Learning-by-Doing at Home: The Childhood Origins of Girls' Domestic Work

Ameek Singh*
London School of Economics

(Job Market Paper)
(Latest Version)

October 28, 2026

Abstract

The output of household work is rarely observed, particularly for individual members, making gendered comparative advantage difficult to measure. I develop a model-based test using household consumption per minute of domestic work: lower opportunity cost alone cannot raise this measure, whereas productive capacity can. I exploit rainfall-induced variation in household-production demand in rural India to study how such capacity forms. A one-standard-deviation rainfall increase at ages 4–7 raises girls' domestic work at ages 12–17 by 5.3 minutes daily. The estimate for boys is 1.0 minute and insignificant; infancy rainfall has no persistent effect for either gender. Effects concentrate among eldest daughters. Households whose eldest daughter experienced favourable early-childhood rainfall obtain 4–10 percent more consumption per minute, depending on current rainfall. Despite lower study time, I find little evidence of adverse effects on enrolment, grade progression, or measured learning. The results are consistent with early domestic-work engagement building task-specific productive capacity.

JEL Codes: D13, J24, J16, J22, J13, O15, O12.

* I am deeply grateful to my supervisors, Tim Besley and Rocco Macchiavello, for their guidance and encouragement throughout this project. It would not exist without Gabriela Deschamps, whose persistence as an early collaborator kept it alive. Ronny Condor provided excellent research assistance. I also thank seminar participants and faculty at the Department of Economics at LSE for many helpful comments. I gratefully acknowledge financial support from the Hayek Fund, the Hub for Equal Representation in the Economy at LSE, DAISI at J-PAL, and the LSE Department of Economics studentship and Job Market Fund. All remaining errors are my own.

1 Introduction

Economics has long treated the gender division of household work as an allocation shaped by wages, preferences, norms, and comparative advantage in home production. Quantitative accounts of adult time allocation represent the observed division of work using gendered productivity, preference, or disutility wedges, and comparative advantage in home and service production (Ngai and Petrongolo, 2017). These objects are inferred from adult time use or calibrated to reproduce adult allocations, rather than traced to their origins (Gottlieb et al., 2024). Yet time-use data record minutes spent on domestic tasks, not the quantity or quality of the output they produce. They cannot establish whether women are more productive at domestic work than men, or why one woman completes a task more efficiently than another. There is consequently little direct evidence for the domestic comparative advantage that these accounts represent, or for its source.

That gap matters because different sources of domestic specialisation have different implications. A capacity acquired through experience can reduce the time required for a given task, while making its holder especially valuable when household resources expand and must be converted into consumption at home. Habit, constrained choices, and norms can produce similar adult allocations without producing the same response to labour-saving technology, income growth, or market opportunities. Existing work shows that technology, marketisation, and sectoral change can alter gendered time allocation (Dinkelman and Ngai, 2022; Ngai and Petrongolo, 2017), but adult allocations alone cannot reveal whether a domestic-work advantage is productive skill, habit, or a consequence of constrained choices.

Making progress requires an external event that changes actual engagement in domestic work, rather than simply being correlated with a girl's initial capacity or her family's gender norms. Rural agricultural households provide such a setting. Where final consumption relies partly on labour-intensive household production, a favourable agricultural income shock can expand the household's demand for consumption and, with it, the domestic labour needed to produce that consumption at home. Whether it does so, and which child is recruited, are empirical questions.

Evidence on Learning-by-Doing. We find evidence of learning-by-doing in domestic work: early engagement builds domestic-work-specific skill. In rural India, temporary rainfall plenty during ages four to seven leaves girls with a durable domestic-work role in adolescence. The paper's central claim is therefore not simply that a favourable childhood shock increases girls' work, but that learning-by-doing in household production can form part of the domestic comparative advantage observed later in life.

We use variation in local rainfall during childhood as a source of shocks to agricultural income and household production. A one-standard-deviation increase in rainfall when a girl is aged four to seven raises her domestic work at ages 12–17 by 5.3 minutes per day, about 8 percent of the gender gap in the estimation sample; there is no comparable response for boys

or for girls exposed outside this early-childhood window. The estimated effect increases with the girl's age at exposure, from infancy through early childhood, and peaks when rainfall occurs at ages four to seven. It is also most pronounced among girls who later do substantial domestic work.

A persistent association with a household shock does not itself establish learning-by-doing: every child in the household shares the resulting change in resources and domestic-production demand. The short-run evidence shows that rainfall plenty raises demand for household domestic work—work performed principally by women and girls—alongside agricultural output and consumption. The persistent response is likewise driven by rainfall plenty rather than drought. We therefore ask whether it is concentrated on the girl whose position made her most likely to be recruited into domestic work. It is concentrated among eldest daughters rather than younger sisters. An eldest daughter has no older girl available to assist adults and is more likely to face the additional needs associated with younger siblings; a younger sister has an older, typically more physically capable girl already available to the household. This pattern is therefore consistent with the child who was recruited when domestic labour demand rose having accumulated experience through doing the work.

Household Model and a Test for Productive Capacity. We then develop a two-period household-allocation model to turn the unobserved-output problem into an empirical test. The model allows childhood engagement to have two different later effects: it can build domestic-work-specific productive capacity, or lower the opportunity cost of assigning the daughter's time to domestic work. It yields two main implications. First, even a complete set of later time-use effects cannot distinguish these channels: under plausible parameter values, either can generate the same allocation response. Second, a pure reduction in the opportunity cost cannot raise consumption per minute of domestic work, or make that ratio respond more positively to a new non-labour-income shock; productive capacity can.

The eldest-daughter evidence makes this test feasible. Because the childhood rainfall effect is concentrated on the daughter most likely to have been recruited into domestic work, we use her rainfall history as a source of variation in the household's domestic-production capacity and examine the associated adjustment of household consumption and work allocation. This does not attribute the household's output to one girl. It uses a child-specific exposure to study whether the household that relied on her during childhood subsequently produces more consumption from its domestic work, and reallocates tasks differently among its members.

We find robust evidence for this test. Households with an eldest daughter exposed to early-childhood plenty obtain about 4 percent more consumption per minute of domestic work at average current rainfall, and their advantage grows to roughly 10 percent when current rainfall is one standard deviation higher. This model-derived household-level test is the paper's central empirical contribution: it provides evidence of productive capacity where individual domestic output is not observed. The same households show a broader reorganisation away from direct domestic tasks toward upstream domestic production

and paid work by other members, consistent with an expanding household-production environment.

Robustness and Other Effects. Two further challenges must be addressed before interpreting this pattern as learning-by-doing. The first is causal: the rainfall history could proxy for another district–cohort change, a selected adolescent sample, or a particular way of constructing rainfall or diary time. We show that the long-run relationship survives alternative rainfall constructions, diary coding rules, survey weights, sample definitions, selection tests, spatial-inference procedures, and a directionally consistent check in the conventional labour-force survey. We also assess alternative explanations for a valid persistent relationship, including later fertility and child-care demand, health, and persistent household- and place-level consequences of historical rainfall. In particular, early-childhood rainfall plenty is not associated with lower enrolment, slower grade progression, or weaker measured learning achievement.

Finally, we ask whether the acquired domestic-work role is part of a broader change in adolescent gender roles. Domestic work is only one domain in which girls' choices may be restricted. We therefore combine the TUS measure of discretionary time outside the home with the NFHS measures of health-care access, permitted mobility, financial resources, and observed activities. The childhood history associated with more domestic work is not accompanied by less discretionary movement, lower reported health-care access, more restricted mobility, or lower observed access to resources. Together, the education, health, and gender-role evidence does not support a broader deepening of gendered disadvantage in girls' adolescent lives beyond domestic work. It instead locates the durable effect in a specific domestic-work role.

The paper makes two contributions. First, it advances understanding of the early-childhood origins of gendered domestic roles. It provides micro evidence that engagement in labour-intensive household production can build domestic-work-specific skill, a channel through which a comparative advantage often treated as an adult primitive may emerge. The resulting durable domestic-work role is not accompanied by a commensurate broadening of observed disadvantage in educational outcomes, mobility, permissions, or access to resources. Second, it develops a way to use time-use data to study otherwise unobserved productive capacity: a household-level consumption–time test that identifies a productive component even when individual domestic output and output quality are unobserved. Section 2 situates these contributions in the relevant literatures.

The remainder of the paper proceeds as follows. Section 3 describes the time-use, rainfall, and auxiliary data and sets out the empirical strategy. Section 4 presents the main time-use results. Section 5 examines identification threats, alternative specifications, measurement, inference, and external validation. Section 6 develops the model. Section 7 tests the household-production and learning-by-doing mechanisms. Section 8 examines whether the domestic-work effect is accompanied by broader changes in adolescent gender roles.

Section 9 concludes.

2 Related Literature

This paper asks whether girls acquire skills in domestic work through childhood experience. It builds on the household-production framework, in which households combine goods and time to produce consumption and services (Becker, 1965; Gronau, 1977). Models of the gender division of work use comparative advantage and preference differences, alongside sectoral demand, to explain how adults divide their time between home and market work (Ngai and Petrongolo, 2017; Gottlieb et al., 2024). We build on this work by asking whether childhood experience can create differences in domestic skills and thereby help explain later specialisation. Answering this question requires evidence of skill, but time-use data record time spent on domestic work, not the output produced by each person. A girl may do more housework because she is better at it, or because she gives up less by doing it. Our model provides a way to distinguish these explanations using household consumption and domestic-work time, even when her individual output is unobserved.

The demand for domestic work also depends on household technology and the availability of purchased substitutes. Labour-saving durables can reduce the time needed at home (Greenwood et al., 2005); infrastructure and the shift towards purchased goods and services can also change women's time use (Dinkelman and Ngai, 2022). We study an expansion of household production in rural India, where consumption still requires substantial work at home. Favourable rainfall increases agricultural output, consumption, and total household domestic work, providing variation in the demand for domestic labour. This demand response helps explain why acquiring a domestic skill need not reduce the girl's work: she may need less time for a given task, yet be assigned more work when household demand rises.

The rainfall literature shows why better agricultural conditions need not release children's time from work. Shah and Steinberg (2017) show that favourable rainfall in India can raise the opportunity cost of schooling when children are old enough to work. Bai and Wang (2020a) distinguish the income and price effects of economic shocks on child labour and schooling. We add another reason why a favourable shock can increase work: more household consumption can require more domestic labour. We then ask whether this additional work leaves a lasting effect. The adolescent response is concentrated among eldest daughters exposed at ages four to seven, with no comparable infancy response. Together, these patterns help assess explanations based only on changes in household resources, although evidence that the girl acquired a skill requires the additional model-based test.

Research on child labour also asks what children give up when they work. Households may rely on child labour when resources are scarce and children's work can substitute for adult work (Basu and Van, 1998). This can reduce investment in the child's future (Baland

and Robinson, 2000). Domestic work needs separate attention because it can increase while a child remains enrolled and outside measured employment. Indian time-use evidence shows different gender and age patterns in chores, market work, and schooling (Kis-Katos, 2007). Our results illustrate why both time use and educational outcomes matter: study time falls, but there is limited evidence of lower enrolment, slower grade progression, or weaker measured learning. Our model also allows children to acquire a specific skill through work, alongside any costs from the activities that work displaces.

Which child does the work is therefore central to our argument. Productive support targeted to women in Nicaragua increased older girls' specialisation in domestic and non-agricultural work (Del Carpio and Macours, 2010). Mother–daughter substitution under PROGRESA reduced first teenage daughters' childcare and increased their school participation (Dubois and Rubio-Codina, 2012). These studies show how changes in family production and adult work can redistribute children's tasks. Building on this evidence, we ask whether children who take on additional domestic work learn to do it more productively. Both the short-run increase and the persistent response in our data are concentrated among eldest daughters. Although birth order is not a randomly assigned treatment, this concentration links the persistent response to the daughter most likely to have taken on additional work when household demand rose.

Skills acquired through work must also be distinguished from other forms of human capital. Pitt et al. (2012) study how gender-specific physical capacity, formed through health and nutrition, affects work, schooling, and occupations in rural Bangladesh. Our explanation concerns skills acquired through practice in domestic tasks, but rainfall could also affect later work through health and nutrition. We therefore examine health and schooling as alternative explanations for the persistent domestic-work response before interpreting it as evidence of learning through practice.

Research on gender inequality gives us another reason to examine these alternative explanations. Economic returns, institutions, and norms shape parents' investments and girls' opportunities (Jayachandran, 2015). These forces could sustain domestic roles even without a difference in skill. Other studies show that particular constraints can be changed: reducing travel costs raises girls' school participation (Muralidharan and Prakash, 2017), exposure to female leaders changes aspirations, education, and household work (Beaman et al., 2012), and school-based interventions can change gender attitudes (Dhar et al., 2022). These findings also show why domestic work, schooling, and agency should be examined separately. We examine whether the rainfall history associated with more domestic work also predicts losses in these other areas, finding limited evidence of broader losses in educational achievement, health, mobility, or observed agency. This supports an interpretation centred on domestic work, although it does not establish that additional work is costless or directly measure beliefs or bargaining power.

Finally, research on skill formation helps explain how an early experience could matter years

later. Skills acquired early can make later investment more productive (Cunha and Heckman, 2007). Weather studies also find that the ages at which exposure matters differ across later health and cognitive outcomes (Maccini and Yang, 2009; Webb, 2024). Our age profile is consistent with learning beginning when children can take on domestic tasks, but cannot distinguish whether practice is especially productive at those ages or early experience is reinforced by later practice.

Learning-by-doing means that experience improves productive performance (Arrow, 1962; Foster and Rosenzweig, 1995). To test this in domestic work, persistence alone is insufficient: a girl could continue doing more work because it has become habitual or because her alternatives have become less valuable. Because individual output is unobserved, our model derives a test using household consumption and domestic-work time. Under its assumptions, a lower opportunity cost alone cannot raise household consumption per minute of domestic work or make this ratio respond more positively to a new resource shock. We find this pattern in the data, providing evidence consistent with acquired domestic skill while allowing habit, norms, and continued task assignment to matter as well.

3 Data and Empirical Strategy

Our empirical design compares time allocation among same-age adolescents observed five years apart in two Time Use Survey waves. To implement this design, we combine each adolescent's TUS diary with her district-level rainfall history from the Global Precipitation Climatology Centre's (GPCC) Monitoring product, which provides a monthly gridded record from 1982 to the present. We use the geoBoundaries district shapefile to aggregate the GPCC rainfall grid into a district-level series at the same district boundaries used for the time-use data. Four auxiliary surveys extend the analysis to margins the diary does not cover: NFHS for adolescent health and agency, ASER for learning outcomes, PLFS for non-diary measures of work and studying, and ICRISAT for district-level agricultural output. We next describe these data sources and how we process them for the analysis.

3.1 India Time-Use Surveys: TUS 2019 and 2024

Our primary outcomes come from India's Time-Use Surveys (TUS), conducted by the National Statistical Office in 2019 and 2024. These are nationally representative surveys recording detailed 24-hour diaries for every household member aged 6 and above, a rare feature among time-use surveys, which typically cover only adults or a subset of household members. This makes the TUS unusually well suited to studying gender gaps in children's time use and their evolution over childhood.[1]

Within each round, the survey's two-stage design balances day-of-week and seasonal coverage

[1] Most time-use surveys collect data for children above the age of 12 or 15, and often only for a random sample of household members. India's TUS is particularly useful for studying the evolution of gender gaps in time use since early childhood.

by construction: every rural or urban block selected as a first-stage cluster (the survey's primary sampling unit) canvasses two sample households on each day of the week, and the survey period is partitioned into four equal-allocation quarterly sub-rounds spanning the year.[2] Within each sample household, every eligible member is asked to complete a single 24-hour diary; the diary records both the day of the week it covers and whether the respondent considered it a typical day or one departing from her usual routine.[3]

The diary is collected by face-to-face recall of the previous 24 hours, from 04:00 on the day before the interview to 04:00 on the day of the interview, divided into forty-eight 30-minute slots. Each slot may contain up to three activities; for slots with more than one activity, the survey additionally records whether they were performed simultaneously and which one the respondent reports as the "major" activity.

3.1.1 Outcome Construction

Our main individual-level outcomes are the sum of slot-level minutes spent on each activity category over the diary day. We compute these from the slot-level activity log in two steps: for each slot reporting multiple activities, the slot's 30 minutes are split equally across the activities reported;[4] the resulting per-slot minutes are then aggregated to person-day totals using the ICATUS 2016 classification, which contains 165 unique activity codes harmonised into a smaller set of aggregates. The equal-split rule means the minimum observed time on any activity is 10 minutes.

This construction preserves the 1,440-minute diary budget. Multi-activity reporting is common enough to matter for measurement: roughly one quarter of girls' domestic-work minutes occur in slots with more than one activity (Table A1, appendix), so Section 5 reports robustness to alternative slot-coding rules.

At the highest level, we organise an individual's daily time allocation into four exhaustive, non-overlapping categories that sum to 24 hours:

- **Domestic work** (h_d): time spent on household chores, including meal preparation, cleaning and laundry, care for household members, household maintenance, and provisioning. We also classify tasks such as growing vegetables or tending livestock as *Domestic Work* if they are aimed at the household's own consumption; similar tasks aimed at income generation are classified as *Paid Work*. The economic aim of the activity is marked in the original ICATUS activity codes under which they are recorded.
- **Studying** (h_s): time spent in formal and informal education, as well as studying at home. We also include time spent commuting to school or other educational institutions

[2] 14 households per primary sampling unit, with two canvassed on each day of the week; four 3-month sub-rounds with equal allocation of primary sampling units.

[3] Exception types include festival days, public holidays, family events or ceremonies, sick days, and other respondent-flagged categories. Within the headline rural 12–17 sample across the 26 main states, the resulting day-type composition is 75.6% Monday–Saturday normal-day diaries, 10.8% Sunday diaries, and 13.6% Monday–Saturday non-normal diaries.

[4] As a robustness check, we recompute all aggregates using only the declared "major" activity per slot; results are very similar.

in total time on studying.

- **Paid work** (h_m): time devoted to activities aimed at income generation — wage and salaried work, self-employment, casual labour, and unpaid help in household enterprises (including own farm or tending to livestock).
- **Leisure and self-care** (h_l): residual time including sleep, personal care, socialising, media consumption, religious and community participation, and entertainment.

For domestic work, our primary outcome, we further disaggregate into: (a) food preparation, (b) cleaning and laundry, (c) domestic production, (d) care work, and (e) shopping and budgeting; we retain these task-level components alongside the aggregate. Where the focal margin is discretionary leisure rather than residual non-work time, we also use a narrower *Entertainment and Hobbies* measure that excludes sleep, eating, and personal care from the leisure aggregate.

All main minutes-per-day outcomes are winsorised at the state × gender 99th percentile. Figure 1, Panel B, maps the district-level distribution of girls' domestic work. Several of the paper's analyses look at how time use and the rainfall-shock effect differ by day type: school day versus non-school day, and normal versus non-routine. For analyses that condition on the time of day rather than on day-level totals (the time-of-day profiles in Section 4 and the slot-level decompositions in the appendix), we retain the diary at its native 30-minute resolution. Each individual's diary day yields 48 slots, and each slot's time on a given activity is computed by the same proportional rule.

3.1.2 Parent and Sibling Identification

Several analyses require reconstructing parent–child links and sibling age ranks. The TUS roster records each member's relationship to the household head, age, gender, and marital status, rather than direct family links. We use these fields to identify eldest daughters and construct the family-identified samples used in the mechanism analysis; the headline regression does not require them.

Here, an eldest daughter is the oldest co-resident daughter aged 17 or younger in a reconstructed sibling set. If an older sister has left the household, an observed eldest daughter is in fact a younger sister. Absent departure that is selectively related to rainfall, this misclassification attenuates the eldest-daughter result; Section 7.2 examines it directly.

We rank siblings by observed age within each reconstructed set of children aged 17 or younger. Analyses that require an unambiguous focal family impose the additional single-family restriction; the main eldest-daughter and sibling-order results use the broader family-identified sample. The linkage procedure, its retention rates, and robustness checks addressing older-sister attrition and family composition are described in Section 7.2, where these measures are first used.

3.1.3 *Estimation Sample and Coverage*

The headline estimation sample is rural children aged 12–17 who completed a TUS diary, pooled across the 2019 and 2024 rounds, in 26 main states (listed in Table 7). The sample contains 59,466 children, split roughly evenly across rounds.[5]

The restrictions are chosen to match the estimand rather than to maximise sample size. The rural restriction follows from the mechanism we study: early-childhood agricultural plenty is expected to operate through rural household production. The urban sub-sample is therefore not a robustness alternative for the headline; it is used separately as a mechanism check in Section 7. The lower age bound of 12 places the outcome in adolescence, after the gender gap in domestic work has become meaningfully large (Figure 1, Panel A). We report age-by-age heterogeneity for the main outcomes in Section 4 to show that the headline is not driven by a single age. The upper bound of 17 keeps the sample below the ages at which marriage and out-migration to the husband's household would begin to select the observed sample of girls.

The state restriction is driven by sample-size requirements. Several of India's smallest union territories and northeastern states contribute too few rural girls aged 12–17 for stable estimates: each fewer than 200 across the two pooled rounds, and several fewer than 50 (Table 7). We exclude ten such units; they account for 1.9% of the rural children 12–17 in the pooled sample (1.96% of the rural girls).

3.1.4 *Descriptive Patterns in Children's Time Use*

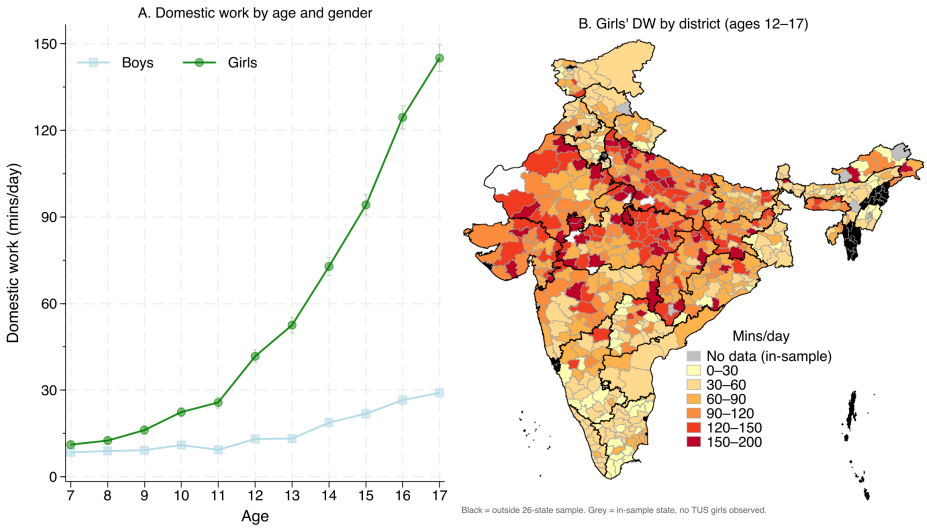
The motivating fact for the paper is a pronounced gender gap in adolescent domestic work that emerges over childhood. Figure 1, Panel A, traces the within-cell mean of daily domestic work by age and gender for rural children with completed time diaries. The two series begin to diverge from age 9, by age 14 girls record nearly four times as much daily domestic work as boys, and by age 17 roughly five times as much. Panel B maps the district-level distribution of girls' domestic work over the rural 12–17 sample: mean values are highest in the agricultural belt of the north and centre and lowest in parts of the south and the eastern coast.

Table 8 broadens the comparison to other categories of time use across the full pooled rural sample, separately for ages 7–11 and 12–17. Among ages 12–17, 55% of girls cross the participation margin in domestic work against 20% of boys. Within domestic work, food preparation accounts for the bulk of the gap: girls average 40.9 minutes per day against 1.9 for boys, with 38% of girls vs. 3% of boys recording any time on food preparation. Cleaning shows the same direction at smaller magnitude. Care work and domestic production show statistically significant but much smaller differences than food preparation, and budgeting and shopping is the one sub-category that is statistically ungendered.

Studying time, by contrast, is comparable across genders at both age bands (361 vs. 362

[5] Analyses that require linking a child to her parents (Section 3.1.2), notably the eldest-daughter regressions and the mother time-use regressions in Section 7, are run on the family-identified sub-sample described in Section 3.1.2. The sub-sample retains 95.2% of children at the broader 0–17 cut-off and 95.8% at the 12–17 cut-off used here.

