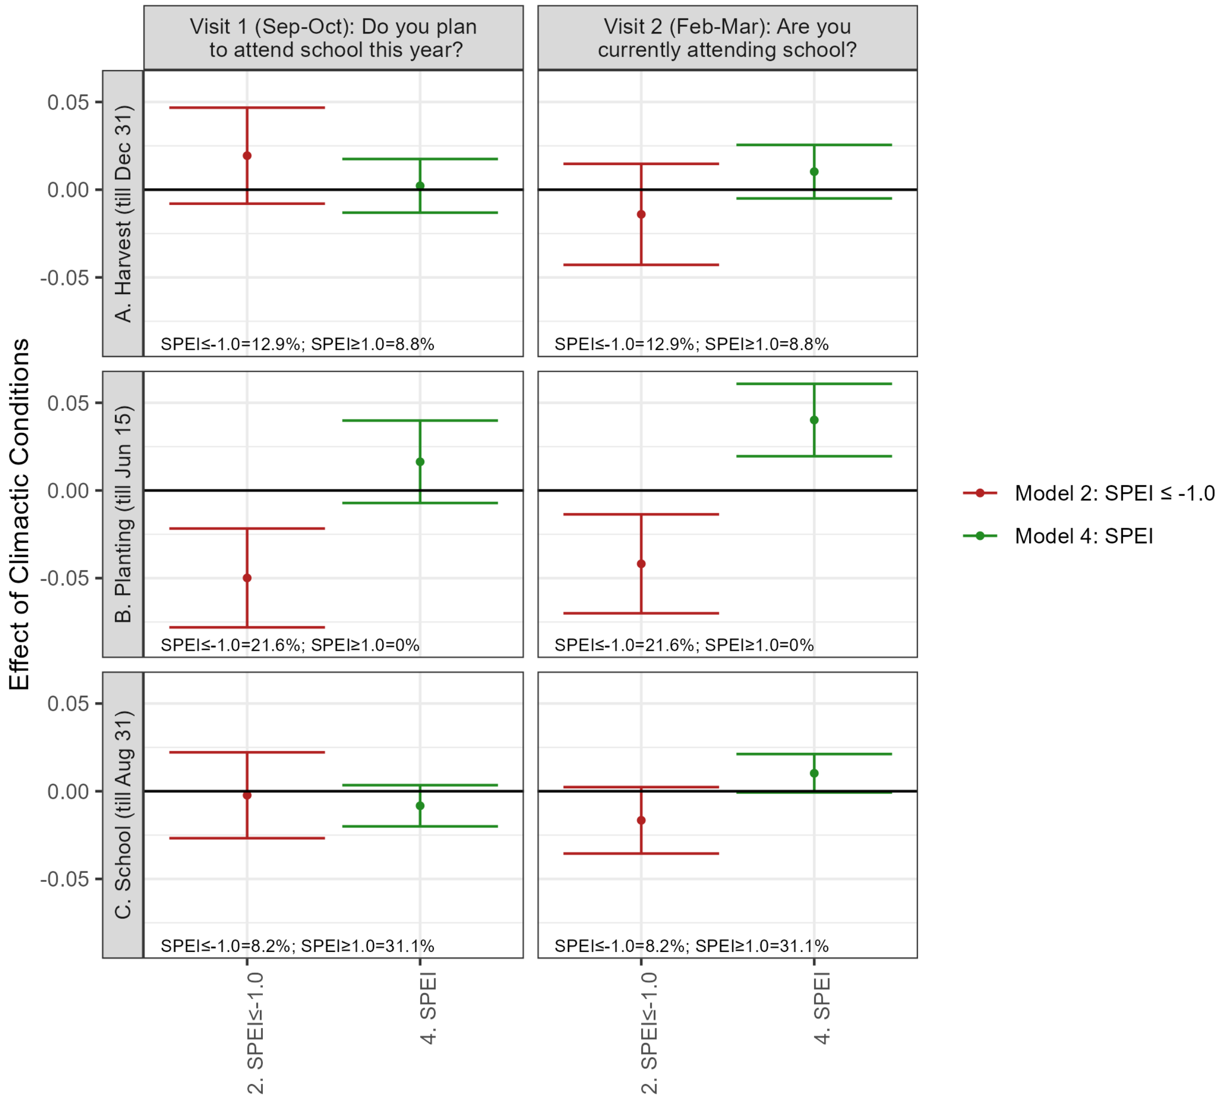

Table 1: Descriptive Statistics

	Wave 1 Visit 1		Wave 1 Visit 2		Wave 2 Visit 1		Wave 2 Visit 2	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
Visit 1: Do you plan to attend school this year?	0.87	0.34			0.85	0.36		
Visit 2: Are you currently attending school?			0.84	0.36			0.84	0.37
Sex (M=1; F=2)	1.45	0.50	1.46	0.50	1.45	0.50	1.45	0.50
Age (range:6-15)	9.27	2.44	9.27	2.44	10.85	2.51	11.16	2.47
SPEI03: Harvest (till Dec 31)	0.03	0.86	0.03	0.86	-0.30	0.75	-0.30	0.75
SPEI03: Planting (till Jun 15)	-0.57	0.58	-0.57	0.58	-0.31	0.58	-0.31	0.58
SPEI03: School (till Aug 31)	0.22	0.77	0.22	0.77	0.67	0.98	0.67	0.98
Observations	8380				8380			

Notes: Descriptive statistics of analysis sample in Visit 1 and Visit 2. Attendance questions ("Do you plan to attend school this year?" and "Are you currently attending school") were asked in only one visit each. SPEI3 Graphs show the distribution of SPEI values by observations in the dataset, using the 3-month SPEI. Harvest season is Oct 1 to Dec 31, planting season is Mar 15 to Jun 15, and pre-AY (academic year) season is Jun 1 to Aug 31.