FIGURE 1: *The Adolescent Gender Gap in Domestic Work*



Note: **Sample.** Panel A uses rural children aged 7–17 with completed time diaries. Panel B uses rural girls aged 12–17 with completed time diaries in the main 26 states. **What is plotted.** Panel A plots age-by-gender mean domestic-work minutes per day; vertical bars are 95% confidence intervals for the within-cell means. Panel B maps pooled district means of girls' domestic-work minutes per day. District polygons outside the main 26-state estimation sample are shaded black. District polygons within the 26 states but with no eligible diary observation are shaded light grey.

minutes at ages 7–11; 351 vs. 348 at ages 12–17). The girls' domestic-work gap does not come out of studying in the descriptive aggregate. It comes out of paid work, which moves in the opposite direction (boys aged 12–17 average 31.9 minutes per day against 12.8 for girls), and out of leisure and self-care, where girls aged 12–17 average about 46 minutes per day less than boys. The leisure gap lands almost entirely in entertainment and active hobbies (boys 206.6 vs. girls 159.5 minutes per day, a 47-minute differential): sleep, personal care, and socialising are essentially balanced across genders (sleep 569 vs. 568; personal care 168.8 vs. 167.5; socialising 87.9 vs. 93.8, with girls slightly higher). The descriptive picture is therefore not one of girls sleeping less or socialising less, but of girls having less discretionary entertainment and hobby time.

3.2 Rainfall Data and Shock Construction

To capture rainfall variation during childhood, we use the Global Precipitation Climatology Centre (GPCC) Monitoring Product. This dataset provides monthly gridded precipitation estimates at $1^\circ \times 1^\circ$ resolution from 1982 to the present, based on quality-controlled rain-gauge observations from a global network of meteorological stations. Section 5 assesses whether the results depend on this product or on the rule used to assign its grid cells to districts.

Of the two gridded GPCC products covering India, we use the *Monitoring* series ($1^\circ \times 1^\circ$,

updated monthly through 2024) as our primary rainfall measure. The higher-resolution *GPCC Full* series ($0.25° \times 0.25°$, full historical station archive) has not been updated for years after 2022, and so does not cover the recent end of the early-childhood window for the younger cohorts of the 2024 TUS round. The Monitoring product covers both life-stage windows for every cohort in both rounds; we report sensitivity to the GPCC Full product, for cohorts where both are available, in the appendix.

The GPCC product is on a regular spatial grid, while our outcome data and the rest of the paper's analysis sit at the district level. Section 3.2.1 below describes the polygons we use and the mapping from grid points to districts; here we take that mapping as given and construct the rainfall variables at the district-year level.

Annual rainfall shocks. For each district, we sum monthly rainfall to obtain annual totals R_{dy} and define the annual rainfall shock as the standardised deviation from the long-term mean:

$$Z_{dy} = \frac{R_{dy} - \mu_d}{\sigma_d}, \tag{1}$$

where μ_d and σ_d are the mean and standard deviation of annual rainfall for district d over 1982–2024.

Two life-stage rainfall windows. We focus on two early-life windows that distinguish infancy from ages at which children begin to take on more household activity. The *Early Childhood Shock* (EC) is the average district rainfall z-score over ages 4–7. The *Infancy Shock* (IC) is the average over ages -1 to 2. Formally, for an individual of age a observed in survey year t:

$$\text{ECS}_{dat} = \frac{1}{4} \sum_{y \in \mathcal{Y}_{at}^{EC}} Z_{dy}, \qquad \text{ICS}_{dat} = \frac{1}{4} \sum_{y \in \mathcal{Y}_{at}^{IC}} Z_{dy}, \tag{2}$$

where $\mathcal{Y}_{at}^{EC}$ and $\mathcal{Y}_{at}^{IC}$ are the sets of four calendar years corresponding to the early-childhood and infancy windows. Both averages are re-standardised to unit standard deviation in the estimation sample.[6]

Three properties of the rainfall series matter for the empirical strategy. First, annual district rainfall has little serial correlation: within-district partial autocorrelations of annual rainfall z-scores have absolute value below 0.10 through lag 5 (Table 9, appendix). Second, the two life-stage averages are only modestly correlated in the estimation sample ($\text{Corr}(\text{IC}, \text{EC}) = -0.08$), allowing both to enter the same regression. Third, Figures 4 and 5 describe the realised differences used by the two TUS rounds. The district maps show that the Early Childhood and Infancy differences have distinct spatial patterns. The density figure shows that the 2019–2024 differences for both windows fall within the long-run distribution of district-level five-year rainfall differences, rather than representing an unusually one-sided episode. Section 5 uses these facts when assessing the baseline comparison.

[6] An additional window covering ages 7–10 (the early-school years) can be constructed analogously but is excluded from the headline specification (see Section 3.4.2).

Concurrent rainfall. Several mechanism and discussion analyses use a separate survey-period rainfall measure rather than the historical life-stage shocks. For each diary observation, concurrent rainfall is the mean of the four standardised quarterly rainfall shocks in the 4 quarters (12 months) ending in the survey quarter. The one-year lag is defined analogously over the preceding four quarters. These rolling twelve-month shocks are quarter-aligned and re-standardised within the relevant estimation sample. They describe contemporaneous household responses to rainfall and interactions between current rainfall and childhood exposure; they are distinct from the annual life-stage averages in equation 2.

3.2.1 District Polygons

Three datasets supply district boundaries used in the paper: the geoBoundaries-IND-ADM2 administrative shapefile (Runfola et al., 2020), the TUS 2019 sampling frame, and the TUS 2024 sampling frame. The three do not fully agree, because some districts have been split, merged, or renamed between rounds, so we construct a single set of polygons onto which all three are mapped. We construct this polygon set once, before undertaking the analysis, and use it throughout: to harmonise the two TUS rounds, aggregate historical rainfall, and place every auxiliary dataset on the same geographic units.

The polygons sit at a slightly coarser level than the district list of any single dataset, chosen so that every administrative district in any of the three datasets fits entirely inside one polygon. TUS reports data at the district level and no finer, so the polygons can only aggregate districts upward, never split them. A merge is required whenever a 2024 child district draws territory from more than one 2019 parent — as in Andhra Pradesh — and the affected parents and children are then combined into a single polygon. The same logic applies whenever the geoBoundaries shapefile reshuffles territory across pre-existing districts.

Most districts do not need to be coarsened. The three lists contain 734 polygons (geoBoundaries-IND-ADM2), 683 districts (TUS 2019), and 744 districts (TUS 2024); after the necessary merges, we are left with 662 polygons. 647 of these polygons cover exactly one 2019 administrative district, and only 15 combine two to four neighbouring districts. These 15 merges are concentrated in five states: Andhra Pradesh, Gujarat, Uttar Pradesh, Arunachal Pradesh, and Mizoram.[7]

We then map the remaining datasets onto these polygons. The GPCC rainfall product reports rainfall on a regular grid of $1^\circ \times 1^\circ$ cells (about 110 km on a side at India's latitude). This is coarser than the typical Indian district, so many polygons contain no grid cell whose centre falls inside their boundary. To assign a rainfall reading to every polygon, we widen each polygon by a 30 km buffer and compute its rainfall as a weighted average of every grid cell

[7] The 15 merges are: (Uttar Pradesh) Rae Bareli + Sultanpur; (Arunachal Pradesh) Upper Subansiri + Lower Subansiri, and West Siang + East Siang; (Mizoram) Aizawl + Champhai; (Gujarat) Sabar Kantha + Arvalli, Ahmedabad + Bhavnagar + Botad, Surendranagar + Rajkot + Morbi, Jamnagar + DevBhumi-Dwarka, Junagadh + Gir Somnath, Kheda + Panch Mahals + Mahisagar, and Vadodara + Chhota Udepur; (Andhra Pradesh) Srikakulam + Vizianagaram + Visakhapatnam + East Godavari, West Godavari + Krishna, Guntur + Prakasam, and Sri Potti Sriramulu Nellore + Y.S.R. (Cuddapah) + Chittoor.

whose centre lies inside the widened polygon, with cells in successively wider buffer rings down-weighted as they sit further from the polygon's true boundary.[8] Without the buffer, 423 of the 662 polygons — nearly two-thirds — have no grid cell whose centre falls strictly inside their boundary. The 30 km buffer covers all but 45 of these polygons; the radial alternative described below picks up the remaining 45 by extending each polygon's catchment further outward. Figure A1 illustrates the construction for Ludhiana district.

As sensitivity checks, we report two alternative constructions: one that uses only cells strictly inside the polygon (defined for the 239 polygons that contain at least one interior cell), and one that averages cells within a fixed radius of the polygon's centre, where the radius is 50 km plus the radius of a circle with the same area as the polygon. Where all three are defined, the pairwise correlations of the resulting annual rainfall z-scores are 0.999 (buffered vs. interior-only), 0.985 (buffered vs. radial), and 0.983 (interior-only vs. radial); the appendix reports the headline rainfall coefficients under each alternative construction.

The auxiliary surveys (PLFS, NFHS, ASER, and ICRISAT) are placed onto the polygons using the same logic. For each external district, we trace its territory onto the polygons. When the territory falls cleanly inside a single polygon, the district is assigned to it. When the territory straddles two or more polygons, those polygons cannot be reconstructed cleanly for that period and are excluded from the panel, together with any further external district whose territory touches an excluded polygon.

All district-level choropleth maps in the paper use the polygons defined above. Polygons outside the 26-state estimation sample (Section 3.1.3) are shaded black, and polygons inside it with missing outcomes (purely urban polygons and a handful of remote districts) are shaded grey.

3.3 Auxiliary Datasets

We use five auxiliary datasets: NFHS for adolescent health and indicators of mobility and autonomy, ASER for foundational reading and arithmetic competence, PLFS for non-diary measures of domestic work and studying, HCES for food consumption, and ICRISAT for district-level agricultural output. In each, we restrict to the same 26 main states used in the TUS headline sample.

National Family Health Survey (NFHS). We use NFHS for adolescent health and observed indicators of mobility, permissions, and autonomy. Section 8 defines these outcomes when it examines whether the domestic-work pattern is accompanied by broader differences in adolescent gender roles.

NFHS-4 (2015–16) and NFHS-5 (2019–21) are the Indian rounds of the Demographic and Health Survey programme. The two rounds together interview around 1.4 million women aged 15–49 across all states and union territories. A household questionnaire collects the

[8] Cells inside the polygon (or within 1 km of its boundary) receive weight 1; cells in the 1–2 km ring 1/2; in 2–4 km 1/4; in 4–15 km 1/15; and in 15–30 km 1/30.

household roster, basic demographics, and asset ownership; a women's questionnaire is administered to every eligible woman in the sampled households and collects detailed information on health, fertility, marriage, decision-making, and attitudes. A subset of the agency items used here sits within a longer empowerment block administered to a state-rotating sub-sample of households rather than the universal women's sample, so estimation sample sizes differ across the agency outcomes.

Annual Status of Education Report (ASER). We use ASER to measure foundational reading and arithmetic competence in an independent learning-assessment instrument, complementing the schooling-time measures recovered from TUS.

ASER is an annual rural household-based survey running since 2005, conducted by the ASER Centre (previously by Pratham). The sample covers all rural districts of India, and every child aged 5–16 in a sampled household is administered short reading and mathematics tests in the home regardless of school-enrolment status. The paper uses the 2018 and 2022 rounds, which are closest to the cohorts observed in TUS 2019 and 2024.

The reading test rates each child on a five-level ladder of competence: below the alphabet, recognition of letters, recognition of common words, ability to read a Grade-I-level paragraph, and ability to read a Grade-II-level short story; the highest level achieved is coded as the child's reading score. The mathematics test follows the same five-level structure across number recognition (1–9), two-digit recognition (11–99), two-digit subtraction with borrowing (Grade I), and three-digit-by-one-digit division (Grade II).

From the two ladders we construct three nested learning indicators, matching the column labels in the ASER results table: *Reading & math (Grade I)*, equal to 1 when the child passes the Grade-I task in both reading and mathematics; *Both Grade I & at least one Grade II*, equal to 1 when she additionally passes the Grade-II task in at least one of the two domains; and *Reading & math (Grade II)*, equal to 1 when she passes the Grade-II task in both. The three indicators trace the share of children reaching progressively higher rungs of the joint reading-and-mathematics ladder. We additionally use *enrolment* (currently enrolled in any school) and an *on-track-for-age* indicator equal to 1 when the child is enrolled in a class within six years of her age. Children not tested on a component are coded as failing that component.

Periodic Labour Force Survey (PLFS). PLFS is India's national labour-force survey, in the field since 2017–18. We use the 2019, 2021, and 2023 rounds to stay close to the TUS cohorts while avoiding the COVID-affected 2020 and 2022 rounds. We construct four weekly outcomes: domestic-work days per week, an indicator that domestic duties are the respondent's principal weekly status, studying days per week, and an indicator that studying is the respondent's principal weekly status. These are non-diary measures from an independent survey, so they validate signs and gender patterns rather than magnitudes in TUS minutes per day.

Other surveys. The Household Consumption Expenditure Survey (HCES) records monthly physical quantities of food items consumed, alongside expenditure. We use its rural 2009–11 rounds to distinguish food ingredients and staples from prepared food and meals eaten away from home when examining the contemporaneous rainfall response in Section 7. Annual district-level agricultural data for 1990–2017 come from the International Crops Research Institute for the Semi-Arid Tropics (ICRISAT). They report production, yield, and area harvested by major crop, and are used to measure the year-to-year agricultural-output response to rainfall in Section 7.

The next section sets out the main identification strategy for the paper.

3.4 Empirical Strategy

3.4.1 What We Identify

The empirical object is the persistent effect of district rainfall during ages 4–7 on time use in adolescence, and how that effect differs between girls and boys. Formally, these are the coefficients β_E (the effect on girls, the base group) and β_E^g (the boys-versus-girls differential) on the Early Childhood Shock in the specification below. The outcome is observed nearly a decade after the rainfall window, so the design must separate realised early-life rainfall from slow-moving district features, cohort composition, and survey-wave changes in adolescent time use.

The resulting coefficient is a reduced-form effect of rainfall during ages 4–7 on adolescent time allocation. The estimating equation does not determine why this relationship persists. Section 5 examines whether it reflects the intended rainfall comparison; Sections 6 and 7 then assess whether the evidence is consistent with domestic-work-specific skill formation or other long-run explanations.

Rainfall is a standard source of year-to-year variation in agricultural conditions in development economics, used to study outcomes such as wages, schooling, health, and labour supply (Maccini and Yang, 2009; Almond and Currie, 2011). Our implementation compares children observed at the same age in the same district but born in different years, and therefore exposed to different rainfall during the same life stage. We apply this approach to early-childhood exposure and adolescent time use.

3.4.2 Baseline Specification

The two TUS rounds allow us to compare same-age cohorts observed five years apart within a district. A 13-year-old observed in 2019 and a 13-year-old observed in 2024 were born five years apart and faced different rainfall histories at the same early-life ages. The baseline specification therefore holds district, age, and survey quarter fixed while using the rainfall history attached to each cohort's childhood.

Its identifying assumption is that the realised difference in childhood rainfall between these

cohorts is not systematically related to other district-level changes that independently affect adolescent time use. The fixed effects absorb time-invariant district features, district-specific age and seasonal profiles, and national gender-specific changes in adolescent time use. The remaining variation is the difference in realised childhood rainfall for same-age cohorts within a district. The rainfall diagnostics described above show that these realised differences are not simply a continuation of district rainfall or a common shift across the two life stages; Section 5 examines the remaining threats to this interpretation.

The main specification implements this design with a gender-interacted fixed-effect structure. For individual i in district d, observed in quarter q of survey year t, at age a:

$$y_{idqt} = \sum_{P \in \{I,\, E\}} \left[\beta_P \, S^P_{dat} + \beta^g_P \, (G_i \times S^P_{dat}) \right] + \gamma^g_{daq} + \delta^g_t + \varepsilon_{idqt}, \tag{3}$$

where y_{idqt} is the time-use outcome (minutes per day, winsorised at the state × gender 99th percentile), G_i is a gender indicator (1 if male), S^P_{dat} is the life-stage shock for stage $P \in \{\text{Infancy } (I),\ \text{Early Childhood } (E)\}$, and γ^g_{daq} and δ^g_t are the fixed-effect blocks described below. The parameters of primary interest are β_E (the effect of the Early Childhood Shock on girls, the base group) and β^g_E (the boys-versus-girls differential); the total effect on boys is $\beta_E + \beta^g_E$.

The infancy coefficients β_I and β^g_I provide a developmentally distinct early-life comparison. The two windows are not placebos: they capture different ages at exposure, and either could in principle have lasting consequences. Their comparison becomes informative because the paper's proposed account concerns the greater opportunity for girls to engage in household tasks at ages four to seven. We average rainfall over four years because the estimand is exposure over a developmental stage rather than the effect of a single annual realisation. This grouping is not essential to the result: the rolling-window analysis estimates separate two-year exposure windows and locates the response across the same early-childhood years.[9]

The fixed-effect structure has two blocks. The first, district × age × quarter × gender $\left(\gamma^g_{daq}\right)$, absorbs the slow-moving district-age-gender-seasonal regularities most likely to confound the rainfall shock — the technology of home production, the seasonality of farm work, district-specific age profiles of girls' domestic work, and the gendered expectations around children's time use that vary across communities. The second, survey-year × gender $\left(\delta^g_t\right)$, absorbs national-level shifts in adolescent time use between 2019 and 2024 (schooling expansion, smartphone access, labour-market entry) that may themselves be gender-specific. Gender-interacting both blocks is algebraically equivalent to running separate regressions for girls and boys while preserving a pooled specification in which β^g_P can be tested directly. This is not merely a convenient way to report gender differences: with few survey waves, common fixed effects can mechanically load persistent gender differences onto a gender-specific

[9] A third window covering ages 7–10 can be constructed analogously; the rolling-window age profile reported in Section 4.1 examines these years without imposing the headline four-year grouping.

treatment coefficient. When later analyses compare rainfall effects across other observed groups, we either estimate the groups separately or interact the full fixed-effect structure with the group indicator. In both cases, the treatment effect is identified relative to group-specific district-age-quarter and survey-year baselines, rather than from pooled group differences. Appendix B, especially Sections B.3–B.5, gives the derivation. Standard errors are clustered at the district × age × survey-year level, the level at which the life-stage shocks vary.

The resulting estimate of β_E is the average effect for rural adolescents aged 12–17 in the 26 main states, identified from differences in realised early-childhood rainfall for cohorts whose ages 4–7 fall between roughly 2008 and 2018.

3.4.3 Identification Challenges and Alternative Specifications

The baseline comparison may be threatened if other district-level changes follow the realised childhood-rainfall differences between cohorts, if mortality, fertility, or migration change which adolescents are observed, or if household members reallocate tasks in response to one another's exposure. Further issues arise from constructing shocks from a gridded rainfall product and from spatial correlation, which determines the relevant uncertainty in the estimates. Section 5 addresses each of these issues with complementary evidence and alternative specifications.

The repeated cross-sections also permit two supplementary comparisons. The first compares age cohorts within the same district and survey round. This comparison is informative, but supplies less independent variation for the stage-level shocks used in the baseline specification: adjacent cohorts' four-year exposure averages share three of their four annual rainfall realisations. Comparing cohorts farther apart increases the independent rainfall variation, but also increases differences in age and cohort-specific institutional exposure. A cross-district specification would compare same-age cohorts across districts within a round. Although shocks are measured relative to each district's long-run rainfall distribution, realised rainfall shocks remain spatially correlated across neighbouring districts; the within-district comparison therefore has greater statistical power. Section 5 reports the alternatives as consistency checks; the baseline remains the comparison of same-age cohorts observed five years apart within districts.

The next section reports the headline estimates against this design.

4 Results

Building on the baseline specification in Section 3.4, we document the persistent relationship between rainfall during ages 4–7 and adolescent time use. We first establish the headline domestic-work effect, its extensive and intensive margins, and its composition across tasks (Section 4.1). We then examine which other uses of time move with it (Section 4.2), its timing within the day (Section 4.2), and its reach across the domestic-work distribution and household backgrounds (Section 4.3).

4.1 The Headline Effect

Higher rainfall during ages four to seven raises girls' domestic work in adolescence, a decade later. A one-standard-deviation increase in Early Childhood rainfall raises rural girls' domestic work by 5.3 minutes per day (Table 1). This is 6% of their 89-minute daily mean and widens the adolescent domestic-work gap by 4.3 minutes, while the implied effect for boys is close to zero. The result is therefore not a general increase in children's domestic work: it is an additional, persistent gender difference in girls' time allocation.

TABLE 1: *Early Childhood Shock and Adolescent Domestic Work*

	Domestic Work				
	Minutes per Day			**Extensive Margin**	**Intensive Margin**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Infancy Rain Shock	-0.6		1.0	0.016***	1.1
	(1.3)		(1.3)	(0.006)	(2.3)
× *Boys*	-1.3		-2.6*	-0.013*	-3.7
	(1.3)		(1.4)	(0.007)	(3.8)
Early Childhood Rain Shock		4.9***	5.3***	0.022***	6.6***
		(1.6)	(1.6)	(0.006)	(2.5)
× *Boys*		-3.4**	-4.3**	0.002	-7.8*
		(1.6)	(1.7)	(0.008)	(4.0)
Observations	59,466	59,466	59,466	59,466	21,751
Girls Mean	89.2	89.2	89.2	0.552	167.0
Boys Mean	19.4	19.4	19.4	0.202	100.9
Boys share in sample	0.53	0.53	0.53	0.53	0.29
[Girls] p: Infancy =Early Ch.			[0.02]	[0.39]	[0.05]
[Boys] p: Infancy =Early Ch.			[0.00]	[0.00]	[0.72]
District-Age-Qtr-Gender FEs	Y	Y	Y	Y	Y
Year-Gender FE	Y	Y	Y	Y	Y

Notes: **Sample.** Rural children aged 12–17, TUS 2019 + 2024 pooled, main 26 states, both genders. **What it shows.** Five columns. Cols (1)–(2) enter the Infancy and Early Childhood Shocks separately; cols (3)–(5) enter both jointly. The outcome is domestic-work minutes per day in cols (1)–(3) on the full sample, an indicator for any domestic work in col (4), and minutes per day conditional on positive domestic work in col (5). The Girls Mean and Boys Mean rows report the gender-specific within-sample mean of the column outcome. The [p] rows test, within each gender, equality of the two shock coefficients in cols (3)–(5). **Domestic work.** Total unpaid household-service time in minutes per day, winsorised at the state ×gender 99th percentile. **Infancy Shock.** Mean district rainfall z-score over ages −1 to 2, re-standardised within the estimation sample. **Early Childhood Shock.** Mean district rainfall z-score over ages 4 to 7, re-standardised within the estimation sample. **Estimation.** Both shocks are interacted with gender. Shock rows report girls' effects and ×Boys rows report boys–girls differentials. District-age-quarter-gender and survey-year-gender fixed effects are included. **Inference.** Standard errors clustered at district ×age ×survey-year, shown in parentheses; $***p<0.01$, $**p<0.05$, $*p<0.10$.

The persistent response is concentrated in exposure at ages four to seven, rather than

infancy. The girls' Infancy coefficient is only 1.0 minute per day, against 5.3 minutes for Early Childhood rainfall, and neither early-life window has a meaningful effect on boys' domestic work. The Early Childhood coefficient is unchanged when the two histories enter the regression together, and the within-gender tests reject equality of the two life-stage effects. This distinction matters for the rest of the paper: the result is not a generic association between favourable early-life rainfall and later domestic work, but is tied to a particular period of childhood.

The Early Childhood effect reflects both entry into domestic work and more time among girls already doing it; the latter separates it from infancy. Exposure at ages four to seven raises the probability that a girl records any domestic work by 2.2 percentage points and raises minutes among those already doing some by 6.6 minutes per day (Table 1). At baseline, these two margins each account for roughly half of the five-minute headline effect. Infancy has a similar, modest association with whether a girl reports any domestic work, but no corresponding effect on conditional minutes. The distinctive feature of the Early Childhood response is therefore a deeper later engagement with domestic work, not participation alone.[10]

Composition across tasks. Food preparation accounts for the largest part of the increase, but several tasks move together. Food preparation contributes 2.9 of the 5.3 additional daily minutes, close to its large baseline share of girls' domestic work (Table 10). Care work and shopping or budgeting also rise, while the estimates for cleaning and domestic production are positive but imprecise. The result is therefore not an increase confined to preparing food: girls' later domestic-work response spans the tasks through which households organise daily consumption and care. The food-preparation effect is also absent for boys, matching the headline gender pattern.

Girls exposed at ages four to seven also perform a wider range of domestic activities. The shock raises the number of distinct domestic-work activities that girls record by 0.145 codes and reduces the share of time concentrated in their single most-used activity by 1.7 percentage points (Tables 11 and 12). The broader repertoire is visible within food preparation and cleaning, and neither pattern appears for boys or following the Infancy Shock. Together with the task-level results, this indicates a broader domestic-work repertoire rather than longer time on one existing chore.

The next question is what the additional domestic work displaces from girls' time.

4.2 What Domestic Work Displaces

The increase in girls' domestic work documented in Section 4.1 must come at the expense of some other use of girls' time. The descriptive accounting in Section 3.1.4 offers a prior: among rural adolescents the gender gap in domestic work is offset in raw means by girls' lower paid work and lower leisure, while studying time is close to ungendered at every age.

[10] The descriptive decomposition evaluates $\partial \mathbb{E}[Y]/\partial S = \partial P(Y > 0)/\partial S \times \mathbb{E}[Y \mid Y > 0] + P(Y > 0) \times \partial \mathbb{E}[Y \mid Y > 0]/\partial S$ at baseline values.

This subsection asks which use gives way, and whether the answer changes with the school calendar. Day type matters because the same daily demand for domestic work may be met out of different margins on school days and non-school days. We therefore separate Sunday — a weekly school holiday in nearly every Indian state[11] — from Monday-to-Saturday diary days, and then examine weekday non-normal days as a second comparison.

Among girls' other major time uses, studying is the only one with a coefficient of meaningful size. Table 13 reports the Early Childhood effect on girls' three other major time uses. Of studying, leisure, and paid work, only studying carries a sizeable coefficient: -4.7 minutes per day, with a standard error of 2.9, not significant at conventional levels. Leisure (-0.9) and paid work ($+0.3$) are small and null.[12] The pooled daily total does not reveal on which days the studying adjustment occurs, which is the question the day-type comparison resolves.

Variation across day types. On ordinary school days, the additional domestic work is accompanied by less studying; on Sundays, studying returns and leisure gives way instead. The Early Childhood Shock raises domestic work by 5.4 minutes on Monday-to-Saturday diary days and by a statistically indistinguishable amount on Sundays (Table 2). Studying moves differently: it is 6.9 minutes lower on ordinary school days but about 10 minutes higher on Sundays, with the Sunday recovery drawn from leisure. Thus the additional domestic work persists across day types, while the time cost of meeting it depends on whether school occupies the day. The separate Infancy pattern is small and imprecise, and the result survives the specification that allows the day-type baseline to vary by district and quarter.

The school-day adjustment is in attendance and its commute, while food preparation accounts for the added domestic work. The weekday reduction in studying is concentrated in recorded attendance and the associated commute, not in homework; the Sunday recovery occurs on those same margins (Table A2). On the work side, food preparation accounts for the weekday increase and does not increase further on Sundays. This decomposition points to an adjustment around the school-day schedule, rather than a general reduction in study carried out at home.

The same reallocation appears when an ordinary weekday is disrupted. Sundays could reflect a planned weekly routine, so we also construct a household-level non-normal-day measure for Monday-to-Saturday diaries. It equals one when at least one quarter of all household diarists, including the focal girl, classify the day as non-normal — for example, a festival, public holiday, family event, illness, or another non-ordinary day. The indicator is therefore not the girl's own day-type alone, although her report contributes to the household share.

[11] Sunday is the weekly school holiday across essentially all Indian states. The rural government-school week otherwise runs Monday to Saturday; Saturday is a full or a reduced (half- or activity-) day depending on the state and school management, and only a minority of (largely private or metropolitan) schools observe a five-day week. Because Saturday is a school day for most of the rural sample, we group Monday to Saturday as school days and test only the Sunday differential, treating Sunday as the clean non-school day; the few Saturday-holiday or half-day cases attenuate the school-day contrast we estimate rather than manufacture it.

[12] Boys' studying coefficient is comparable in magnitude (-3.4 minutes per day, not significant), but their released time appears to flow into paid work rather than domestic work: the boys-girls differential on paid work is $+2.8$ minutes per day, significant at the 10% level, against effectively no movement on boys' domestic work.

TABLE 2: *Study Time and Leisure by Day Type*

	Domestic Work		Studying		Leisure	
	Mins *(mins/day)*	**≥ 120** *(share)*	**Mins** *(mins/day)*	**≥ 360** *(share)*	**Mins** *(mins/day)*	**≥ 960** *(share)*
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Infancy Rain Shock	0.9	0.003	2.3	0.007	-5.0**	-0.003
	(1.3)	(0.005)	(2.6)	(0.005)	(2.2)	(0.006)
× *Sunday*	1.1	0.001	-2.5	-0.021**	1.8	-0.006
	(3.2)	(0.012)	(4.8)	(0.011)	(4.6)	(0.011)
Early Childhood Rain Shock	5.4***	0.015***	-6.9**	-0.012*	1.4	0.008
	(1.7)	(0.006)	(2.9)	(0.006)	(2.5)	(0.007)
× *Sunday*	0.9	0.010	16.7***	0.044***	-18.2***	-0.043***
	(3.2)	(0.011)	(4.7)	(0.010)	(4.6)	(0.011)
Observations	27,977	27,977	27,977	27,977	27,977	27,977
Weekday Mean	84.6	0.280	367.0	0.660	974.3	0.481
Sunday Mean	127.4	0.444	164.2	0.165	1128.3	0.827
[Weekday] p: Infancy =Early Ch.	[0.01]	[0.06]	[0.01]	[0.01]	[0.03]	[0.17]
[Sunday] p: Infancy =Early Ch.	[0.35]	[0.19]	[0.11]	[0.00]	[0.02]	[0.06]
District-Age-Qtr FEs	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Sunday-FE	Y	Y	Y	Y	Y	Y

Notes: **Sample.** Rural girls aged 12–17, TUS 2019 and 2024 pooled, in the main 26 states; all diary days. **What it shows.** Domestic work, studying, and leisure, each measured in minutes per day and by an indicator for at least 120, 360, and 960 minutes, respectively. Shock rows report effects on Monday–Saturday diaries; ×Sunday rows report the additional effect on Sunday diaries. **Outcomes.** Domestic work is unpaid household services. Studying is formal study time. Leisure combines sleep, eating, personal care, passive media, and social contact. **Estimation and inference.** Regressions interact each rainfall shock with an indicator for Sunday and include rural ×district ×age ×quarter and survey-year fixed effects. Standard errors are clustered at district ×age ×survey-year and shown in parentheses. $***p<0.01$, $**p<0.05$, $*p<0.10$.

We repeat the exercise at a stricter threshold requiring at least three quarters of household diarists to report a non-normal day.

The non-normal-day interaction does not materially change the Early Childhood effect on domestic work. However, as on Sundays, a disruption to the household's ordinary routine changes which use of time adjusts. Relative to ordinary Monday-to-Saturday diaries, the Early Childhood study-time coefficient is about 20 minutes higher on a household-level non-normal day. This more than offsets the 6–9-minute study deficit on ordinary weekdays, leaving a positive implied study effect; leisure falls by a similar amount (Table A3). The school-day pattern is therefore not limited to the regular Sunday cycle: when the household day is less likely to follow its ordinary schedule, study time returns without a further increase in domestic work.

Timing within the day. The day-type results show when study time is lower, but not where the added domestic work occurs within the day. If it is concentrated in meal production or other daily routines, it should appear around those routines; a more general reassignment of chores would be more evenly spread. Because school occupies fixed daytime hours, the hourly profile also shows whether domestic work and studying move at the same time or whether the adjustment operates across the day as a whole.

We estimate the headline specification separately within each 30-minute slot of the day, separately for weekdays and Sundays, first for domestic work and then for studying. Each figure plots the within-sample mean profile across the day together with the Early Childhood and Infancy shock coefficients.

Across the day, the extra domestic work lands in the meal hours while the lost study spans the school-day part of the day. On weekdays the Early Childhood Shock raises girls' domestic work in two windows: the morning around breakfast (about 0.3 minutes per 30-minute slot at 06:00–07:00, significant at the 5% level) and, more sharply, the evening dinner hours, where it reaches roughly 0.5 minutes per slot between 18:30 and 19:30 (1% level); the Sunday profile adds further in the evening (Figure A4). Studying moves in the opposite direction, and in a different part of the day: the weekday loss runs from mid-morning through mid-afternoon, and reappears in the evening dinner window (Figure A5). The Infancy coefficient is flat and indistinguishable from zero at almost every hour on both day-types, the one exception a positive domestic-work blip in the early Sunday morning.

Read together, the day-type and hourly comparisons narrow the descriptive account. The additional work is not assigned only on days when the girl has the most free time: baseline domestic work is substantially higher on Sundays and non-normal weekdays, yet the rainfall-induced increment is approximately constant across day types. Nor is it a minute-for-minute replacement of studying by domestic work in the same hours. Domestic work rises around breakfast and dinner, while much of the weekday studying difference occurs during the school-day portion of the day. Study time returns when the regular school schedule is relaxed. This pattern is consistent with a daily household-production requirement, but the model and mechanism sections assess the source and persistence of that requirement.

4.3 Distributional Reach

The headline is an average, and the same 5.3-minute effect could be generated by different changes beneath it. It could reflect a few girls beginning light domestic work, a broad shift in minutes, or larger increases among girls already doing a great deal. We distinguish these cases using two complementary tools: the effect on the probability of clearing each point on a ladder of domestic-work thresholds, and quantile estimates across the distribution under the location-shift approach of Canay (2011).[13] Both exercises describe where the response

[13] The quantile estimates follow Canay (2011): the fixed effects are removed in a first step under a location-shift assumption, which treats them as shifting the whole conditional distribution by a constant, and quantile regressions of the de-meaned outcome on the two shocks are then estimated across the distribution.

appears in the outcome distribution; they do not identify endogenous selection into which girls perform domestic work.

Across the distribution, the Early Childhood effect reaches the heaviest domestic workers, while the Infancy effect stays near the bottom. Section 4.1 showed that the effect appears on both the extensive and the intensive margin, and that the intensive margin is what separates the Early Childhood shock from the Infancy shock. Where in the distribution does that fall? Figure 6 traces the effect on the probability of clearing each domestic-work threshold, from any domestic work at all up to five hours a day. The Early Childhood effect runs along the whole ladder. It raises the probability of doing any domestic work by 2.2 percentage points, and it stays positive and statistically significant at every rung above that, including the highest: the probability of clearing three hundred minutes a day rises by about one percentage point. The Infancy effect behaves differently. At the foot of the ladder it raises participation by a comparable amount (1.6 percentage points at any domestic work, statistically indistinguishable from Early Childhood there), but it fades to zero by the middle of the range and is absent at the higher thresholds. The quantile estimates tell the same story from the other side (Figure 7): the Early Childhood effect on girls' domestic work is modest near the median and grows across the upper tail, reaching roughly twelve minutes a day at the seventy-fifth percentile and sixteen at the ninetieth.

Taken together, the two patterns weigh against the reading that the effect simply pulls a few girls across the participation margin into light domestic work. A pure-entry response could raise the probability of doing any domestic work without changing the time use of girls already engaged in it. Instead, the Early Childhood effect remains positive at high domestic-work thresholds and grows in the upper tail of the positive-work distribution. The more direct reading is that the shock raises domestic work across the distribution, and most of all among girls already doing a great deal; the modest participation response, which it shares with the Infancy shock, is not what marks the Early Childhood shock out.

Household-background heterogeneity. We next ask whether the headline coefficient differs across household backgrounds. Table A4 reports girls-only regressions that interact each childhood shock with one binary characteristic at a time. We treat this exercise as exploratory. The household characteristics are correlated with one another, subgroup sizes and baseline domestic work differ, and each interaction may stand in for several features of the household environment. The estimates can locate where the response is more or less visible, but they do not separately identify why those households differ.

The Early Childhood coefficient remains positive across the six reference groups, ranging from 4.0 to 6.7 minutes per day. Two interaction differentials are negative and precisely estimated: -7.7 minutes for daughters of mothers educated above secondary school and -5.8 minutes in joint families, leaving small implied effects in both groups.[14] The remaining interactions do

[14] Households are classified using the standard Indian sociological typology of Karve (1953); Kolenda (1968), distinguishing *nuclear* (parents and unmarried children), *extended* (nuclear plus a widowed parent-in-law),

not provide evidence of differences by land ownership, maternal age, grandparent presence, or religion.

Girls in the high-mother-education subgroup have much lower baseline domestic work (42 minutes per day, against 90 in the reference group), so its attenuation occurs in households with different domestic-work demands as well as different maternal education. Joint-family status changes the availability of other household workers, but also household composition, resources, and task organisation. We therefore use these estimates to describe where the response is weaker, not as separate evidence for a particular mechanism.

Summary and discussion. The headline estimate is a durable, girl-specific change in adolescent time allocation: rainfall during ages four to seven raises girls' later domestic work by 5.3 minutes per day, or 6% of their daily mean, with no comparable response among boys. Two features of the result make it more than an average difference in daily minutes. The first concerns the depth of girls' engagement in domestic work. The second concerns the way its time cost is met around the school schedule. A minutes effect alone does not determine its economic importance: if early task experience also raised the productivity of domestic work, it could understate the change in the girl's role in household production. We examine that possibility below.

First, the response is a deepening of domestic-work engagement, not merely a marginal-entry result. It appears in minutes among girls already doing domestic work, at high domestic-work thresholds, and in the upper tail of the domestic-work distribution; it also spans several domestic tasks. The Infancy Shock shares a modest participation response, but not this deeper and broader pattern. This makes the later mechanism question sharp. If early exposure merely lowers the opportunity cost of assigning a girl's time to domestic work, one might expect it principally to draw additional girls into the activity. If domestic-work skill compounds with experience, the largest later response may instead appear among girls already carrying substantial responsibility. The reduced form does not distinguish these possibilities, but it identifies the empirical pattern that the later analysis must explain.

Second, the additional work looks like a routine household requirement rather than an activity concentrated on one particular day. The Early Childhood effect on domestic work persists on ordinary school days, Sundays, and household-level non-normal weekdays. What changes is the margin from which its time cost is met: studying is lower when school is in session, but recovers when the schedule is relaxed, at the expense of leisure and without a further increase in domestic work. A simple reduction in the returns to studying does not naturally predict this recovery on non-school days while the domestic-work response remains unchanged. The pattern is instead consistent with routine domestic tasks that cannot be fully deferred or stored across days, while study time is more movable. These findings

stem (two conjugal units of consecutive generations), and *joint* (three or more conjugal units, often of the same generation) households by their internal generational composition. Both differentials remain significant under a coarse Bonferroni adjustment across the subgroup tests.

motivate the mechanism questions that follow. Before interpreting them, Section 5 examines the validity of the reduced-form comparison. Sections 6 and 7 then develop and test the household-allocation account.

5 Identification and Robustness

The baseline specification compares girls of the same age in the same rural district across the two TUS waves, who faced different rainfall histories at the same childhood ages. This comparison could be misleading if the district-level differences in childhood rainfall also track other changes between the two waves that independently affect adolescent time use. Early rainfall could also affect which children remain in the observed adolescent cohort, while the measured coefficient could depend on how gridded rainfall and multi-activity diaries are constructed. We therefore examine confounding changes within districts, selection into the observed sample, interference across siblings, measurement choices, and spatial inference.

Confounders. The fixed effects absorb time-invariant district characteristics and gender-specific survey-year shifts. The remaining concern is that another district-level change between 2019 and 2024 happens to be correlated with the realised Early Childhood shock difference and independently affects girls' domestic work. Figure 5 compares the 2019–2024 cohort differences with the historical distribution of five-year-apart differences in district rainfall. The historical distribution is centred near zero, whereas the realised Early Childhood and Infancy distributions have different calendar-window means. Those mean shifts are absorbed by the survey-year effects. What remains for identification is the substantial two-sided district-level variation in each realised distribution: the TUS comparison is not a uniform contrast between wet and dry cohorts. Figure 4 provides a complementary check. The Early Childhood and Infancy differences show little spatial correspondence, consistent with distinct realised weather variation across the two childhood windows rather than a common shift in district weather. These diagnostics do not rule out confounding, but they make its required form specific: an omitted district-level change would need to be correlated with the particular spatial pattern of the realised Early Childhood differences, and independently raise girls' domestic work.

Table 14 first asks whether the estimated relationship is explained by household margins that may both vary with district-level shock differences and predict adolescent time use. It adds parental age and education, land, caste, religion, household size, number of children, and gender interactions with these controls. The girls' Early Childhood Shock coefficient remains between 5.23 and 5.39 minutes per day. Figure 8 broadens this exercise, varying the fixed-effect structure across 20 alternatives and cumulatively adding six control sets, including household-production conditions such as cooking fuel, lighting, washing, sweeping, and dwelling structure. Its 174 estimates remain concentrated and almost entirely positive, with substantial mass above the 5.32-minute headline. The cross-sectional branch of this exercise

includes state × survey-year fixed effects. These are demanding comparisons: they remove all shocks common to districts within a state and survey year, leaving identification from differences across districts within the same state-year. The continued positive estimates therefore weigh against state-wide policy changes and other common state shocks, including changes in education policy, rural employment programmes, mid-day meals, and fiscal responses to COVID. The companion figure shows that the pooled time-use pattern reported in Section 4.2 is also stable across these specifications: studying is the only other major activity with a meaningful negative coefficient in most specifications, while leisure and paid work remaining small and variable across specifications (Figure A6).

Leave-one-state-out estimates provide a separate influence check, rather than a further test of state-level confounding. Figure A7 re-estimates the headline after dropping each of the 26 states in turn. The girls' Early Childhood coefficient ranges from 3.6 to 6.4 minutes per day around the full-sample estimate of 5.3 and remains positive throughout, including when the largest migration-source and migration-destination states are excluded. Figure A8 provides the corresponding descriptive geographic evidence: the separately estimated state and NSS-region coefficients are heterogeneous, but positive associations are not confined to one state or one region. The map is not a formal test of each local coefficient. Taken together with the shock map, it nevertheless strengthens the geographic argument. The realised Early Childhood shocks are positive in some districts and negative in others, while positive domestic-work effects appear across several states and NSS regions. An omitted factor that raises girls' domestic work could generate this pattern only if it were high where the realised Early Childhood shock is high and low where it is low, across these separate locations. A common state-wide or regional change cannot generate that district-specific alignment.

Taken together, the evidence narrows the form that any remaining confounder could take. It is difficult to identify a plausible omitted factor that satisfies these conditions. Such a factor would have to be district-level and time-varying, correlated with the particular realised spatial pattern of the Early Childhood Shock but not with the distinct Infancy pattern, and independently affect girls' domestic work after common state-by-year shocks are removed. It would also have to be absent from the observed household and production margins. This does not rule out a remaining confounder.

Selection. Early-life rainfall can affect fertility, mortality, and migration (Maccini and Yang, 2009; Almond and Currie, 2011). If these responses altered the composition of the cohort observed at ages 12–17, the Early Childhood coefficient could partly compare different adolescents rather than different childhood rainfall histories. Table A5 examines the observable margins on which this selection should appear: the cohort's female share and its size relative to other age cohorts or to the same age cohort in the other survey round. Each test is also repeated for the eldest-in-household sample used in the sibling analysis below.

Columns 1 and 2 regress an indicator that a child is female on the two life-stage shocks. A response would be consistent with selective fertility, gender-selective mortality, or gender-

differential migration before the child enters the adolescent sample. The Early Childhood coefficients are 0.001 in the full cohort and 0.002 for eldest children, against female shares of 0.470 and 0.467; neither is statistically distinguishable from zero. The Infancy coefficients are also close to zero. The rainfall histories therefore do not predict the sex composition of the observed adolescent cohort.

Columns 3–4 compare each age cohort's share of rural girls in its district-round with the shares of the other ages; columns 5–6 compare the same district-age-quarter cell across the two TUS rounds. These two constructions address different forms of selection. The first asks whether the affected cohort is unusually large or small relative to other adolescents observed in the same district-round. The second asks whether it is unusually represented in one survey round relative to the other. The displayed means of 0.742 and 0.751 in the latter construction are not means for a balanced two-round comparison: 51–57% of district-age-quarter cells appear in only one TUS round and therefore contribute a value of one in their observed round. Those cells are absorbed by the district-age-quarter fixed effects; among cells observed in both rounds, the corresponding mean is about 0.58. The Early Childhood coefficients are effectively zero in cols 3–4 and equal -0.008 and -0.007 in cols 5–6; none is statistically distinguishable from zero. The Infancy Shock provides a useful contrast: its coefficients are negative and statistically significant in cols 3–5 (-0.005, -0.009, and -0.017), although the eldest across-round estimate attenuates to -0.011 and is significant only at the 10% level. The data do not distinguish a modest birth-year demographic or migration response from an artefact of the cohort-share construction. But the contrast matters for the headline result: selection evidence, to the extent it appears, is confined to Infancy rather than the Early Childhood window.

These tests cannot rule out selection that leaves both sex composition and cohort size unchanged. For example, rainfall could change which households migrate or provide usable diaries while offsetting flows leave the number of girls and the female share unchanged. They do show that an account based on differential survival, fertility, or migration would have to alter an unmeasured feature of the observed cohort without producing a corresponding shift in any of the three composition margins, and only for the Early Childhood Shock. The evidence therefore weighs against broad cohort-composition change as an explanation for the headline coefficient.

SUTVA Concerns. When an early shock changes one girl's capacity or willingness to perform domestic tasks, the household may reallocate work toward her and away from her sisters. A child's observed domestic work then depends not only on her own rainfall history but also on the histories of other children in the household. Including each of those children as a separate focal observation creates a problem similar in nature to well-known SUTVA concerns in randomised trials: the increase for the girl who takes on tasks can be paired with a reduction for sisters who are indirectly relieved of them, making the estimated effect of her own rainfall history larger than the direct effect holding her siblings' task allocation fixed.

Table 15 addresses this concern by retaining only the eldest child under 18 in each household, whether that child is a girl or a boy. Each household therefore contributes one focal observation, so a sibling whose work has been reallocated is not also used as an additional treated or comparison child in the estimating sample. This changes the object of estimation: it describes the association for eldest children rather than a fully representative sample of all adolescents. The girls' Early Childhood coefficient is 6.40 minutes per day, compared with 5.32 minutes in the headline sample. The coefficients should not be read as a direct equality test, since eldest children are older, have different birth order, and report more domestic work on average.

For completeness, column (3) adds household × gender fixed effects to the full sample and identifies the coefficient from comparisons among same-gender children within a household. Its girls' Early Childhood point estimate is similarly positive, at 6.31 minutes per day, but imprecise under this highly saturated specification. This is an extreme within-household comparison, not the preferred estimate. Together with the eldest-child restriction, it shows that the positive result is not confined to the particular across-household composition of the headline sample.

Measurement and Sample Construction.

Rainfall Assignment. The preceding exercises concern the children and households entering the regression. We next ask whether the result depends on the measurement choices required to assign gridded rainfall to districts and to convert multi-activity diaries into daily domestic-work minutes.

The immediate rainfall-assignment issue is that GPCC is reported on a $1^\circ \times 1^\circ$ grid while outcomes are observed in districts, so the grid does not align neatly with district boundaries (Section 3.2.1). The headline uses the Monitoring product, which covers the relevant exposure years through 2024, and assigns nearby grid cells with a buffered, inverse-distance rule. We therefore ask whether the result depends on that product and assignment rule.

Table 16 first varies the rainfall-shock construction while holding the diary rule fixed. The headline GPCC Monitoring buffer construction gives a girls' Early Childhood coefficient of 5.31 minutes per day. The estimate is 5.76 when rainfall is built from grid cells whose centroids fall strictly inside the district polygon, 5.66 under radial-distance weighting, and 5.24 when the GPCC Full quality-controlled product replaces the Monitoring product under the inside-polygon rule. The same table also considers alternative ways to construct the rainfall shock over the agricultural year. As expected, the April–September monsoon months carry most of the relevant variation. The monsoon shock raises girls' domestic work by 6.5 minutes per day, while the post-monsoon shock is small and statistically insignificant at 1.3 minutes. A linear tail coding of annual rainfall, following Kaur (2019), gives a standardised Early Childhood coefficient of 3.6 minutes per day and a raw-scale coefficient of 11.2. The close positive estimates show that the result does not rest on one rainfall product or grid-to-district

rule. The seasonal and functional-form columns address the agricultural interpretation and are considered in Section 7, rather than serving as additional measurement checks.

Diary Construction and Survey Weights. The TUS records up to three activities in each 30-minute slot, and roughly one quarter of girls' recorded domestic-work minutes occur in slots with more than one activity (Section 3.1.1). No accounting rule reveals how each half-hour was divided. The baseline splits such slots equally, preserving the daily time budget; the alternatives assign the slot to the reported major activity or distinguish sequential from simultaneous activities.

Table A6 then varies the diary-slot coding and survey weighting. Under the headline equal-split rule, the girls' Early Childhood coefficient is 5.3 minutes per day. Assigning each multi-activity slot to the major activity gives 5.4, and the hybrid rule that preserves equal splitting for sequential activities but assigns simultaneous activities to the major activity gives 5.1. Applying the TUS person weights yields coefficients of 4.7, 4.7, and 4.5 under the same three coding rules. Each estimate remains positive and statistically significant. The result is therefore not an artefact of proportional slot splitting or of estimating the headline specification without survey weights. The four-year shock-window choice is addressed separately by the rolling-window age profile in Section 7.

Replacing the state-by-gender p99 winsorised outcome with raw daily domestic-work minutes also leaves the girls' Early Childhood coefficient at 5.3 minutes per day. Standard errors are marginally larger with the untrimmed upper tail, but the qualitative inference is unchanged.

Age Composition. In the realised sample, the Early Childhood Shock is positively correlated with age within some survey cells. This does not bias the baseline coefficient, which compares children within the same age cells. It does mean that the pooled estimate gives more weight to ages that are more prevalent in the realised sample. Table A7 therefore either gives each age equal weight within district-quarter-survey-year cells or estimates age-specific effects before averaging them equally. The resulting Early Childhood coefficients are 6.47 and 5.23 minutes per day, bracketing the 5.32-minute headline. The positive relationship is thus similar under two alternative ways of aggregating the age-specific comparisons.

External Validation. The TUS result concerns the allocation of a girl's day, including adjustments within school and non-school hours. The Periodic Labour Force Survey (PLFS) records principal weekly activities and days rather than a daily time diary, so it cannot be expected to reproduce those within-day adjustments or the magnitude in minutes. It can nevertheless show whether the broader gendered pattern appears when domestic duties and studying are measured with a different, conventional labour-force instrument.

Table A8 re-estimates the specification for 85,041 rural children aged 12–17 in the 2019–2023 PLFS rounds. The PLFS is India's main national labour-force survey and the conventional source for measures of principal activity and work participation. Girls' Early Childhood Shock raises days in domestic duties by 0.10 per week, while the boys–girls differential is

−0.17 days; principal domestic-duty status moves in the same direction. Girls' studying days and student status both decline. These estimates are directionally consistent with the TUS evidence. They are a limited check on whether the broad gendered pattern is specific to the TUS diary, rather than a second estimate of the daily time-use response.

Inference: Spatial Dependence and Window Selection. Rainfall shocks are spatially correlated, so standard errors must allow for correlation across nearby observations. Columns (1)–(6) of Table A9 retain the 5.3-minute point estimate and vary the standard-error calculation. The coefficient remains statistically significant with district, state, and NSS-region clustering and with Conley standard errors at 200- and 500-km cutoffs. A state-level wild-cluster bootstrap gives a nominal p-value of 0.031. Together with the leave-one-state-out estimates, these results show that, for the maintained exposure windows, the positive result does not depend on the baseline clustering rule or one state.

Column (7) addresses selection among eight four-year exposure windows, starting at ages 0 through 7, while holding Infancy exposure (ages −1 to 2) and its gender interaction fixed in every specification. The Bonferroni correction is applied separately across the eight girls' coefficients and across the eight boys–girls differentials. Under baseline clustering, the adjusted p-value is 0.0098 for girls' Early Childhood coefficient and 0.0910 for the boys–girls differential. Applying the same eight-window correction to the state wild-bootstrap test gives an adjusted p-value of 0.251. The girls' coefficient therefore remains statistically significant after the window correction under baseline clustering, but not when that correction is combined with state-level bootstrap inference.

Inference: Using Placebo Rain Shocks. The baseline comparison uses the children's actual birth years: within each TUS wave, birth year changes one-for-one with age, and same-age children observed in 2024 belong to cohorts born five years after those observed in 2019. The placebo preserves this within-wave age structure but changes its calendar placement independently across the two waves. For each pair of anchor years, it assigns a placebo birth year to age-17 children in each wave and then assigns successively later birth years to younger children; it recomputes their infancy and early-childhood shocks from the same district-year rainfall series, re-standardises them, and re-estimates the headline regression. The anchor range is 1983–2012, which permits all required exposure years for the full age 12–17 sample to lie within the 1982–2024 rainfall record. We enumerate the 870 ordered pairs of distinct anchors (30×29); equal anchors are excluded because they assign the same placebo exposure to same-age children in both waves and leave no cross-wave variation after the fixed effects are absorbed.

The girls' Early Childhood coefficient of 5.30 minutes per day lies in the right tail of that placebo distribution ($p = 0.032$, one-sided), whereas the girls' Infancy coefficient and both boys' coefficients are not distinguishable from their respective placebo distributions (Figure A9). The placebo exercise therefore recovers the particular pre-school-girl pattern in the main estimates without relying on a cluster-robust variance formula.

The positive result survives checks on the cross-wave comparison, observed cohort selection, within-household reallocation, the construction of rainfall and time-use measures, and spatial inference. It also appears directionally in the PLFS, a separate conventional labour-force survey. These exercises do not themselves determine why early exposure has a persistent effect. Having established the empirical pattern, the next section develops a household-allocation framework to distinguish the mechanisms that can generate it.

6 Model

The reduced-form evidence raises a question that an ordinary income effect does not answer. Why can a favourable shock to non-labour income lead a household to use more domestic work from a girl whose hour is less productive than an adult's? In rural households, additional non-labour income can increase the materials and time needed for home-produced consumption. Whether that demand is met by adult or child time depends not only on relative productivity, but also on the availability and opportunity cost of adult time.

We use a minimal two-period household-allocation model to make this logic explicit. The model also separates two reasons why a childhood episode might persist into adolescence. Early domestic work may raise the daughter's efficiency in the domestic tasks that produce home consumption, or it may lower the opportunity cost of assigning her hour to those tasks. These channels look similar in time-use data, but they differ when household output is considered alongside household time. We use that distinction to derive a one-sided empirical implication for domestic-work-specific productive capacity. The model is not intended as a structural account of every way rainfall affects wages, paid work, agricultural output, or income. Instead, Y represents the non-labour-income component of the shock. Appendix C records the common identities and proofs for the propositions.

6.1 Model Setup

Timing and household members. There are two periods. Period 1 is early childhood, when an unanticipated shock to non-labour income can change the household's demand for domestic work; period 2 is adolescence, when we observe time use. In each period the household solves a static allocation problem. We do not model a household that assigns a child to chores in order to invest in a later state. Rather, period-1 domestic work is an unintended consequence of the period-1 allocation, and the reduced-form laws of motion below describe what that experience may carry into adolescence.

The household is represented by a mother and a daughter. In the model, "mother" denotes the adult household member whose time is principally allocated to domestic work, while retaining the option of paid work. She may represent one adult woman or a composite of adult women. Other household members are not assigned separate time choices in this partial model; their earnings, farm returns, and transfers enter non-labour income Y_t. The mother divides her time between paid work $h^p_{m,t}$, domestic work $h^p_{d,t}$, and leisure $h^p_{l,t}$. The

daughter divides hers between domestic work $h^c_{d,t}$ and residual time $h^c_{o,t}$:

$$h^p_{m,t} + h^p_{d,t} + h^p_{l,t} = 1, \qquad h^c_{d,t} + h^c_{o,t} = 1. \tag{4}$$

Residual time includes studying, school-related time, leisure, and other non-domestic activities. The interpretation of this aggregation is discussed below.

Home production and non-labour income. The household consumes two forms of final consumption: a ready-purchased market good, $C_{m,t}$, and a home-produced good, $C_{h,t}$. The latter converts materials into final consumption using effective domestic work:

$$C_{h,t} = \min\left\{M_{r,t},\ bD_t\right\}, \qquad D_t = h^p_{d,t} + \gamma_t h^c_{d,t}, \qquad b > 0. \tag{5}$$

Here $M_{r,t}$ denotes materials or intermediate inputs, such as food ingredients and other purchased inputs used in household production, and D_t denotes effective domestic work. The daughter supplies γ_t effective units per hour relative to the mother, with $0 < \gamma_t \leq 1$. We refer to γ_t as the daughter's *domestic-production skill*: a task-specific capability that raises the effective domestic input generated by an hour of her work. Holding materials and required home-produced consumption fixed, a higher γ_t reduces the time required from her to meet a given domestic-production requirement. It does not refer to general human capital, market productivity, or the quality of household output.

At an optimum, $M_{r,t} = bD_t$. The interpretation of this production technology and its relation to market consumption are discussed below.

The price of the market good is p_m, the price of materials is p_r, and w_t is the return to an hour of the mother's paid work. We hold p_m, p_r, and w_t fixed in the non-labour-income comparative static, isolating a change in Y_t. The household faces the cash budget

$$p_m C_{m,t} + p_r M_{r,t} = w_t h^p_{m,t} + Y_t. \tag{6}$$

The term Y_t denotes non-labour income relative to the mother's paid-work choice. It can include other members' paid earnings, farm profits, transfers, and asset income. In a rural agricultural household, it also includes the market value of farm output retained for home consumption, which can be used as household-production material rather than sold.

Preferences. The household values total consumption rather than either component alone:

$$C_t = \mathcal{G}(C_{m,t}, C_{h,t}). \tag{7}$$

The aggregator $\mathcal{G}$ combines market goods with final home-produced consumption. We assume that it is regular and constant returns to scale. The interpretation of this restriction is discussed below.

Household utility is defined over final consumption, the mother's leisure, and the daughter's

residual time. The maintained specification uses CES preferences over these three objects,

$$u_t = \left[a_C C_t^{\frac{\sigma-1}{\sigma}} + a_l \left(h_{l,t}^p\right)^{\frac{\sigma-1}{\sigma}} + a_{o,t} \left(h_{o,t}^c\right)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \qquad \sigma \geq 1. \tag{8}$$

The weights a_C, a_l, and $a_{o,t}$ are positive. The role of the maintained elasticity restriction $\sigma \geq 1$ is discussed below.

Persistence. Early domestic engagement can then carry forward in two distinct ways:

$$\log \frac{\gamma_2}{\gamma_1} = A_\gamma h_{d,1}^c, \qquad \log \frac{a_{o,2}}{a_{o,1}} = -A_\psi h_{d,1}^c, \qquad A_\gamma, A_\psi \geq 0. \tag{9}$$

The parameter A_γ captures learning-by-doing in domestic-production skill. By contrast, A_ψ leaves what an hour of domestic work produces unchanged and lowers the opportunity cost to the household of assigning the daughter's time to it. Their interpretation and the empirical mapping are discussed below.

6.2 Discussion of Assumptions

Household representation. The mother represents the adult household member whose time is principally allocated to domestic work, while retaining the option of paid work. Fathers and other members are not assigned separate time choices. Their earnings, farm returns, and transfers enter Y_t; in settings where men are principally allocated to farm or market work, this includes rainfall-sensitive male earnings. Treating rainfall as a shock to Y_t therefore isolates its non-labour-income channel while holding fixed the mother's return to paid work. The model does not describe wage, price, or labour-demand responses to rainfall.

Home production and consumption. The fixed-proportions technology concerns the conversion of materials and effective work into core home-produced consumption. Ingredients require preparation before they become meals, and preparation without ingredients does not add to final consumption. This is most directly a representation of food preparation, cleaning, and related tasks. It does not fix the household's overall labour intensity of consumption: the household can purchase C_m ready for use, choose its mix of market and home-produced goods, and differ in b, for example with access to household durables or more convenient inputs. Constant returns of $\mathcal{G}$ permit a general price-dependent mix of market and home-produced final consumption. Its substantive restriction is homotheticity: at given relative prices, a rise in non-labour income does not itself shift expenditure toward prepared goods, appliances, or other labour-saving technologies. The corner result also maintains that the household does not substitute so sharply toward market consumption as home production becomes relatively more costly that additional materials cease to require domestic work. This is a restriction on price responses within $\mathcal{G}$, separate from constant returns; Appendix C.3 gives its exact form. The one-sided output-per-domestic-work result rests on the production identity itself.

Time and persistence. Residual time combines studying, school-related time, leisure, and other non-domestic activities. The reduced-form evidence motivates this simplification: domestic work displaces studying on ordinary school days, but studying recovers on less constrained days and leisure adjusts instead; there are no broad losses in enrolment, attainment, or measured learning. The model consequently treats these uses as a common alternative to domestic work, without modelling a separate within-day allocation or dynamic return to study time. A rise in γ is an increase in domestic-production skill: an hour of the daughter's domestic work produces more final home consumption relative to an adult hour. A fall in a_o lowers the weight on her residual time in the household allocation; it can represent lower returns to studying or other alternative uses, or lower disutility from domestic work. Health belongs in γ only when it raises task-specific domestic productivity. Effects on usable time, energy, or returns to study would require additional state variables.

The adult paid-work margin. When the mother supplies paid work, the market wage prices adult time. An increase in non-labour income then expands the household budget at fixed prices and can be accommodated by reducing paid work. When she supplies no paid work, the market wage remains the return available on re-entry but does not price a marginal hour within the household; the corner allocation is governed by an internal shadow value of adult time that responds to non-labour income. Appendix C.1 derives the two allocation problems. Proposition 1A and Proposition 1B use the positive-paid-work and corner derivations in Appendix C.2 and Appendix C.3, respectively.

Substitutability across consumption and time. The elasticity σ governs substitution among total consumption, the mother's leisure, and the daughter's residual time. It is not an assumption about substitution between market and home-produced consumption; that margin is governed separately by $\mathcal{G}$. The maintained condition $\sigma \geq 1$ rules out strong complementarity among consumption and the two time margins: the household can substitute among them at least as readily as under Cobb–Douglas preferences. We use this restriction only for the corner response and the productive-capacity persistence result; the remaining results do not rely on it.

Empirical mapping. The Early Childhood rainfall exposure is a reduced-form proxy for a period-1 shock to non-labour income Y_1, while concurrent rainfall proxies for Y_2. This isolates the model's non-labour-income mechanism rather than asserting that rainfall affects households only through Y_t. The Early Childhood Plenty × concurrent-rainfall regression is the finite-treatment analogue of the cross-period comparative static.

6.3 Wage and Non-Labour-Income Shocks

Relative productivity alone does not settle whether the daughter performs additional domestic work. Her hour may produce less than an adult's, but it can still be the relevant source of domestic work when the adult's time has a high market or household opportunity cost. The first result distinguishes a change in the return to the mother's paid work from a

change in non-labour income at a fixed wage.

Proposition 1A (wage shocks). *Suppose the mother supplies paid work. Under the maintained CES allocation, holding Y fixed, an increase in her wage weakly raises the daughter's domestic work:*

$$\frac{\partial h_d^c}{\partial w} = \frac{h_o^c}{w}\left[\frac{Y}{F} + (\sigma - 1)s_C\varphi\right] \geq 0,$$

with strict inequality if $Y > 0$ or $\sigma > 1$. Here s_C is the consumption expenditure share and $\varphi \in (0, 1)$ is the exposure of the consumption price to the wage, both defined in Appendix C.1.

Proof. See Appendix C.2.

An increase in the mother's wage raises the opportunity cost of using her time in domestic production. In the maintained allocation, it also raises the implicit price of the daughter's residual time. The first term in brackets reflects the declining value of a fixed amount of non-labour income relative to the household's time endowment; the second is the substitution response.

Proposition 1B (non-labour-income shocks). *Hold w fixed and suppose consumption, the mother's leisure, and the daughter's residual time are normal goods. Once non-labour income is sufficiently high that the mother supplies no paid work, an increase in Y raises the daughter's domestic work in the maintained CES allocation when*

$$Y > \bar{Y} \equiv (1 + \gamma)w\frac{c}{1 - c},$$

under the maintained CES allocation and home-production demand condition discussed above.

Proof. See Appendix C.2 for the positive-paid-work case and Appendix C.3 for the corner condition.

While the mother supplies paid work, an increase in non-labour income has the familiar income effect. The relevant signs are

$$\frac{\partial h_m^p}{\partial Y} < 0, \qquad \frac{\partial h_d^p}{\partial Y} > 0, \qquad \frac{\partial h_l^p}{\partial Y} > 0, \qquad \frac{\partial h_d^c}{\partial Y} < 0.$$

The wage fixes the opportunity cost of adult time, and the household can accommodate a larger desired level of home production by moving the mother out of paid work. The daughter is consequently released toward her residual time. Once paid work is zero, this adjustment is unavailable. The threshold $\bar{Y}$ identifies that regime: non-labour income is high relative to the market value of household time, so the notional interior allocation would require negative paid work. It does not, by itself, determine the daughter's response at the corner.

At the corner, additional non-labour income raises the household's internal value of adult time. More materials generate final home consumption only when they are combined

with effective domestic work. Under the maintained conditions, the household therefore economises on adult time, and the daughter's domestic work rises. A less productive daughter can consequently supply additional domestic work when the household economises on adult time. The result is local to this corner allocation, rather than a prediction that favourable shocks universally raise girls' work.

6.4 Persistence and the Limits of Time-Use Evidence

The period-1 episode becomes persistent only if the additional childhood domestic work changes something that matters in adolescence. Let Z_1 denote any temporary period-1 shock—to non-labour income, the mother's wage, or both—that has no direct carryover to period-2 wages, non-labour income, or preferences, and affects those period-2 states only by changing the daughter's childhood domestic work. Write

$$\theta_Z \equiv \frac{\partial h^c_{d,1}}{\partial Z_1} > 0$$

for its childhood-engagement first stage. Any shock with $\theta_Z > 0$ can induce persistence through either channel. The laws of motion imply

$$\frac{\partial \log \gamma_2}{\partial Z_1} = A_\gamma \theta_Z, \qquad \frac{\partial(-\log a_{o,2})}{\partial Z_1} = A_\psi \theta_Z.$$

Productive capacity changes what the daughter's hour produces. A lower opportunity cost changes what the household gives up when it assigns that hour to domestic work. The distinction is central, but it is not visible from her later minutes alone.

Proposition 2A (persistence through two channels). *Suppose $\theta_Z > 0$ and consider an interior period-2 allocation. The daughter's later domestic work can rise through either productive capacity or a lower opportunity cost of domestic work:*

$$\frac{\partial h^c_{d,2}}{\partial Z_1} = \theta_Z \left(A_\gamma \Pi_\gamma + A_\psi \Pi_\psi \right),$$

where $\Pi_\psi > 0$ and $\Pi_\gamma > 0$ under the maintained substitutability condition.

Proof. See Appendix C.4.

The lower-opportunity-cost channel directly makes later domestic assignment less costly. Productive capacity has two opposing effects: it raises the effective output of the daughter's hour, but it also makes the household effectively richer, which tends to release her time. Under the maintained substitutability condition, the first effect dominates in the relevant period-2 comparative static. The coefficients Π_γ and Π_ψ collect these two period-2 responses.

Proposition 2B (local non-identification from time-use effects). *Let*

$$T \equiv (h_d^c, h_o^c, h_m^p, h_d^p, h_l^p)$$

denote the full vector of observed time-use margins. On a non-empty open set of parameter values, a productive-capacity-only economy and a lower-opportunity-cost-only economy, facing the same period-2 wage and non-labour income, can generate the same local effect vector $\partial T / \partial Z_1$. Consequently, neither the signs of the time-use effects nor any linear or nonlinear statistic constructed solely from that effect vector identifies whether persistence operates through γ or a_o.

Proof. See Appendix C.5.

This is stronger than the observation that both channels can raise the daughter's domestic work. Even the full pattern of local effects on the mother's paid work, domestic work, and leisure can be the same. The result concerns effects around a given allocation. It does not assert equality of time-use levels, finite changes, or unrestricted time-use distributions; those additional objects could impose further restrictions. Appendix C.5 gives the construction. The consumption-per-domestic-work result is useful precisely because it adds an output object that is absent from T.

6.5 A One-Sided Test for Productive Capacity

Time-use data reveal who works, but not what each worker's hour produces. Individual domestic output is rarely observed, so neither consumption alone nor domestic-work time alone separates the two channels. Let $H_d \equiv h_d^p + h_d^c$ denote household domestic-work hours, let $x \equiv h_d^c / H_d$ denote the daughter's share of those hours, and define monetary consumption as

$$X_C \equiv p_m C_m + p_r M_r. \tag{10}$$

It includes ready-purchased market consumption and the value of materials used in home production, including retained farm output when valued at p_r. At an optimum, $M_r = bD_t$. With constant returns in $\mathcal{G}$, the interior allocation implies

$$\frac{X_C}{H_d} = (p_m \rho + p_r) b \left[1 + (\gamma_2 - 1) x\right]. \tag{11}$$

where $\rho \equiv C_m / M_r$ is fixed at given relative prices. The multiplier is positive. The Appendix derives this identity and the signs used below.

Proposition 3 (one-sided test using monetary consumption). *Consider an interior period-2 allocation in which early engagement has increased the daughter's later domestic work, $\gamma_2 \leq 1$, and the maintained constant-returns consumption aggregator applies. (a) If the childhood episode affects*

only the opportunity cost of assigning her domestic work ($A_Y = 0$), then

$$\frac{\partial (X_C / H_d)}{\partial Y_1} \leq 0.$$

(b) Under the same conditions, it cannot make the response of X_C / H_d to current non-labour income more positive:

$$\frac{\partial^2 (X_C / H_d)}{\partial Y_1 \partial Y_2} \leq 0.$$

Thus, a positive level response or a positive interaction implies $A_Y > 0$. The converse does not hold.

Proof. See Appendix C.6.

The test is deliberately one-sided. Productive capacity may be present without generating a positive empirical ratio response because the household can reallocate domestic work at the same time. A positive level response or positive interaction with current non-labour income nevertheless rules out a pure lower-opportunity-cost account within the maintained model.

Section 7 implements the proposition using log reported household consumption per minute of core domestic work. The constant-returns mapping from the model's monetary-consumption object to this ratio is assessed there.

6.6 Implications for the Mechanism Evidence

The model gives an order to the mechanism evidence. We first ask whether favourable rainfall raises household domestic-production demand alongside output and consumption, rather than only relaxing the budget constraint. This is the household response required for a childhood shock to increase the daughter's actual engagement in domestic work.

We then ask whether the persistent response is concentrated on the daughter who was most likely to perform that additional work. The eldest-daughter comparison and the developmental timing results are evidence on this engagement margin. They do not, on their own, establish what the engagement produced.

Finally, we combine household monetary consumption with core domestic-work time. Proposition 3 implies that a positive response of consumption per unit of core domestic work, or a more positive response to a current non-labour-income shock, supports a productive-capacity component. A null or negative response remains consistent with either channel. The next section follows this sequence. Appendix C gives the proofs.

7 Mechanism

Sections 4 and 4.2 show that favourable rainfall in the pre-school years is associated with more domestic work by girls many years later, with the additional work displacing studying on school days, and leisure on others. The mechanism must therefore explain why a temporary resource shock draws a young girl into household production, and how it can have a persistent effect a decade later. Our central interpretation is learning-by-doing: favourable rainfall can

expand the household's production and consumption activity, increasing the domestic work required at the time; early engagement in that work may then build domestic-production skill. We do not observe task engagement in early childhood or individual domestic output. Instead, we assess whether the evidence has the pattern implied by this interpretation: whether rainfall plenty explains more of the persistent effect than episodes of low rainfall in early childhood; whether the response is concentrated on eldest daughters, who are most likely to absorb additional domestic work; and whether later household consumption is higher per unit of core domestic-work time. We also consider explanations based on later changes in family composition, schooling and health, and local economic conditions. Taken together, the evidence is consistent with learning-by-doing and offers little support for those alternative accounts, while not establishing every channel through which early exposure may persist.

7.1 Agricultural Surplus and Domestic-Labour Demand

7.1.1 Rainfall and Home Production

Favourable rainfall can raise agricultural production, household income, and the returns to household members' work. The mechanism studied here begins with a further consequence of agricultural production in rural households: part of consumption is produced at home. Household consumption and the allocation of domestic labour cannot then be considered independently. A positive shock may relax the cash constraint, but it can also increase the materials and domestic work required to produce the final goods the household consumes.

The distinction depends on the form in which additional resources are used. If they finance goods that are ready to consume, higher income can release household time. But additional home-grown crops, livestock products, and other raw household inputs must be cooked, processed, cleaned, and combined with domestic work before they become final consumption. The same rainfall shock can therefore raise purchasing power while also increasing the household's requirement for domestic labour. The relevant question is not whether higher income always raises domestic work. It is whether, where home production is substantial, an expansion of agricultural resources and consumption is accompanied by an expansion of the domestic labour used to turn those resources into final consumption.

The first empirical exercise asks whether rainfall shocks in the survey period generate the household response the proposed mechanism requires. Concurrent rainfall is defined as the average of the four standardised quarterly rainfall shocks in the twelve months ending in the survey quarter; lagged rainfall is the analogous average over the preceding four quarters. This construction attaches to each diary a rainfall shock measured over a comparable twelve-month horizon, while district-by-quarter fixed effects absorb seasonality in rainfall and time use, so the estimates compare rainfall variation over the same seasonal horizon. The exercise is a contemporaneous analogue of the proposed historical channel, not the long-run estimate itself. It asks whether, when rainfall raises resources, household domestic production also

expands.

Household Responses to Current Rainfall Tables 17 and 3 show the required household expansion. A one-standard-deviation increase in concurrent rainfall raises agricultural production by 30.2 thousand tonnes and monthly household consumption by Rs 218. It also raises the home-grown share of consumption by 0.5 percentage points and total household domestic work by 10.3 minutes per day. Lagged rainfall raises total household domestic work by 12.9 minutes, including 6.8 minutes of food preparation and 3.9 minutes of cleaning; concurrent rainfall also raises domestic production for own use by 7.5 minutes. Thus, the rainfall response is not limited to resources measured in income or consumption expenditure: it includes more home-grown consumption and more domestic labour within the household.

TABLE 3: *Household Response to Rainfall: Consumption and Time Use*

	Household Consumption *(Rs per Month)*		**Household Total Domestic Work** *(mins per day)*				**HH Paid Work** *(mins per day)*
	Consumption *(log)*	**Home-grown** *(share)*	**Total**	**Food Prep**	**Cleaning**	**Domestic Production**	**Total**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**	**(7)**
Concurrent Rainfall	0.021***	0.005***	10.3**	2.8	2.2*	7.5**	-3.5
	(0.007)	(0.002)	(4.2)	(1.8)	(1.2)	(3.0)	(4.9)
Lagged Rainfall	0.013*	0.004**	12.9***	6.8***	3.9***	2.8	-4.1
	(0.007)	(0.002)	(4.2)	(1.8)	(1.2)	(4.1)	(5.6)
Observations	61,547	61,547	61,547	61,547	61,547	61,547	61,547
Outcome Mean	9.044	0.114	664.0	294.7	136.6	125.6	564.3
District FE	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y

Notes: **Sample.** TUS 2019 + 2024 pooled, rural households with an identified eldest child aged 17 or younger and uniquely identifiable parent–child relationships, one observation per household. **What it shows.** Eight columns in two groups. *Consumption (cols 1–3):* total monthly household consumption in rupees (col 1) and in logs (col 2), and the home-grown share of consumption (col 3: home-grown consumption, missing recoded to zero, divided by total consumption, in $[0, 1]$). *Time use (cols 4–8):* household-total domestic work in minutes per day (col 4) and in logs (col 5), household-total paid work in minutes per day (col 6) and in logs (col 7), and the household-enterprise share of paid work (col 8). Log and share columns condition on positive values: 0.06% of households drop for domestic work and 10.96% for paid work, and the level column beside each log keeps them. **Household totals.** Each diary respondent's minutes, winsorised at the state ×gender 99th percentile, summed over all diary respondents in the household; consumption is recorded at the household level and enters directly, not winsorised. **Concurrent and lagged rainfall.** Quarter-aligned rolling twelve-month averages of standardised district rainfall — concurrent covers the twelve months ending in the diary quarter, lagged the twelve months before — each re-standardised within the estimation sample. **Estimation.** District-quarter and survey-year fixed effects; coefficients are per one-standard-deviation increase. **Inference.** Standard errors clustered at district ×survey-year, shown in parentheses; $***p<0.01$, $**p<0.05$, $*p<0.10$.

Unlike expenditure, the Household Consumption Expenditure Survey (HCES) measures monthly quantities of food items consumed, so an expenditure increase driven only by prices need not be read as an increase in food consumption. It also distinguishes ingredients from prepared food. In rural households, concurrent rainfall raises the quantities of food ingredients and the broad staples basket, while purchases of packed or served food and meals eaten away from home fall (Table 18). HCES is an earlier, separate survey and does not observe who prepares the food. It nevertheless shows that rainfall is associated with a shift

toward food inputs more likely to require household preparation, rather than simply higher expenditure on ready-to-eat food.

We use the rural–urban comparison as a scope check on the proposed agricultural home-production channel. The Early Childhood Shock raises domestic work among rural girls, while the rural–urban differential offsets that relationship and leaves the implied urban coefficient indistinguishable from zero (Table A10). Rural and urban households differ in many ways, so this does not identify a particular rural technology. It is nevertheless consistent with an account centred on agricultural home production rather than a generic consequence of rainfall.

One explanation for additional domestic work by girls is that favourable rainfall draws adult women into agricultural or paid work, leaving girls to meet a domestic-work requirement. Table A11 does not support that particular account. The paid-work coefficients for adult women and mothers are small and imprecise at both the concurrent and lagged rainfall horizons, and the table does not separately identify agricultural work.

The detectable adult response is instead an increase in domestic work following lagged rainfall: 8.3 minutes per day for adult women and 8.1 for mothers. Daughters' domestic work rises by 4.2 minutes and their studying falls by 8.8 minutes in response to concurrent rainfall. These estimates need not describe a single synchronous within-household adjustment. They nevertheless show that the proposed mechanism need not rely on rainfall moving adult women out of domestic work: favourable rainfall is also associated with more domestic work by adult women, alongside the daughter's time-use response.

7.1.2 Surplus Rather than Scarcity

The home-production reading has a directional implication. If the mechanism is scarcity or distress, the lower tail of the rainfall distribution should matter: drought years should push girls toward household work as families cope with lower income. If the mechanism is surplus home production, the upper tail should matter instead: unusually wet years should raise agricultural output and home-grown consumption, increasing the domestic labour needed to transform those resources into meals and other final consumption goods. To distinguish these cases, we split each four-year life-stage rainfall shock into two indicators. *Plenty* equals one when the standardised rainfall shock is at least +1.15 SD, and *Drought* equals one when it is at most −1.15 SD; the middle of the distribution is the omitted category. The same split is applied separately to the Infancy and Early Childhood windows.

Evidence from the Rainfall Tails Table 19 reports an upper-tail pattern. Early-childhood plenty raises girls' adolescent domestic work by 25.1 minutes per day and lowers studying by 37.1 minutes, both significant at the 1% level. Early-childhood drought, by contrast, has no detectable effect on domestic work and only a small, statistically insignificant effect on studying. Drought may reduce the quality or quantity of consumption, and it may damage health, but it need not reduce the household's minimum cooking, cleaning, and

care requirements. By contrast, plenty can expand the stock of home-grown output and consumption to be processed, so it can create additional domestic-production pressure. The coefficient pattern is therefore concentrated in the part of the rainfall distribution where the home-production reading implies additional work, rather than the part where a distress interpretation would naturally load.

The infancy window behaves differently. Infancy plenty does not reproduce the early-childhood pattern, while infancy drought reduces girls' later domestic work by 11.6 minutes per day. The equality tests reject common Infancy and Early Childhood coefficients within both tails for domestic work, reinforcing that the timing of the rainfall matters. The infancy pattern is consistent with a distinct deprivation or health channel, as the health evidence below also suggests; the time-use data alone do not identify that channel. The pre-school pattern, by contrast, is concentrated in unusually wet years, as the surplus interpretation predicts. This opposite-tail pattern prevents the mechanism section from treating all childhood rainfall as one object and ties the persistent domestic-work response specifically to pre-school plenty.

A flexible functional-form check reaches the same conclusion without dividing the shocks at an arbitrary threshold. Figure A10 estimates domestic work as a cubic function of the two childhood shocks. Over the range of the data with substantial support, the predicted response of eldest daughters rises as the Early Childhood Shock becomes more favourable; it does not show a corresponding increase at the drought end of the distribution. The flexible estimates therefore support the tail comparison's central distinction: the persistent response is associated with rainfall plenty in the pre-school window, rather than with low rainfall.

The asymmetric-tail quantile estimates add a further distinction between the two positive rainfall associations. Following the Canay (2011) approach described in Section 4.3, Figure A11 compares Early-Childhood Plenty with Infancy Not-Drought among girls with positive domestic work. The Early-Childhood-Plenty effect grows towards the upper part of the domestic-work distribution. By contrast, the Infancy-Not-Drought effect is positive only at low quantiles and falls to zero or below at higher quantiles. The figure does not itself identify the source of either response, but it shows that the two exposures do not have the same distributional signature: pre-school plenty is associated with heavier later domestic-work engagement, whereas the infancy association is confined to girls doing relatively little domestic work.

These results establish the production environment required by the proposed account: favourable rainfall can expand agricultural resources, home-grown consumption, and household domestic labour together. They do not show that girls exposed in early childhood performed the additional work themselves. Section 7.2 takes up that engagement question.

7.2 Engagement and Learning-by-Doing

If favourable rainfall expands household production, the resulting domestic work need not be divided equally across daughters. When the eldest daughter is a child, there is, on average,

no older female child available to help adult women meet that requirement. She is also more likely to have younger siblings, whose care and associated household needs can add to the work the household must perform. These features make her more likely to be recruited when domestic work demand expands.

The comparison with younger daughters follows from the same logic. A younger daughter has at least one older girl in the household, who is generally further advanced physically and can be drawn into domestic work first. The eldest-daughter analysis therefore does not treat birth order as an intrinsically causal trait. It asks whether the persistent rainfall response is concentrated on the child role that, at the time of a household demand shock, is most likely to supply additional domestic work.

The corresponding short-run evidence shows that the same role responds when current household conditions change. In the family-identified eldest-versus-younger frame, a one-standard-deviation increase in concurrent rainfall raises an eldest daughter's core domestic work by 5.6 minutes per day; the younger-sister differential is negative, leaving their implied response close to zero (Appendix Table 20). This shows that contemporaneous rainfall draws eldest daughters into core domestic work, the same margin on which the childhood response later appears.

7.2.1 Sibling-Order Evidence

The survey does not record parent–child links directly, which is particularly important in Indian joint-family households where several nuclear families may share a roster. To reduce the risk of assigning a child to the wrong parental family, we retain her only when at least one parent can be uniquely identified and the other is either uniquely identified or absent. Daughter rank is then defined within that identified parental family, rather than from the ordering of all children in the household. The roster-linking procedure uses relationship-to-head and marital-status information; households with multiple candidates for either observed parent are dropped.[15] Among rural children aged 17 or younger pooled across rounds ($N = 111{,}103$), the mother is uniquely identified for 93.5% and joint identification succeeds for 95.2%.

Identifying Daughter Rank Figure 9 documents the allocation that motivates this comparison. Within each displayed age group and sibling configuration, eldest daughters perform substantially more domestic work than second daughters, while the corresponding studying differences are much smaller. The figure is descriptive: it neither identifies a rainfall effect nor explains why households allocate work by daughter rank. Its purpose is to show that eldest-daughter status corresponds to a large, observable difference in the domestic work that girls perform before turning to the rainfall response.

[15] If the child is the head's own child the candidates are the head and spouse; if the head's grandchild, the head's married children of the appropriate gender; and analogously for the head's siblings. Cases with two or more candidate mothers or candidate fathers are dropped.

Rainfall Responses by Daughter Rank With the family-identified sample in hand, Table 4 turns the descriptive rank difference into a rainfall-response test. The question is not whether eldest daughters perform more domestic work on average. It is whether the effect of childhood rainfall is concentrated on the daughter role that is most likely to be recruited when household domestic-work demand expands. The table addresses that question using two complementary designs.

The first three columns use the panel-by-eldest design. They do not compare the levels of domestic work performed by an eldest daughter and a younger daughter. Instead, they estimate the effect of childhood rainfall separately within each daughter category, using changes across the two survey waves in the same district–age–quarter cells and allowing separate cell and survey-year fixed effects for eldest and younger daughters. The base shock coefficient is therefore the estimated effect among eldest daughters, and the $\times$ Younger coefficient is the difference between that effect and the separately estimated effect among younger daughters. Column (1) gives the unweighted estimate: the Early Childhood coefficient is 5.97 minutes per day for eldest daughters, while the younger-daughter differential is -7.80, implying an effect of -1.83 minutes for younger sisters (SE 4.11; $p = 0.66$).

Columns (2) and (3) address one specific composition issue: eldest and younger daughters are observed at different age distributions and therefore place different implicit weights on the district–age comparisons. Column (2) reweights younger daughters to match the eldest daughters' age distribution within ages 12–17; column (3) repeats the exercise after extending the sample to ages 10–17. The eldest coefficient remains between 5.43 and 5.97 minutes per day, while the younger-daughter differential remains negative in each specification.

Columns (4) and (5) use a different, cross-sectional comparison. Within district–age–quarter–rank cells, girls with the same overall birth rank can differ in whether they are the firstborn girl because their older siblings are boys rather than girls. The design therefore compares girls whose eldest-daughter status differs because of older-sibling gender composition, rather than because one girl is earlier in the overall birth order. Column (4) returns an eldest-daughter coefficient of 5.96 minutes. Column (5) further saturates the rainfall shock within district–age–quarter–survey-year–rank cells, so it identifies only the eldest–younger differential; that differential remains negative at -7.71. Across the designs, the eldest–younger differential is negative but less precisely estimated than the positive eldest-daughter response, and the implied younger-sister coefficient is statistically indistinguishable from zero in the panel estimates.

The null domestic-work response among younger daughters is not, by itself, decisive. Younger daughters may simply remain outside the household role through which additional domestic work is assigned, even if early rainfall had some separate effect on their time allocation. The day-type comparison speaks to that possibility. Table 21 shows the familiar pattern for eldest daughters: less studying on ordinary weekdays, when school time is scarce, and recovery on

Sundays, with leisure adjusting in the opposite direction. For younger daughters, there is no corresponding weekday loss or Sunday recovery in studying, and no comparable day-type pattern in leisure. The point is not that the table precisely estimates every eldest–younger difference. Rather, the entire study–leisure adjustment is absent for the daughters who do not show the domestic-work response. This makes a direct effect of early rainfall on schooling or leisure, independent of the household's domestic-work requirement, less plausible.

The corresponding exercise for boys provides a useful limit on the rank interpretation. After reweighting younger sons to the age distribution of eldest sons, the Early Childhood effect on domestic work is 11.52 minutes per day lower for younger sons than for eldest sons (Appendix Table 22). There is no statistically distinguishable eldest–younger differential in boys' paid work or studying under either weighting. This is not a direct test of why household work is allocated by rank, but it indicates that the rank contrast is not a general shift in boys' time allocation across all activities.

Composition and Household Structure Daughter rank and sibling composition are not randomly assigned. Selective fertility, survival, and household composition could make eldest and younger daughters come from systematically different families. In addition, an older sister may have left the household, so that an observed eldest daughter was not the oldest girl in her full sibship. Absent selective departure related to the rainfall response, this places a younger-role child in the eldest category and attenuates the rank contrast. The younger-mother columns in Table 23 narrow the scope for this concern, while the low-son-preference columns address one source of endogenous sibling composition.

We address these composition concerns by re-estimating the rank comparison among daughters of younger mothers, where an older sister is less likely to have left the household, and in low-son-preference districts, where endogenous sibling composition is less likely to drive the comparison. The eldest-daughter coefficient remains positive and significant in every specification, between 4.15 and 8.52 minutes per day (Table 23). The younger-sister differentials remain negative but are imprecise in three of four columns. These restrictions reinforce the positive eldest-daughter response; they do not warrant a stronger conclusion about the size of the younger-sister response.

Eldest daughters are more likely than younger sisters to come from households with only one daughter at adolescent ages and from smaller households. We therefore re-estimate the rank comparison after absorbing fixed effects for the household's sibling-gender configuration, restricting to households with at least two girls aged 12–17, and dropping every household containing a twin in that age range (Appendix Table A12). The positive response for eldest daughters and the negative younger-sister differential remain under each check; the estimates are less precise in the two-girl sample, as expected from its smaller size.

The same cubic figure adds a rank comparison to the tail results above. Over the part of the Early Childhood rainfall distribution with substantial support, the predicted domestic-work response rises for eldest daughters as the shock becomes more favourable, whereas the

TABLE 4: *Eldest Daughter Versus Younger Sisters: Fixed-Effect Build-Up*

	Panel ×Eldest				Cross-sectional
	Unweighted	Reweighted	Reweighted Ages 10–17	Panel ×Rank	Saturated
	(1)	(2)	(3)	(4)	(5)
Infancy Rain Shock	1.22	1.22	1.62	0.82	
	(1.62)	(1.62)	(1.45)	(1.65)	
× *Younger*	-1.87	-2.30	-0.67	-2.33	1.27
	(3.52)	(4.29)	(3.58)	(3.42)	(4.40)
Early Childhood Rain Shock	5.97***	5.97***	5.43***	5.96***	
	(1.91)	(1.91)	(1.65)	(1.92)	
× *Younger*	-7.80*	-8.97*	-8.25**	-5.29	-7.71*
	(4.44)	(5.32)	(3.74)	(3.53)	(4.13)
Observations	26,781	26,781	34,935	26,781	26,781
Eldest Mean	98.4	98.4	85.6	99.7	101.8
Younger Mean	43.6	55.1	44.4	45.7	46.3
Younger share in sample	0.16	0.16	0.20	0.16	0.16
[Eldest] p: Infancy =Early Ch.	[0.02]	[0.02]	[0.03]	[0.02]	
[Younger] p: Infancy =Early Ch.	[0.80]	[0.74]	[0.34]	[0.64]	[0.17]
District-Age-Qtr-Eldest FE	Y	Y	Y	.	.
Year-Eldest FE	Y	Y	Y	.	.
Overall Sibling Rank FE	Y	Y	Y	.	.
District-Age-Qtr-Rank FE	.	.	.	Y	.
Year-Rank FE	.	.	.	Y	.
District-Age-Qtr-Year-Rank FE	.	.	.	.	Y

Notes: **Sample.** Rural girls, TUS 2019 + 2024 pooled, the main 26 states, family-identified. Cols (1)–(2) and (4)–(5) cover ages 12–17; col (3) extends to ages 10–17. The table compares eldest daughters to younger daughters within this pool. **What it shows.** Two identification strategies. Cols (1)–(3): a panel-by-eldest design with district-age-quarter-eldest + survey-year-eldest fixed effects and an overall-rank control. Cols (4)–(5): a cross-sectional rank design that identifies the contrast off cross-cohort variation in older-sibling gender composition; col (5) fully saturates the rainfall shock at the cell level. **Eldest daughter.** The household's firstborn girl (gender-rank = 1 among children aged 0–17), even with older brothers present. **Estimation.** Each column interacts both shocks with a younger-daughter indicator; reported coefficients are the effect on eldest girls and the ×Younger row is the younger-daughter differential. **Inference.** Standard errors clustered at district ×age ×survey-year, shown in parentheses; $***p<0.01$, $**p<0.05$, $*p<0.10$.

younger-sister curve does not show the same systematic increase. The confidence intervals widen markedly at the extremes.

The joint-family comparison provides related but non-conclusive context. The Early Childhood response is substantially smaller in joint-family households: the 6.7-minute response in other households is offset by a -5.8-minute joint-family differential (Table A4). This pattern is consistent with co-resident adults sharing domestic work and reducing the need to recruit an eldest daughter. However, girls in joint-family households also perform somewhat less domestic work on average at the time of observation (81 rather than 92 minutes per day). The attenuation may therefore reflect lower demand for girls' domestic work as well as a larger supply of potential helpers. It supports, but does not establish, the recruitment interpretation.

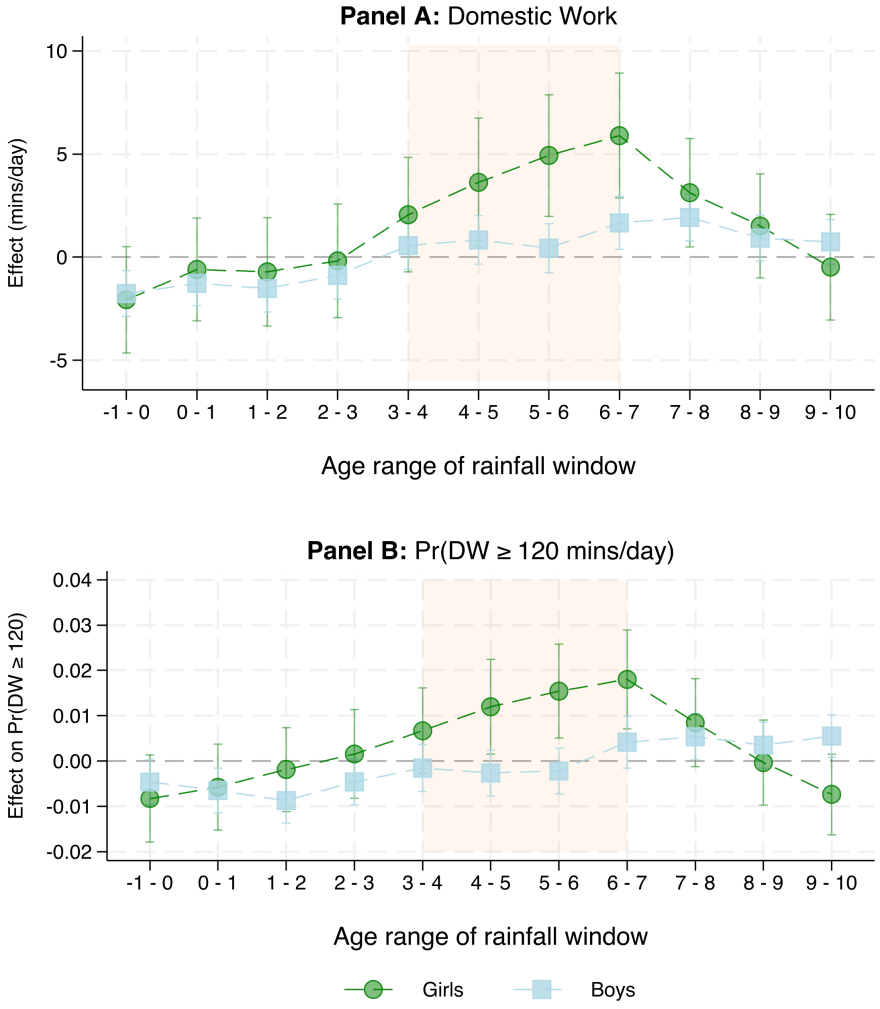
Across the panel and cross-sectional rank designs, the positive Early Childhood response is found for eldest daughters, while the implied response for younger sisters is close to zero and the associated weekday–Sunday adjustment in studying and leisure is absent. The result survives restrictions designed to reduce older-sister attrition and endogenous sibling composition, and is attenuated where co-resident adults may share domestic tasks. Taken together, these patterns are consistent with a household demand shock being absorbed principally by the eldest daughter, rather than with an effect of early rainfall on girls' time use that is independent of their role in household production. They therefore suggest that the daughter's engagement in domestic tasks during childhood, rather than exposure to rainfall alone, is relevant to the later response.

7.2.2 Developmental Timing

Having identified the child role most likely to absorb the additional work, we next ask whether the timing of the rainfall response is consistent with engagement beginning when girls first become able to contribute to domestic tasks. We use rolling two-year exposure windows to ask whether the persistent association is largest at those ages. Figure 2 replaces the fixed four-year shock with those windows and plots the estimated domestic-work coefficient against age at exposure. The profile is small and statistically indistinguishable from zero in infancy, rises through ages four to seven, and returns toward zero by the early school years. Adjacent windows have statistically indistinguishable coefficients, so the four-to-seven definition summarises a smooth profile rather than identifying a discrete developmental threshold. The pattern is consistent with early domestic-task engagement being relevant during the pre-school years; it does not measure that engagement directly.

Figure 3 asks when the later association is visible. Holding the exposure window fixed at ages four to seven, the domestic-work coefficients are positive from the earliest observed outcome window, ages 8–9, through ages 16–17. They are imprecisely estimated at every age, and their size varies across the rolling windows, so the figure does not establish a monotonic age gradient. Its more limited implication is that the association does not first appear only in

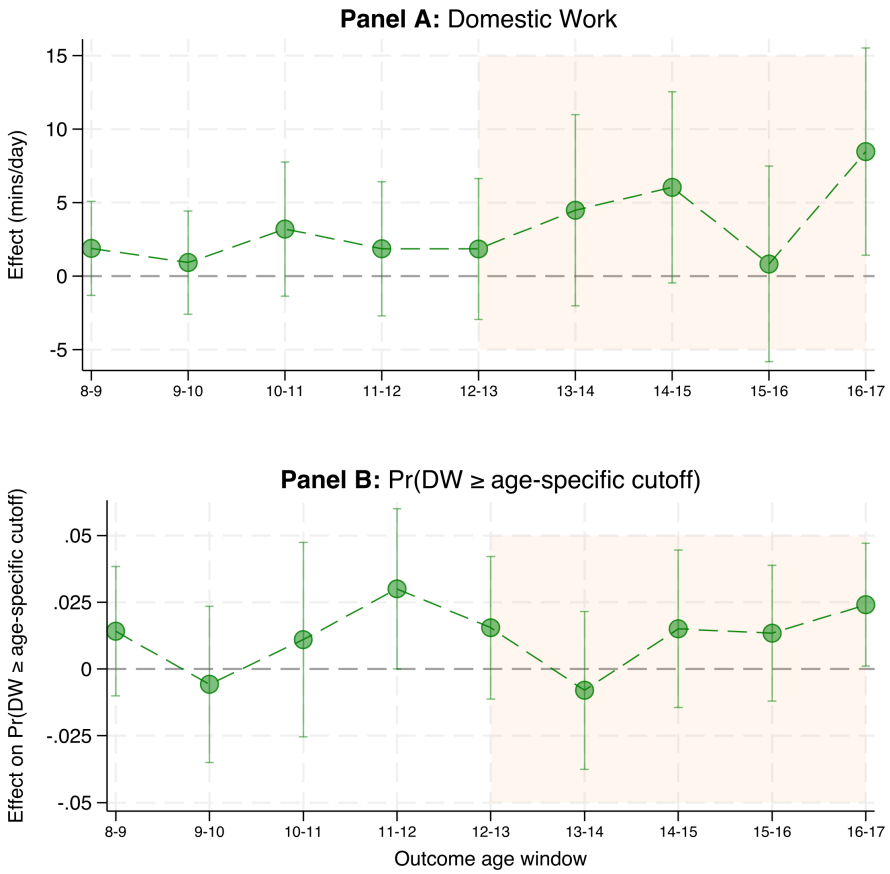
FIGURE 2: *Which Years of Childhood: Rolling-Window Age Profile*



Note: **Sample.** Rural children aged 12–17, TUS 2019 + 2024 pooled, both genders, the main 25 states (West Bengal excluded; qualitative pattern unchanged when West Bengal is included). **What is plotted.** Each marker is a separate regression where the only rainfall variable is the mean district rainfall z-score over a single rolling two-year window; the horizontal axis tracks the age window for the rainfall shock. Shaded band marks ages 4–7. Panel A: domestic work in minutes per day. Panel B: indicator for ≥120 minutes per day. **Rolling rainfall shock.** Mean district rainfall z-score over the two-year span, re-standardised within the estimation sample. **Estimation.** Each regression interacts the rolling window with gender (markers are the effect on girls); headline fixed-effect structure (district-age-quarter-gender + survey-year-gender); clustered at district ×age ×survey-year.

late adolescence. It is consistent with exposed eldest daughters performing more domestic work through childhood and adolescence, rather than with a one-off response that emerges many years after the exposure period.

FIGURE 3: *Early Childhood Rainfall and Eldest Daughters' Domestic Work across Adolescence*



Note: **Sample.** Rural eldest daughters aged 8–17 with completed time diaries, TUS 2019 + 2024 pooled, the main 26 states, family-identified; age range extended below the headline 12–17 sample so the full adolescent profile is visible. **What is plotted.** Each marker is the Early Childhood Shock coefficient from a separate regression on a rolling outcome-age window. Shaded band marks ages 12–17. Panel A reports domestic-work minutes per day. Panel B reports an age-specific participation indicator: 30 minutes per day for windows centred at ages 8.5–13.5 and 60 minutes per day for windows centred at ages 14.5–16.5. **Estimation and inference.** Each regression estimates girls' domestic work on the gender-interacted shocks within the outcome-age band, with district ×age ×quarter and survey-year fixed effects. Standard errors are clustered at district ×age ×survey-year.

Taken together, the two profiles connect the pre-school rainfall window to later domestic work in a way consistent with early engagement: the relevant exposure occurs at ages when children begin contributing to household tasks (Whiting and Edwards, 1988; Lancy, 2015), and the positive association is already visible at the first observed outcome ages. The data are repeated cross-sections rather than a record of each girl's childhood task history. They therefore cannot distinguish cumulative practice from earlier entry into household tasks, or from a household-demand shock that changes the timing of that entry. Nor do they distinguish domestic-work skill from habit formation or another persistent change induced by early engagement.

One possibility is that favourable rainfall leads girls to enter domestic tasks earlier and accumulate more task-specific practice over subsequent years. Another is that the same additional exposure is especially consequential at these ages, when motor coordination, task sequencing, and familiarity with household routines are developing. Both readings imply that a temporary expansion of household work can have lasting consequences when it occurs as girls begin undertaking those tasks.

The rank and timing evidence together establish the empirical premise for the next test: the response is concentrated on the child most likely to have absorbed additional work, and in the period when such engagement could plausibly begin. The next subsection asks whether the later household-production pattern provides evidence that this engagement built domestic-work-specific skill.

7.3 Productive Capacity versus Lower Opportunity Cost

The time-use results alone leave an important ambiguity. Early engagement may lower the opportunity cost of assigning the daughter to domestic work—through habit formation, lower returns to schooling, lower disutility of those tasks, or another change in the value of her alternative time—and each of these channels can raise her later domestic-work minutes. They do not, on their own, raise the household's monetary consumption for a given amount of domestic work. Domestic-work-specific skill can. Proposition 3 therefore gives a one-sided test: a positive response of household monetary consumption per unit of domestic work, either in the standing allocation or when current rainfall again improves, requires a productive-capacity component within the maintained model. A zero or negative response does not distinguish the channels.

The preceding evidence lets us apply that test at the household level. It localises the enduring response in early-childhood rainfall plenty and on the eldest daughter, the child most likely to have been recruited into the additional domestic work. We use her childhood rainfall history as a shock to the household's later allocation, without attributing household consumption or every member's activity to her. The question is narrower: do otherwise comparable households with an eldest daughter exposed to early-childhood plenty obtain more reported consumption from their domestic work, and does that relationship become stronger when current rainfall improves? This combines household consumption and time totals because individual domestic output is not observed.

7.3.1 Mapping the Model to Household Data

The model's object is monetary consumption,

$$X_C = p_m C_m + p_r M_r,$$

per unit of household domestic work, H_d. The diary records detailed domestic-work time for every responding household member, so the household denominator is directly observed.

We use *core domestic work*—food preparation and cleaning—as the preferred denominator because it is most directly involved in converting household inputs into final consumption; the all-domestic-work ratio checks that the result does not depend on that choice.

Reported monthly household consumption is the empirical counterpart of X_C. The required mapping is the constant-returns assumption on $\mathcal{G}$: at given relative prices, it keeps the ratio of ready-purchased market consumption to materials used at home fixed. This is not a claim that the survey separately observes those two components, nor can the data prove exact constant returns. Its substantive implication for the present test is that a rainfall-related rise in reported consumption should not be driven by a switch toward ready-to-consume market goods that require little domestic work.

The available evidence is consistent with that implication. In the HCES, an independent survey that records monthly food quantities, favourable current rainfall raises ingredients and staples while reducing purchases of prepared food and meals eaten away from home (Table 18). In the TUS, concurrent rainfall raises both total consumption and the home-grown share of consumption (Table 3, columns 1–2); the latter measures the source of ingredients rather than final home-produced goods, but it provides related evidence on the materials margin. Thus, while the data do not test constant returns directly, they do not suggest the compositional shift toward prepared market consumption that would undermine this use of reported consumption.

7.3.2 Empirical Approach and Results

The household regressions preserve the source of variation in the baseline design. We retain rural, family-identified single-family households with an eldest daughter aged 12–17 and keep one observation per household. Her age defines the district × age × quarter fixed effects, with survey-year fixed effects. Thus, households with focal eldest daughters of the same age are compared within district and quarter across the two TUS waves; only the outcome has been aggregated to the household level. The main outcome is log monthly consumption per minute of household core domestic work. The specification includes Early-Childhood Plenty, Infancy Not-Drought, concurrent rainfall, and each childhood-history interaction with concurrent rainfall. Tail coding retains power in this household sample while using the two histories that the earlier analysis showed have different empirical signatures.

The main econometric complication is that an eldest daughter's rainfall history may be correlated with that of an age-eligible eldest son, whose own history could affect household consumption and time use. Only 35 percent of the focal-daughter households contain such a son aged 12–17. In the preferred specification, his continuous infancy and early-childhood shocks are re-standardised within the estimation frame, set to zero when no eligible son is present, and accompanied by an absent-son indicator, son-age controls, and, in the concurrent specification, an absent-son-by-current-rainfall interaction. This retains the full focal-daughter sample without treating the absence of a son as a zero shock. It is an

adjustment for correlated household exposure histories, not a comparison of daughters with sons.

At average concurrent rainfall, Early-Childhood Plenty raises log monthly consumption per minute of core domestic work by 0.040, about 4 percent, significant at the ten-percent level (Table 5, column 2). Infancy Not-Drought instead lowers the ratio by 0.077, significant at the five-percent level; the difference between the two histories is significant at the one-percent level. This result already rejects a simple reading in which any relatively favourable early rainfall history produces the same later household allocation.

The current-rainfall interaction gives the stronger test. A one-standard-deviation increase in current rainfall adds 0.058 log points to the consumption-per-core-domestic-work ratio of households with an Early-Childhood-Plenty eldest daughter, significant at the five-percent level. Their total difference is therefore about 0.098 log points, or 10 percent, when current rainfall is one standard deviation above its mean; it is also significantly different from the corresponding infancy response. The analogous infancy interaction is small and imprecise. This is the model's comparative static: when household resources improve again, the households carrying the relevant childhood history obtain more consumption from a given amount of domestic work. It is difficult to reconcile that pattern with an account in which early rainfall did no more than make the adolescent daughter's domestic time cheaper to assign.

The headline result is not tied to the use of core domestic work as the denominator. When all domestic work is used instead, the historical ratio is less precisely estimated but the Early-Childhood-Plenty interaction with current rainfall is 0.073 and significant at the one-percent level. Columns 5–6 then show what drives the core-domestic-work result: Early-Childhood Plenty raises log consumption by 0.052, while its estimated effect on core domestic-work minutes is positive but small (0.012) and imprecise. The corresponding current-rainfall interaction is 0.080 for consumption and 0.022 for core domestic work, again imprecisely estimated. The ratio result is thus not produced by a mechanical fall in domestic-work time. It occurs alongside a rise in consumption and no precisely estimated contraction in the domestic-work input.

7.3.3 Robustness to Correlated Sibling Histories

Appendix Table A14 directly addresses the concern that the focal daughter's coefficient might instead reflect the correlated childhood history of an eligible eldest son. It does not. The Early-Childhood-Plenty interaction is positive in every specification, and the preferred estimate changes little when we either remove the need to impute a son's history or change the way that history enters the regression.

The first two columns set out the cleanest checks. Column 1 omits son histories altogether in the full focal-daughter sample; column 2 restricts the sample to the 12,982 households without an eligible eldest son, where this source of correlated exposure does not arise

TABLE 5: *Household Consumption per Minute of Domestic Work: Childhood and Concurrent Rainfall*

	log(Monthly Consumption (Rs) per minute of DW)						
	Core DW only		All DW		Parts		
	Historical	+Current	Historical	+Current	log**C**	log**Core DW**	Home-grown *(share)*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Infancy Not-Drought	-0.077**	-0.077**	-0.062*	-0.062*	-0.038*	0.040*	0.013**
	(0.030)	(0.030)	(0.032)	(0.032)	(0.022)	(0.023)	(0.005)
Early Childhood Plenty	0.042*	0.040*	0.012	0.011	0.052***	0.012	0.012**
	(0.022)	(0.022)	(0.025)	(0.024)	(0.019)	(0.017)	(0.005)
Concurrent		0.020		0.025	0.027	0.007	0.009
		(0.029)		(0.030)	(0.021)	(0.019)	(0.006)
Infancy ×Concurrent		-0.020		-0.022	-0.013	0.008	-0.002
		(0.029)		(0.030)	(0.021)	(0.020)	(0.006)
Early Childhood ×Concurrent		0.058**		0.073***	0.080***	0.022	-0.002
		(0.027)		(0.028)	(0.023)	(0.022)	(0.006)
[Concurr. =0] p: Infancy =Early Ch.	[0.00]	[0.00]	[0.07]	[0.07]	[0.00]	[0.33]	[0.82]
[Concurr. =1] p: Infancy =Early Ch.		[0.00]		[0.00]	[0.00]	[0.74]	[0.84]
Observations	20,471	20,471	20,486	20,486	20,471	20,471	20,471
Outcome Mean	2.983	2.983	2.653	2.653	9.056	6.072	0.116
District-Age-Qtr FEs	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y
Eldest Boy Shocks (imputed)	Y	Y	Y	Y	Y	Y	Y

Notes: **Sample.** TUS 2019 + 2024 pooled, rural eldest daughters aged 12–17 in family-identified single-family households, main 26 states, all diary days, one observation per household (twins dropped); every column requires positive consumption and positive denominator minutes, and columns 5–7 use column 2's sample. **Outcomes.** Columns 1–4 report log monthly household consumption per minute of household domestic work on the diary day: core domestic work (food preparation plus cleaning; cols 1–2) and all domestic work (cols 3–4). Columns 2 and 4, labelled "+ Current", add concurrent rainfall and the two childhood-history ×current-rainfall interactions. Columns 5–6 give log consumption and log core domestic-work minutes using column 2's sample and its exact specification; their coefficient difference equals the column-2 log-ratio coefficient. Column 7 is the separate, untransformed home-grown share of monthly consumption, with missing home-grown consumption set to zero. Each diary respondent's task minutes are winsorised at the state ×gender 99th percentile before summation across the household; consumption is not winsorised. **Rainfall controls.** Every column controls for the continuous infancy and early-childhood shocks of an eligible eldest son aged 12–17, re-standardised within the estimation sample and set to zero where no such son is present, together with an absent-son indicator and son-age controls. Columns with current rainfall also include son-shock ×current-rainfall and absent-son ×current-rainfall interactions. **Girl shocks.** Infancy Not-Drought (infancy rainfall above −1.15) and Early-Childhood Plenty (early-childhood rainfall at or above +1.15), tail indicators on the re-standardised life-stage shocks, with the middle band omitted; concurrent rainfall is the rolling twelve-month district shock of Table 3, re-standardised within the estimation sample. **Estimation and inference.** District-age-quarter and survey-year fixed effects; tail-indicator coefficients are relative to the middle band, and the concurrent shock is per one-standard-deviation increase. **Inference.** Standard errors clustered at district ×age ×survey-year, shown in parentheses; $***p<0.01$, $**p<0.05$, $*p<0.10$. The two p-rows test Infancy Not-Drought =Early-Childhood Plenty at concurrent = 0and at concurrent = 1.

and no imputation is required. Columns 3–4 instead use the 7,489 households in which an eligible son is observed and enter his rainfall history directly: per-standard-deviation infancy and early-childhood shocks in column 3, and separate Infancy-Not-Drought and Early-Childhood-Plenty indicators in column 4. Columns 5–6 return to the full focal-daughter sample. They set the son's continuous history to zero where he is absent, add an absent-son indicator and its interaction with current rainfall, and use respectively continuous and indicator-coded son histories. Column 5 is the preferred specification used in Table 5.

The focal-daughter interaction is 0.055 without son controls and remains 0.055 when the sample contains no eligible son. It is 0.078–0.079 when the son's history is observed directly, but the standard errors increase substantially when the sample falls by almost two-thirds. In the full frame, it is 0.058 under the preferred continuous son controls and 0.056 when the son's history is instead represented by the two infancy and early-childhood indicators; both are significant at the five-percent level. The son's own historical and current-interaction coefficients are small and statistically indistinguishable from zero when they are observed directly. The result is therefore not an artefact of imputation, of the son's rainfall coding, or of a meaningful response associated with the son rather than the focal daughter.

The same concern about household composition applies to the ratio result. Appendix Table A15 separately estimates households with one and with at least two daughters aged 12–17, excludes households containing a twin, and adds controls for sibling configuration, household composition, and parental characteristics. The Early-Childhood-Plenty interaction with current rainfall remains close to the preferred estimate throughout; precision declines only in the smaller multi-daughter sample.

Interpretation: Expansion of Home Production The positive consumption–time ratio does not imply that the household produces an unchanged consumption bundle with less total work. If domestic-work skill relaxes a bottleneck in the preparation of final consumption, the household can use the released capacity to expand activities that generate the inputs for that consumption, or to redirect time toward paid work that finances them. This makes a larger consumption response than direct domestic-work response economically coherent, even when total household production expands rather than contracts.

The time allocation of other household members is consistent with that interpretation. In households where the eldest daughter experienced Early-Childhood Plenty, other members spend 25.5 additional minutes per day on other domestic work, including activities that produce inputs for later consumption—such as tending livestock, growing vegetables, collecting materials, or making goods for the household—and 24.3 additional minutes on external paid work. Their core domestic-work coefficient is lower but imprecise. More directly, the share of their domestic work devoted to activities upstream of food preparation and cleaning rises by 2.6 percentage points. The corresponding infancy coefficients for other domestic work and external paid work are negative. These changes in how other household members use their time do not provide a second test distinguishing skill from habit. However,

they give the ratio result an economic interpretation: the early-childhood-plenty history is associated with a household that expands and reorganises home production around a daughter who is more productive at core domestic work.

The implications of this result are consequential. Habit, lower disutility of domestic work, lower returns to schooling, and other changes that reduce the opportunity cost of the daughter's time may all operate alongside skill. Yet none of these channels, on its own, can explain a higher household consumption-per-domestic-work ratio, still less a more positive response of that ratio to a new favourable rainfall shock. Under the model's maintained conditions and the empirical mapping set out above, the historical and current-rainfall results provide strong evidence that the persistent response includes domestic-work-specific skill.

This is also a practical empirical contribution. Individual output in domestic production is almost never measured, so time use alone cannot tell whether a girl has become more productive at household tasks or simply less costly to assign to them. The model shows how household consumption and household time-use totals can nevertheless supply a partial test. In this setting, they show that early engagement changed not only the allocation of the eldest daughter's time, but what the household could produce with its domestic work.

7.4 Alternative Explanations

7.4.1 Fertility and Subsequent Childcare

One alternative is that favourable rainfall during the daughter's pre-school years leads the household subsequently to add children. Her additional domestic work would then reflect the feeding and care of a larger family, rather than a domestic-work-specific capacity acquired through early engagement. The descriptive data make this a serious possibility: eldest daughters in households with younger siblings do substantially more domestic work, and the baseline means rise monotonically across the sibling-presence columns of Table 24. This gradient cannot establish the mechanism. Families with different numbers of children also differ in resources, composition, and the organisation of girls' time. We therefore ask the narrower question that follows from the fertility account: does Early Childhood rainfall increase the presence of younger siblings who could have been born during or after the focal daughter's pre-school exposure, and is the domestic-work response concentrated where those siblings are observed?

Subsequent Fertility The relevant sibling definition is determined by the timing of the proposed mechanism. A sibling at least four, or at least six, years younger than the focal daughter could have been born when she was already at the beginning of, or beyond, her pre-school exposure window. A sibling closer in age was born before that window and cannot itself be an outcome of the rainfall realisation studied here. Table 25 therefore estimates the effect of the two rainfall histories on indicators for these age-gap margins, restricting attention to single-family households in which the roster identifies the younger

child unambiguously as a biological sibling. The Early Childhood coefficients are $+0.008$ for a sibling at least four years younger and -0.001 for one at least six years younger, both statistically indistinguishable from zero. The table does not recover a complete fertility history: births followed by death, or siblings living elsewhere, will not appear in the roster. It nevertheless provides no evidence that the shock systematically expands the observed later-child margin needed for the simple child-care account. There is a small positive coefficient on the broader indicator for at least three younger siblings ($+0.009$, significant at 10%), so this exercise does not rule out every association with larger sibling groups. The distinct negative Infancy coefficients are likewise not used to infer a demographic mechanism; the relevant result is the absence of an Early Childhood increase on the age gaps that could plausibly follow the exposure period.

Domestic Work by Sibling Presence Table 24 makes the complementary comparison: does the domestic-work response remain where the relevant younger sibling is absent, and is it larger where that sibling is present? Across the four-year, six-year, and broader sibling definitions, the Early Childhood coefficient is positive and statistically significant for daughters without the relevant observed co-resident sibling. The interaction with sibling presence is small and statistically indistinguishable from zero throughout. Thus the rainfall response is neither absent in households without the later-born sibling margin nor measurably amplified where it is present. This differs from the descriptive sibling gradient: larger families are associated with more domestic work, but that cross-sectional relationship does not describe how the rainfall shock changes the eldest daughter's work.

We obtain the same conclusion when the fixed effects are saturated by sibling-presence status, so that each effect is identified within those cells rather than from their pooled comparison (Table A13). The result is therefore robust to the age-gap definition, the inclusion of sibling-presence interactions, and within-cell fixed-effect comparisons.

Taken together, these exercises weigh against later fertility and subsequent child-care demand as the principal explanation for the persistent response. The shock does not measurably raise the relevant observed later-born sibling margins, and the eldest daughter's domestic work rises even when those siblings are absent. This cannot exclude every unobserved or intensive change in fertility, but such a change alone cannot account for the observed pattern.

7.4.2 Schooling Returns and Educational Outcomes

The fertility exercises above address one source of later domestic-work demand. We next consider one observable route through which early rainfall could instead lower the opportunity cost of domestic work: a reduction in the returns to schooling. The evidence cannot address every reason why school time may become less valuable, but it can answer three direct questions. Does the pre-school shock delay school entry? Are exposed girls more likely subsequently to fall behind or withdraw from school? And does it reduce measured learning or educational attainment? These exercises do not prove the proposed mechanism.

They assess whether the domestic-work response is readily explained by a prior deterioration in school participation or achievement.

School Entry We first ask whether the pre-school shock delays the beginning of a girl's school trajectory. Table 26 combines ASER measures of current enrolment and age-appropriate grade progression with TUS measures of substantial education time, school-day commute, age-appropriate progression, and secondary attainment. Figure 10 traces the age-specific estimates from the first observed school ages onward.

The evidence does not show delayed school entry. At ages 6–8, when the pre-school exposure has only recently occurred, the Early Childhood estimates for formal enrolment are near zero or positive, rather than persistently negative. The TUS panels likewise show no sustained reduction at these ages in reporting at least 120 minutes of education time or a school-day commute. Some early commute estimates are negative, but they are imprecise and are not accompanied by lower formal enrolment or less reported education time. The available measures therefore give no indication that the shock delayed girls' entry into school.

Participation through Adolescence We next ask whether exposed girls subsequently fall behind or withdraw from school. In both sources, *on track for age* means enrolled in or above the modal class for the child's age. The ASER enrolment and on-track profiles remain close to zero or positive through the observed adolescent ages, rather than showing accumulating grade delay. The TUS indicator for at least 120 minutes of education time is also close to zero in the pooled adolescent sample (Table 26). The threshold profile reaches the same conclusion: Early Childhood rainfall does not reduce the probability of doing any studying, and the estimates are modest at lower and intermediate thresholds (Figure A14). The school-day and Sunday reallocation documented in Section 4.2 is therefore an adjustment in when, and how intensively, girls study rather than evidence of broad de-enrolment, grade delay, or school withdrawal.

Learning and Educational Attainment Finally, we ask whether lower achievement could make domestic work the comparatively more valuable use of the daughter's time. ASER tests every sampled rural child, whether or not she is enrolled, on three nested reading-and-mathematics thresholds: mastery of both Standard-I tasks, mastery of both Standard-I tasks plus at least one Standard-II task, and mastery of both Standard-II tasks. The positive Early Childhood coefficients of 1.3–1.6 percentage points in Table 27 are consequently estimated for the full sampled child population, not only for children who remain enrolled. The learning trajectories retain meaningful variation over the observed ages (Figure A15), and the TUS secondary-attainment coefficient is also positive. There is therefore no evidence that the pre-school shock first weakened foundational learning or attainment and thereby made later domestic work relatively more valuable.

Taken together, the three exercises weigh against a schooling-based explanation for the persistent domestic-work response. They do not make the later change in study time costless,

or exclude every change in the return to schooling. They show that it is not preceded by delayed entry, adolescent withdrawal, weaker progression, or lower measured achievement.

7.4.3 Health

Health is a possible parallel route from favourable rainfall to later time use. Better nutrition or physical health could change a girl's general capacity to work, her ability to benefit from school, or both, without constituting domestic-work-specific skill. The NFHS outcomes are measured in adolescence, however, rather than immediately after either exposure window. They cannot establish whether a health difference was created by an early nutritional improvement, later household conditions, the girl's changing role, or some combination of these. The useful question is therefore narrower: can a general health improvement, on its own, explain why only Early Childhood rainfall predicts the persistent domestic-work response?

Table 28 reports four higher-is-better indicators for rural girls aged 15–17: not severely thin, not anaemic, not having low blood pressure, and not stunted, together with their simple average. Both exposure windows predict a higher composite health index: $+0.004$ for the Infancy Shock and $+0.007$ for the Early Childhood Shock (each $p < 0.01$; equality rejected at $p = 0.04$). The stronger Early Childhood association is concentrated in not being anaemic and not being stunted, while the component pattern for Infancy differs. Thus the table is not evidence of broad health damage under either favourable rainfall history.

The cross-window comparison limits what health alone can explain. A general health or nutrition improvement is also present after Infancy rainfall, but Infancy rainfall does not predict the later domestic-work response. Health can therefore be part of the Early Childhood sequence, and the larger Early Childhood association may reflect a more consequential exposure period or later changes in household allocation. It cannot, by itself, account for the developmental timing of domestic work or establish the domestic-work-specific productive capacity considered in Section 7.3. The health results are best read as evidence of a parallel resource-related process, not as a substitute for the engagement and productive-capacity evidence.

7.4.4 Persistent Household and Local Consequences

Rainfall may also have durable consequences for a household or district through investment, resources, or the organisation of production. The relevant question is whether such a lasting change can explain why the later domestic-work response depends on the girl's age when rainfall occurred. The rolling estimates index each rainfall realisation by the age of the child exposed to it. To generate the positive response around ages four to seven, while leaving exposure in infancy and the early school years largely unrelated to later domestic work, a persistent local or household consequence would itself need to have sharply different effects according to the child's age at that same rainfall realisation. A generic change affecting all children then living in the household or district would not naturally generate that pattern.

7.4.5 Summary

The proposed account links findings that are difficult for the alternatives to reproduce jointly. Favourable rainfall raises agricultural output, household consumption, home-grown consumption, food ingredients, and total household domestic work, rather than merely shifting expenditure toward prepared food or drawing adult women out of domestic work. The long-run response is driven by rainfall plenty during ages four to seven, is concentrated on eldest daughters, and is visible from the earliest observed ages after that window. These are the conditions under which a temporary expansion of household production is most likely to increase a young girl's own engagement with domestic tasks. The household results then speak to what that engagement may have produced: higher household consumption per hour of domestic work, both in the long run and when a new favourable rainfall shock raises resources. Under the model, this is the pattern expected if early engagement builds domestic-work skill.

The alternative exercises reinforce that reading. The response remains where later-born siblings are absent and is not larger where they are present; it is not preceded by weaker observed schooling, progression, learning, or study participation; and a generic persistent household or local consequence would need to vary with a child's age at the same historical rainfall realisation. Health can be a parallel resource-related process, but its positive association with both Infancy and Early Childhood rainfall means that it cannot by itself explain the Early Childhood-specific domestic-work result. We do not observe the girl's childhood task inputs or her individual domestic output, and the evidence cannot establish whether all persistence is technical skill rather than another change induced by early engagement. It nevertheless provides convergent evidence that favourable pre-school rainfall draws eldest daughters into household production and that this engagement develops a domestic-work-specific productive capacity which remains relevant in adolescence. Section 8 turns to how households use that capacity when they face new shocks and to its observable adolescent consequences.

8 Domestic-Work Capacity and Adolescent Gender Roles

Section 7 shows that pre-school rainfall plenty is associated with more domestic work in adolescence, principally among eldest daughters, and provides evidence of a domestic-work-specific skill component. A change in girls' domestic-work time, however, may be part of a broader gendered adolescent role. We therefore examine two related questions. First, when current agricultural and household-production conditions improve, how does the eldest daughter's day adjust? Second, is the childhood history associated with later domestic work also accompanied by differences in observed mobility, permissions, or access to resources?

8.1 Current Household Demands and School-Time Adjustment

The historical estimates describe standing differences in the eldest daughter's adolescent time allocation, but they do not show how those differences respond when the household faces a new positive shock. We use concurrent rainfall as a proxy for current agricultural conditions and ask whether the resulting expansion is accompanied by a release of the daughter's school time or by a further reorganisation of her day. The interaction does not provide a second test of the mechanism behind the childhood effect. It instead traces which parts of the daughter's observed day change when a new household-production shock coincides with the relevant childhood history.

Table 6 implements this comparison for rural eldest daughters aged 12–17 in family-identified single-family households. It includes two historical indicators. *Early-Childhood Plenty* is the upper-tail pre-school rainfall exposure that generated the mechanism evidence. *Infancy Not-Drought* contrasts the non-drought part of the infancy distribution with its lower tail, where the infancy response is concentrated. Both histories are interacted with concurrent rainfall, the standardised district rainfall shock over the twelve months ending in the diary quarter. The historical coefficients therefore describe differences at average concurrent rainfall, and the interaction coefficients show how those differences change as current conditions improve. The seven outcomes are daily minutes in domestic work, total studying, school attendance, homework, leisure, sleep, and socialising.

TABLE 6: *Time Use under Childhood and Concurrent Rainfall Shocks: Eldest Daughters*

	Domestic Work (mins/day)	Studying (mins/day)			Leisure (mins/day)		
	Total	Total	Attendance	Homework	Total	Sleep	Socialising
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Infancy Not Drought	24.07***	-24.43**	-15.03**	0.38	-6.10	5.59	3.10
	(5.43)	(10.29)	(7.20)	(4.39)	(7.89)	(4.14)	(4.61)
Early-Childhood Plenty	21.97***	-25.01***	-20.85***	-0.64	0.95	0.29	5.66
	(5.46)	(8.98)	(6.33)	(3.99)	(7.58)	(3.81)	(4.30)
Concurrent Rain	4.12	-0.13	-0.63	1.95	-2.73	3.66	12.21***
	(5.05)	(8.62)	(6.16)	(3.50)	(6.17)	(3.60)	(3.26)
Infancy Not Drought ×Concurrent	-1.00	-5.96	0.37	-6.38*	4.78	-6.89*	-6.26*
	(5.31)	(8.97)	(6.40)	(3.77)	(6.54)	(3.83)	(3.42)
Early-Childhood Plenty ×Concurrent	7.51	-31.10***	-27.69***	-5.24	22.92**	13.18***	7.72*
	(6.03)	(10.82)	(7.34)	(4.40)	(9.02)	(4.31)	(4.60)
Observations	20,653	20,653	20,653	20,653	20,653	20,653	20,653
Outcome Mean	97.7	346.6	196.7	105.8	980.6	564.0	96.8
[Concurr. =0] p: Infancy =Early Ch.	[0.786]	[0.966]	[0.536]	[0.866]	[0.518]	[0.349]	[0.688]
[Concurr. =1] p: Infancy =Early Ch.	[0.602]	[0.189]	[0.011]	[0.988]	[0.092]	[0.063]	[0.062]
District-Age-Qtr FEs	Y	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y	Y

Notes: **Sample.** Rural eldest daughters aged 12–17 in family-identified single-family households, TUS 2019 and 2024 pooled, in the main 26 states. **What it shows.** Seven outcomes in minutes per day: total domestic work; total studying, school attendance, and homework; and total leisure, sleep, and socialising. Regressions include Infancy Not-Drought, Early Childhood Plenty, concurrent rainfall, and interactions between concurrent rainfall and each childhood shock. **Rainfall shocks.** Infancy Not-Drought is an indicator for the Infancy rainfall z-score above -1.15. Early Childhood Plenty is an indicator for the Early Childhood rainfall z-score at or above $+1.15$. Concurrent rainfall is the rolling twelve-month district rainfall z-score ending in the diary quarter. **Estimation and inference.** Regressions include rural ×district ×age ×quarter and survey-year fixed effects. The two p-value rows test equality of Infancy Not-Drought and Early Childhood Plenty when concurrent rainfall equals zero and one. Standard errors are clustered at district ×age ×survey-year and shown in parentheses. $***p<0.01$, $**p<0.05$, $*p<0.10$.

When concurrent rainfall rises, the two histories diverge. The Infancy-Not-Drought interactions are small and imprecise for domestic work, studying, and attendance. For Early-Childhood Plenty, the interaction in total domestic work is also imprecisely estimated. The clear adjustment is instead in school time: studying falls by a further 31 minutes and school attendance by 28 minutes, while leisure and sleep rise. The eldest-son placebo does not reproduce this pattern (Table A16). Thus, current rainfall does not generate a precisely measured additional increase in the exposed girl's total domestic-work minutes. It coincides with a further adjustment away from attendance and studying only for the pre-school-plenty history.

A school-time response without a precisely estimated increase in total daily domestic work is not necessarily paradoxical. The relevant adjustment may be responsibility for tasks at particular times of day, rather than the total amount completed over the day. An eldest daughter may, for example, cover meal preparation or other domestic tasks while an adult reallocates time to farm or paid work; this can conflict directly with school attendance even if other tasks offset it in the daily diary total.

This result offers one explanation for why positive resource shocks and income interventions

do not uniformly release girls from unpaid domestic work or increase their schooling (Edmonds and Schady, 2012; Shah and Steinberg, 2021; Bai and Wang, 2020b; de Hoop et al., 2020). When higher resources or better adult market opportunities expand household production, an eldest daughter's ability to substitute for adults in domestic work can make her time particularly valuable at home. Her school attendance may then be the margin that adjusts, even as the shock relaxes other resource constraints. Because eldest daughters are especially likely to occupy that substitute role, this mechanism may help explain why favourable economic conditions need not narrow gender gaps in schooling among older adolescents.

8.2 Gender Roles Beyond Domestic Work

Domestic work is one domain of a wider gendered adolescent role. We ask whether the childhood history associated with later domestic work is also accompanied by differences beyond the division of chores: discretionary movement outside the home in the TUS, and reported *Autonomy*, *Mobility*, *Empowerment*, and cinema attendance in the NFHS. The available surveys do not directly measure gender norms, bargaining power, or subjective well-being. They nevertheless allow us to assess whether the persistent domestic-work pattern is accompanied by more restrictive gendered patterns in girls' mobility, permissions, and access to resources.

The TUS provides the closest time-use measure of movement. We construct discretionary outside-home time from the 30-minute diary slots spent outside the household premises, after excluding paid work, household production, and formal study. The measure therefore excludes school travel and work-related time outside the home. Rural girls aged 15–17 have substantially less discretionary outside-home time than boys: 40.2 rather than 77.8 minutes on weekdays, and 74.6 rather than 162.8 minutes on Sundays. If the early-childhood history were associated with a broader restriction on movement, this gender gap should widen. Instead, the Early Childhood Shock is associated with 7.6 additional weekday minutes outside the home and a 2.4-percentage-point increase in the probability of any such time; the Sunday estimates are positive but less precise (Table 29). The large baseline gender gap in discretionary movement therefore does not become larger under the childhood exposure associated with domestic work.

8.2.1 Autonomy, Mobility, and Resources in the NFHS

The NFHS groups the first set of outcomes as *Autonomy* and *Mobility*. We use NFHS-4 and NFHS-5 for unmarried rural girls aged 15–17 in the main 26 states. The full women's-questionnaire sample contains 103,107 girls. The three *Autonomy* outcomes summarise five potential barriers to seeking health care: needing permission, lack of money, distance, transport, and not wanting to go alone. We report the mean share of barriers that are not a major problem, and indicators that all five are no major problem or no problem at all. They

are measures of reported access to health care, not direct measures of bargaining power. A smaller state-rotating module, with 16,519 girls, supplies the two *Mobility* outcomes: whether the respondent is allowed to visit a market, health facility, and place outside the village, either at all or alone. All five outcomes are coded so that a higher value indicates fewer reported restrictions.

The *Autonomy* and *Mobility* estimates are modestly positive for the Early Childhood Shock. The mean health-care-access index rises by 0.008, and the indicator that none of the five barriers is a major problem rises by 0.007; the stricter no-problem-at-all indicator is unchanged. The mean mobility indicator for being allowed to visit the three destinations rises by 0.008, while the indicator for being allowed alone to all three is small and imprecise. The corresponding Infancy coefficients are small and statistically indistinguishable from zero. The history associated with later domestic work is thus not associated with lower reported health-care access or more restricted mobility, although these estimates do not show broad new freedom of movement.

The second NFHS table labels its first four outcomes *Empowerment*. They record access to specific resources and financial participation: money for personal spending, a bank account, a mobile phone, knowledge of a microcredit programme, and having taken a loan. In the state-rotating module, we report their mean and indicators for any, at least three, and all five. Monthly cinema attendance, a separate indicator of discretionary activity, comes from the larger full-questionnaire sample.

Early Childhood rainfall raises the mean *Empowerment* index by 0.006, the probability of any item by 0.018, the probability of at least three items by 0.009, and monthly cinema attendance by 0.003; the very rare all-five outcome is unchanged (Table 31). The corresponding Infancy coefficients are negative for the mean index, having at least three items, and cinema attendance; the Early Childhood and Infancy coefficients differ precisely on these margins. Thus, the Early Childhood history associated with later domestic work does not carry the lower resource access, financial participation, and cinema attendance observed for the infancy comparison.

Taken together, the evidence does not support a broader deepening of gendered disadvantage in girls' adolescent lives beyond domestic work. The childhood history associated with more domestic work is not accompanied by less discretionary movement, lower reported health-care access, more restricted *Mobility*, or lower observed *Empowerment*. Read with the schooling and health evidence in Section 7, the durable effect appears concentrated in domestic work rather than extending to a wider deterioration in the adolescent outcomes observed here.

This separation is consistent with Ashraf et al. (2020), who show that adolescent girls can rearrange household chores around schoolwork through negotiation, increasing educational investment without eroding norms of deference. Although their intervention differs from the childhood exposure studied here, it illustrates the narrower point: domestic-work assignment need not move one-for-one with other permissions, resources, or activities. The present

evidence points to such a separation.

9 Conclusion

This paper shows that the production of household consumption can turn a favourable childhood shock into a persistent gendered demand for labour. In rural India, unusually favourable rainfall during ages four to seven raises girls' domestic work in adolescence, with no comparable response for boys. The effect is concentrated among eldest daughters and in years of rainfall plenty rather than drought. This is not the simple pattern implied by an early resource gain that only relaxes household constraints. Favourable rainfall raises agricultural output and consumption, but it also raises the domestic work required to turn household inputs into final consumption. The paper identifies this household-production channel as one route through which a positive shock can recruit girls into work at a formative age and leave a durable mark on their later allocation of time.

The evidence advances the study of gendered time allocation in two ways. First, it moves beyond documenting that girls do more unpaid work to show how a domestic-work role can be formed. The long-run response is attached to the eldest daughter, the child most likely to have absorbed an increase in household labour demand, and is visible from the earliest observed ages after the relevant childhood window. The pattern is not consistent with a generic effect of household rainfall exposure on all girls. Nor is it accompanied by broad deterioration in school enrolment, grade progression, learning, health, mobility, or the observed indicators of agency. The result is therefore more specific than a general worsening of girls' adolescent circumstances: it is a durable domestic-work role that is empirically separable from a broader restrictive gender-role profile in adolescence.

Second, the paper provides a way to study whether domestic-work experience changes productive capacity when individual domestic output is not observed. Time-use data alone cannot distinguish skill from a lower opportunity cost of assigning a girl to household work. The model shows that, under its maintained conditions, the latter channel cannot by itself raise household monetary consumption per unit of domestic work, or make that ratio respond more positively to a new favourable rainfall shock. Households with an eldest daughter exposed to early-childhood rainfall plenty display both patterns. Read alongside the rank and timing evidence, this is the paper's central learning-by-doing result: a temporary increase in the eldest daughter's engagement in domestic tasks builds an enduring, task-specific productive capacity. This is one-sided evidence rather than a full structural decomposition: other changes in the value of the girl's time may operate alongside skill. But it rules out an account in which the persistent response is only habit, lower disutility, or lower returns to alternative uses of her time. More broadly, it demonstrates how household consumption and time-use totals can reveal a productive-capacity channel that conventional individual-level data do not measure.

These findings change the policy question. The issue is not simply whether a household has more resources, or whether a girl is formally enrolled in school. Where consumption remains labour intensive, agricultural growth, income support, or other favourable shocks can leave the domestic-work constraint intact and may increase the household's use of the girl who has become capable of meeting it. Policies that reduce the labour required to produce and maintain consumption—reliable water, fuel and electricity, time-saving household technologies, childcare, and access to services or prepared goods—may therefore complement investments in schooling and income. The relevant design question is whether an intervention changes the household-production constraint as well as its budget constraint.

The first next step is direct observation. Which tasks, forms of engagement, and ages of exposure build domestic-work skill? How does that learning interact with cognitive development and with returns to schooling? And does becoming more capable at household work reinforce gendered restrictions, or provide a girl with a bargaining resource that helps her negotiate the timing and scope of those restrictions? Answering these questions requires panel data that link children's task inputs to household output, cognitive outcomes, and measures of decision-making and bargaining.

A second agenda concerns the temporal technology of household production. The evidence suggests that the constraint is not only the total labour intensity of producing consumption. Much domestic work must be performed on particular days and at particular times: meals, cleaning, care, and the processing of perishable inputs cannot simply be stored and shifted to another day. This helps explain why school-day study time is constrained while it recovers on Sundays and other non-normal days. Future work should examine how infrastructure, refrigeration and storage, prepared-food markets, and services change both the amount and timing of domestic work. That question also connects household production to structural transformation: a growing supply of formal and professional jobs, especially for women, may depend not only on higher wages or schooling but on technologies and services that make domestic production less labour intensive and more temporally flexible. The central lesson is that gendered time allocation is not only a distribution of tasks at a point in time; it can also be part of how gender-specific productive capacity is formed during childhood.

References

Almond, D. and Currie, J. (2011). Killing me softly: The fetal origins hypothesis. *Journal of Economic Perspectives*, 25(3):153–172.

Arrow, K. J. (1962). The economic implications of learning by doing. *The Review of Economic Studies*, 29(3):155–173.

Ashraf, N., Bau, N., Low, C., and McGinn, K. (2020). Negotiating a better future: How interpersonal skills facilitate inter-generational investment. *The Quarterly Journal of Economics*, 135(2):1095–1151.

Bai, J. and Wang, Y. (2020a). Returns to work, child labor and schooling: The income vs. price effects. *Journal of Development Economics*, 145:102466 (article no.). Article-number journal (no page range). Confirmed via IDEAS/RePEc.

Bai, J. and Wang, Y. (2020b). Returns to work, child labor and schooling: The income vs. price effects. *Journal of Development Economics*, 145:102466.

Baland, J.-M. and Robinson, J. A. (2000). Is child labor inefficient? *Journal of Political Economy*, 108(4):663–679.

Basu, K. and Van, P. H. (1998). The Economics of Child Labor. *American Economic Review*, 88(3):412–427.

Beaman, L., Duflo, E., Pande, R., and Topalova, P. (2012). Female leadership raises aspirations and educational attainment for girls: A policy experiment in india. *Science*, 335(6068):582–586.

Becker, G. S. (1965). A theory of the allocation of time. *Economic Journal*, 75(299):493–517.

Canay, I. A. (2011). A simple approach to quantile regression for panel data. *The Econometrics Journal*, 14(3):368–386.

Cunha, F. and Heckman, J. J. (2007). The technology of skill formation. *American Economic Review*, 97(2):31–47.

de Hoop, J., Groppo, V., and Handa, S. (2020). Cash transfers, microentrepreneurial activity, and child work: Evidence from Malawi and Zambia. *The World Bank Economic Review*, 34(3):670–697.

Del Carpio, X. V. and Macours, K. (2010). Leveling the intra-household playing field: Compensation and specialization in child labor allocation. In *Research in Labor Economics, Vol. 31: Child Labor and the Transition between School and Work*, pages 259–295. Emerald.

Dhar, D., Jain, T., and Jayachandran, S. (2022). Reshaping adolescents' gender attitudes: Evidence from a school-based experiment in india. *American Economic Review*, 112(3):899–927.

Dinkelman, T. and Ngai, L. R. (2022). Time use and gender in Africa in times of structural transformation. *Journal of Economic Perspectives*, 36(1):57–80.

Dubois, P. and Rubio-Codina, M. (2012). Child care provision: Semiparametric evidence from a randomized experiment in mexico. *Annales d'Économie et de Statistique*, 105/106:155–184.

Edmonds, E. V. and Schady, N. (2012). Poverty alleviation and child labor. *American Economic Journal: Economic Policy*, 4(4):100–124.

Foster, A. D. and Rosenzweig, M. R. (1995). Learning by doing and learning from others: Human capital and technical change in agriculture. *Journal of Political Economy*, 103(6):1176–1209.

Gottlieb, C., Doss, C., Gollin, D., and Poschke, M. (2024). The gender division of work across countries. *SSRN Electronic Journal*.

Greenwood, J., Seshadri, A., and Yorukoglu, M. (2005). Engines of Liberation. *The Review of Economic Studies*, 72(1):109–133.

Gronau, R. (1977). Leisure, home production, and work—the theory of the allocation of time revisited. *Journal of Political Economy*, 85(6):1099–1123.

Jayachandran, S. (2015). The roots of gender inequality in developing countries. *Annual Review of Economics*, 7:63–88.

Karve, I. (1953). *Kinship Organization in India*. Deccan College Postgraduate and Research Institute, Poona.

Kaur, S. (2019). Nominal wage rigidity in village labor markets. *American Economic Review*, 109(10):3585–3616.

Kis-Katos, K. (2007). The Work–Schooling Trade-off Revisited: Market and Domestic Work of Indian Children. Working paper. Study of children in two North Indian states; publication outlet not verified from the source copy.

Kolenda, P. M. (1968). Region, caste, and family structure: A comparative study of the Indian "joint" family. In Singer, M. and Cohn, B. S., editors, *Structure and Change in Indian Society*, pages 339–396. Aldine Publishing Company, Chicago.

Lancy, D. F. (2015). *The Anthropology of Childhood: Cherubs, Chattel, Changelings*. Cambridge University Press, 2nd edition.

Maccini, S. and Yang, D. (2009). Under the weather: Health, schooling, and economic consequences of early-life rainfall. *American Economic Review*, 99(3):1006–1026.

Muralidharan, K. and Prakash, N. (2017). Cycling to school: Increasing secondary school enrollment for girls in india. *American Economic Journal: Applied Economics*, 9(3):321–350.

Ngai, L. R. and Petrongolo, B. (2017). Gender gaps and the rise of the service economy. *American Economic Journal: Macroeconomics*, 9(4):1–44.

Pitt, M. M., Rosenzweig, M. R., and Hassan, M. N. (2012). Human Capital Investment and the Gender Division of Labor in a Brawn-Based Economy. *American Economic Review*, 102(7):3531–3560.

Runfola, D., Anderson, A., Baier, H., Crittenden, M., Dowker, E., Fuhrig, S., Goodman, S.,

Grimsley, G., Layko, R., Melville, G., Mulder, M., Oberman, R., Panganiban, J., Peck, A., Seitz, L., Shea, S., Slevin, H., Youngerman, R., and Hobbs, L. (2020). geoBoundaries: A global database of political administrative boundaries. *PLOS ONE*, 15(4):e0231866.

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.

Shah, M. and Steinberg, B. M. (2021). Workfare and human capital investment: Evidence from India. *Journal of Human Resources*, 56(2):380–405.

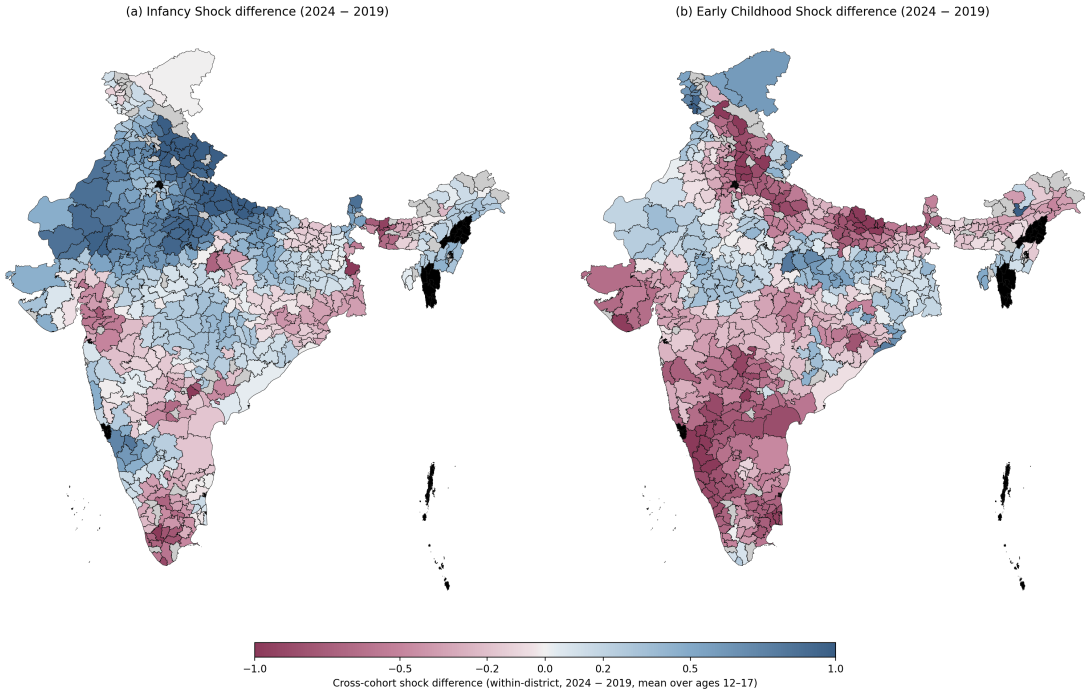
Webb, D. (2024). Critical periods in cognitive and socioemotional development: Evidence from weather shocks in indonesia. *The Economic Journal*, 134(660):1637–1665.

Whiting, B. B. and Edwards, C. P. (1988). *Children of Different Worlds: The Formation of Social Behavior*. Harvard University Press.

A Paper Appendix

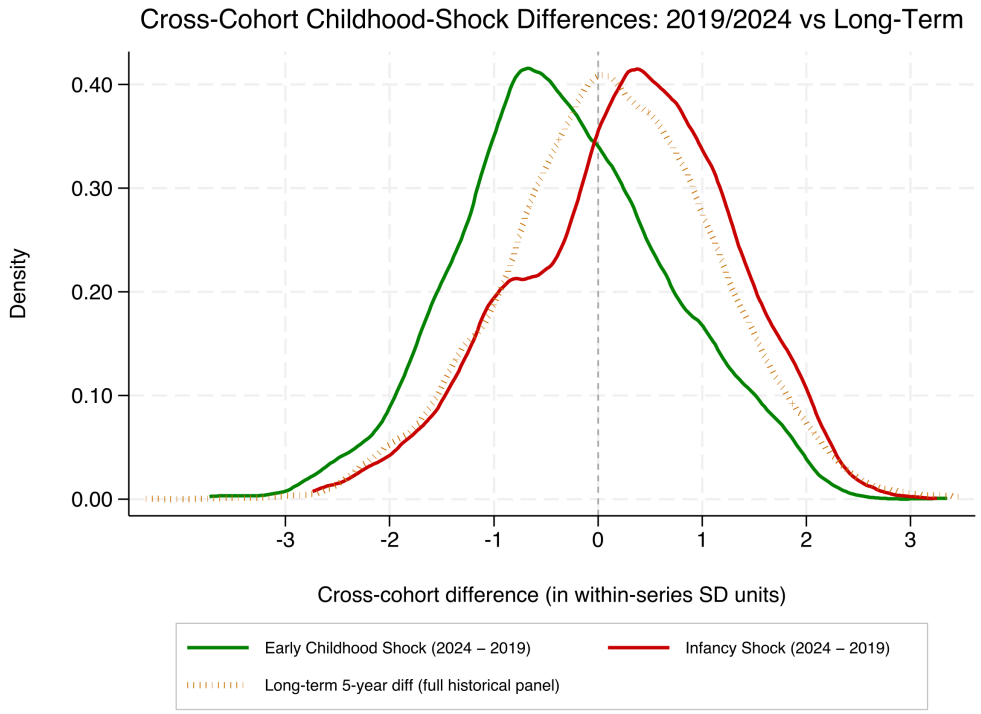
Figures

FIGURE 4: *Geographic Distribution of Cross-Cohort Rainfall-Shock Differences*



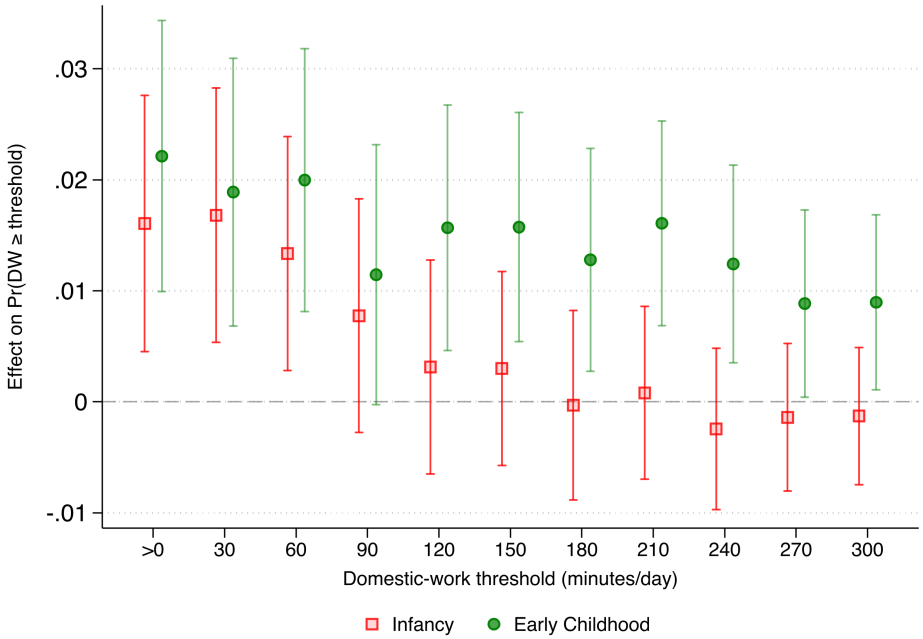
Note: **Sample.** Rural children aged 12–17, TUS 2019 and 2024 pooled, in the main 26 states. **What is plotted.** Two district choropleths of within-district cross-cohort shock differences. For each district d and age $a \in \{12, \ldots, 17\}$, $\Delta_a(d)$ is the mean shock among a-year-olds in 2024 minus the mean shock among a-year-olds in 2019; the district value is the simple mean of these differences across ages observed in both waves. Panel (a) plots Infancy Shock differences and panel (b) plots Early Childhood Shock differences. **Polygons and coverage.** Polygons use the TUS 2019 adjusted district boundaries (Section 3.2.1). Districts outside the main 26-state sample are shaded black. Districts within the main 26 states with no matched-age cell in both TUS waves are shaded light grey.

FIGURE 5: *Cross-Cohort Childhood-Rainfall Differences and the Long-Run Distribution*



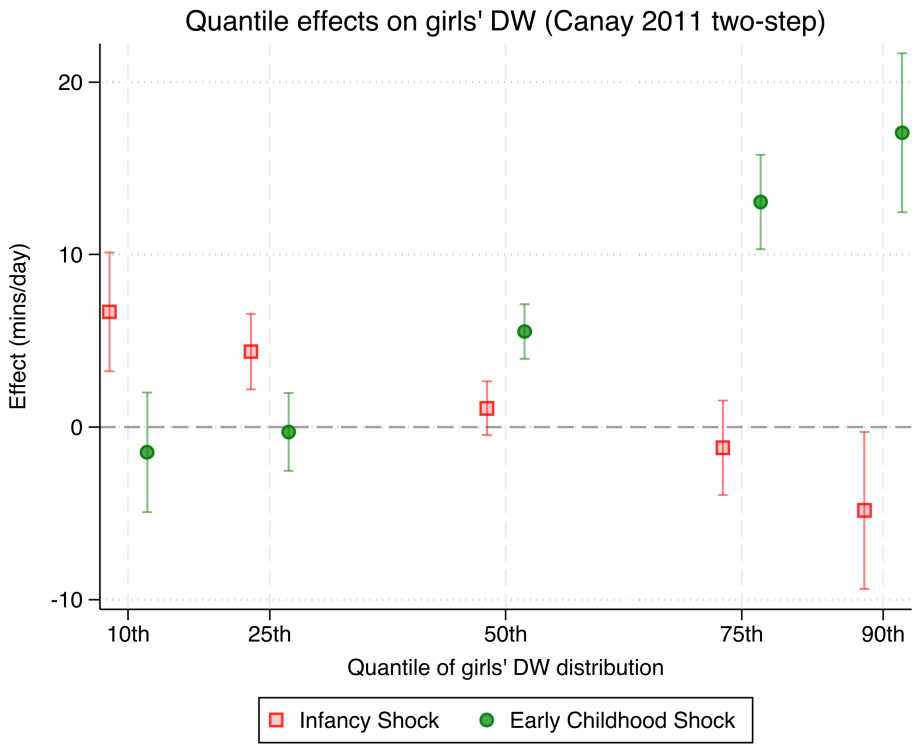
Note: **Sample.** The Early Childhood and Infancy series use rural children aged 12–17 in the main 26 states, TUS 2019 and 2024 pooled, on balanced district-age cells. The long-run series uses all district-year pairs in the GPCC Monitoring panel for which four-year rainfall averages can be calculated in years t and $t-5$. **What is plotted.** Kernel densities of three cross-cohort differences. The Early Childhood and Infancy series are 2024 minus 2019 differences in cell-mean rainfall shocks. The long-run series is the five-year difference in four-year average annual rainfall z-scores. Each series is divided by its own standard deviation before plotting.

FIGURE 6: *Effects of Childhood Rainfall Shocks across Domestic-Work Thresholds*



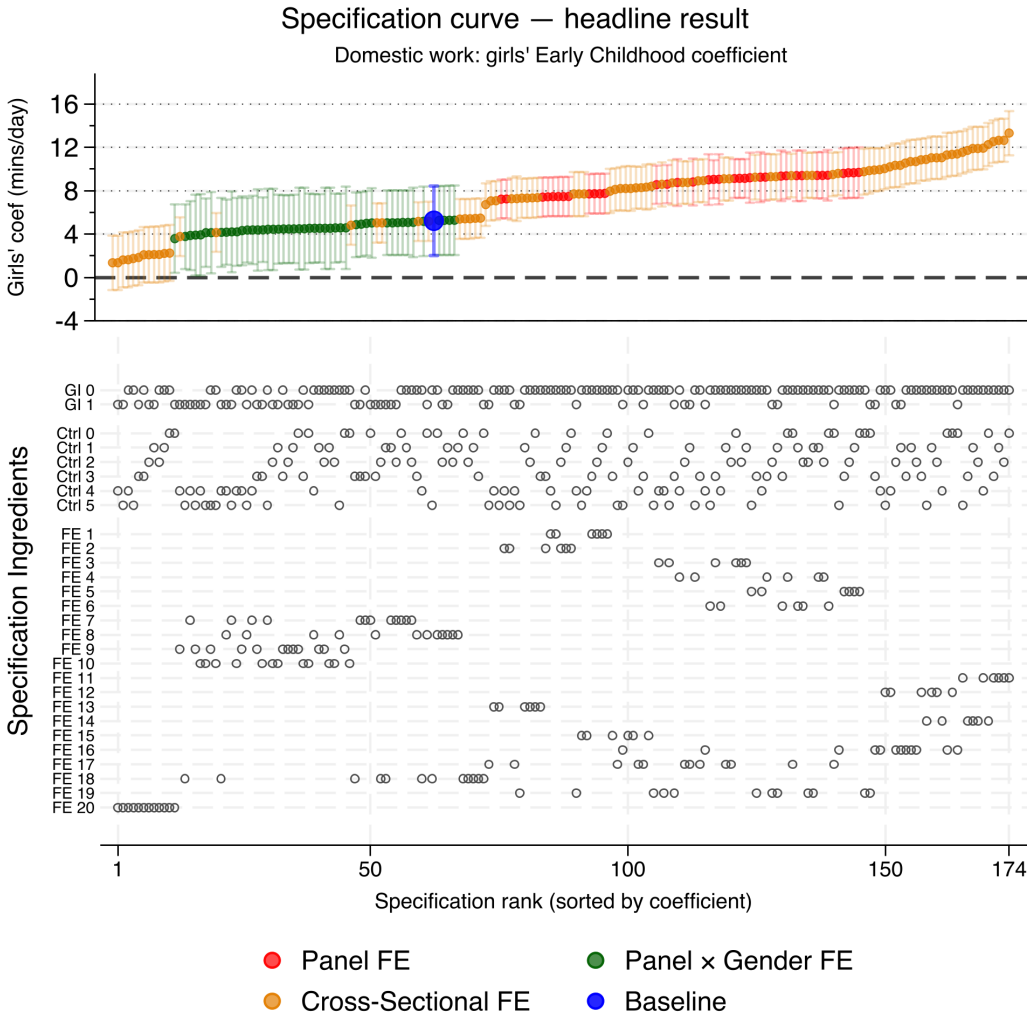
Note: **Sample.** Rural girls aged 12–17 in the main 26 states, TUS 2019 + 2024 pooled; shocks re-standardised within sample. **What is plotted.** For each stated domestic-work threshold (any domestic work, then 30, 60, . . . , 300 minutes per day), the outcome is an indicator that daily domestic work is at least that threshold. Each marker is the girls-only coefficient on the relevant childhood shock, with a 95% confidence interval. **Estimation.** Fixed effects rural ×district ×age ×quarter and survey-year; domestic work is winsorised at the state ×gender 99th percentile before the indicators are constructed. Standard errors clustered at the district ×age ×survey-year level.

FIGURE 7: *Quantile Effects of Childhood Rainfall Shocks on Girls' Domestic Work*



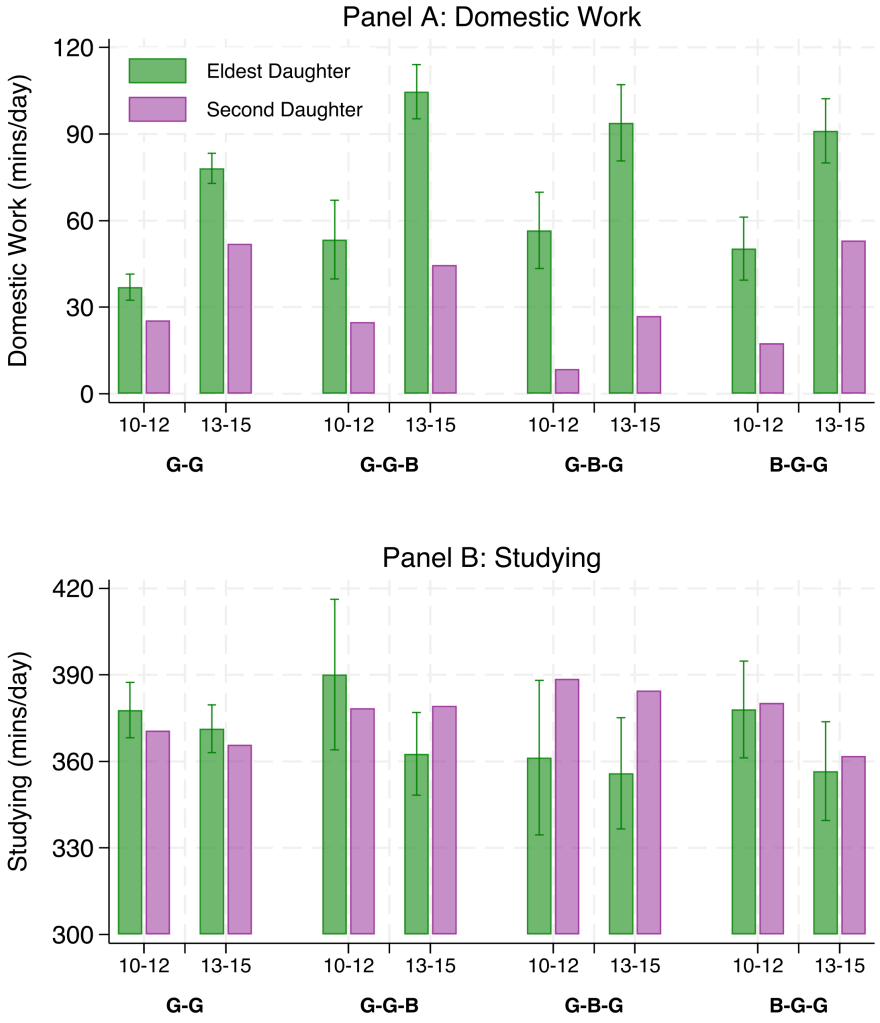
Note: **Sample.** Rural girls aged 12–17 with positive domestic work, TUS 2019 and 2024 pooled, in the main 26 states. **What is plotted.** Early Childhood (green) and Infancy (red) rainfall-shock coefficients for domestic-work minutes per day at the 10th through 90th percentiles of the positive domestic-work distribution, estimated using the Canay (2011) two-step estimator. **Estimation.** In the first step, the outcome is residualised on the headline fixed effects under a location-shift assumption. The second step estimates quantile regressions of the residualised outcome. Shocks are re-standardised within the estimation sample. Fixed effects are rural ×district ×age ×quarter and survey-year; standard errors are clustered at district ×age ×survey-year.

FIGURE 8: *Specification Curve for the Early Childhood Shock Effect on Girls' Domestic Work*



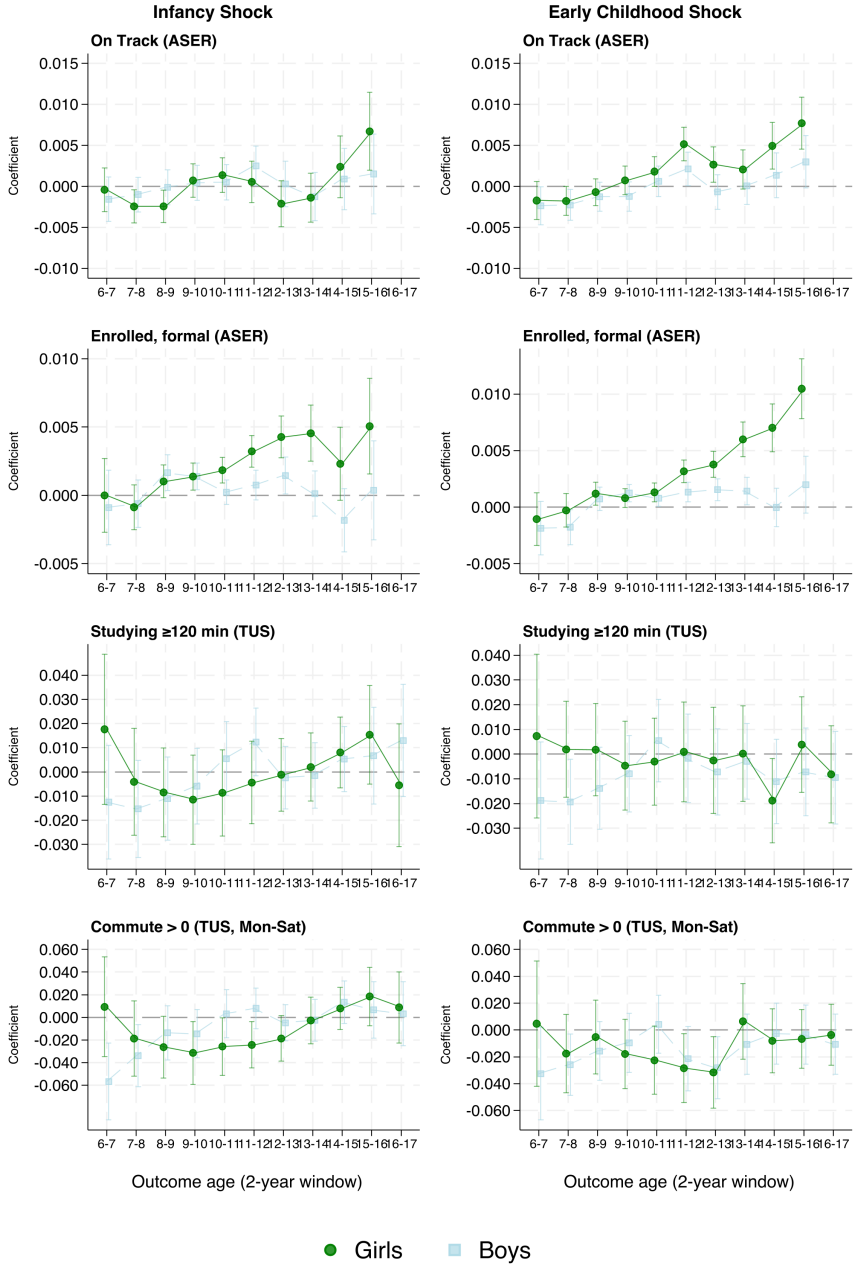
Note: **What is plotted.** Girls' Early Childhood Shock coefficients on domestic work from 174 specifications varying the sample, controls, fixed effects, winsorisation rule, and rainfall-shock construction. The upper panel reports each coefficient in minutes per day with a 95% confidence interval, ordered from smallest to largest. The lower panel records the analytic choices used in each specification.

FIGURE 9: *Domestic Work and Studying by Daughter Rank and Family Type*



Note: **Sample.** Rural girls aged 10–15, TUS 2019 + 2024 pooled, households with identifiable parent–child relationships, restricted to girls who are the firstborn or second-born daughter in the household. **What is plotted.** Predicted within-cell means of the outcome from a regression that fully interacts household family type, age group (10–12 / 13–15), and daughter rank (firstborn vs second-born girl). Panel A: domestic work in minutes per day; Panel B: studying in minutes per day. Within each age-group and family-type cell, the green bar is the eldest daughter and the purple bar is the second daughter. **Domestic work.** Total domestic work in minutes per day, as defined in the paper; winsorised state ×gender p99. **Studying.** All formal study time in minutes per day. **Family-type categorisation.** Households partitioned by sibling gender ordering among children aged ≤17: G-G (two or more girls, no boys); G-G-B (eldest two girls, at least one younger boy); G-B-G (firstborn girl, then a boy, then a later-born girl); B-G-G (a boy older than every girl). **Estimation.** Fixed effects district ×survey-year ×quarter (between-household design); clustered at district ×survey-year. 95% confidence bars on each predicted mean.

FIGURE 10: *Enrolment and Achievement Trajectory by Age (ASER and TUS)*



Note: **Sample.** ASER rows: rural ASER, the main 26 states, birth cohorts 2000–2014, ages 6–16. TUS rows: TUS 2019 + 2024, rural, child-roster restriction, ages 6–17; the commute row restricts further to Mon–Sat diaries. **What is plotted.** Four outcome rows by outcome-age window. Left column: Infancy Shock; right column: Early Childhood Shock. Girls (green) and boys (blue); each marker is a separate regression on a two-year window centred at the marked age. **Outcomes.** *On-track for age* (ASER): enrolled in or above the modal class for the child's age. *Enrolled, formal* (ASER): currently in school in class 1–12. *Studying ≥120*(TUS): indicator for at least 120 daily minutes on the broad 1-digit education total (includes commute). *Commute >0* (Mon–Sat, TUS): indicator for any travel-related-to-learning minutes on Mon–Sat diaries. **Estimation.** ASER fixed effects district ×age ×gender + survey-year ×gender; TUS headline fixed-effect structure; clustered at district ×age ×survey-year. For 6-year-olds the Early Childhood Shock uses ages 4–6.

Tables

TABLE 7: *Sample Coverage by State*

State	Rural girls 12–17 (N)	Mean DW (mins/day)	Headline sample
Uttar Pradesh	5,331	121.4	Y
Bihar	3,001	79.4	Y
Rajasthan	2,182	118.7	Y
Madhya Pradesh	1,793	128.0	Y
Maharashtra	1,790	86.8	Y
West Bengal	1,634	41.2	Y
Gujarat	1,150	118.7	Y
Tamil Nadu	1,067	39.2	Y
Odisha	1,044	69.6	Y
Assam	1,015	42.9	Y
Jharkhand	870	109.9	Y
Karnataka	854	58.6	Y
Chhattisgarh	746	102.0	Y
Andhra Pradesh	659	59.3	Y
Punjab	565	59.1	Y
Kerala	557	41.9	Y
Haryana	509	49.0	Y
Tripura	488	57.8	Y
Manipur	460	32.5	Y
Meghalaya	436	105.8	Y
Telangana	369	64.4	Y
Jammu & Kashmir	295	50.9	Y
Uttarakhand	291	64.2	Y
Himachal Pradesh	256	62.0	Y
Arunachal Pradesh	168	99.3	Y
Sikkim	167	47.1	Y
Nagaland	165	102.7	.
Mizoram	114	85.9	.
Dadra & Nagar Haveli	71	94.6	.
Andaman & Nicobar	49	53.9	.
Goa	45	38.7	.
Daman & Diu	31	67.7	.
Puducherry	31	37.3	.
Lakshadweep	19	17.4	.
Delhi	15	66.0	.
Chandigarh	4	93.8	.

Notes: **Sample.** TUS 2019 and 2024 pooled, rural households with a completed diary and at least one child aged 17 or younger. One row per state. **What it shows.** State-level sample counts, mean domestic-work minutes, and indicators for inclusion in the main 26-state estimation sample and the high-domestic-work sample. **Domestic work.** Mean domestic work is winsorised at the state ×gender 99th percentile.

TABLE 8: *Time Use by Gender and Age Band*

	Age 7 to 11				Age 12 to 17			
	Minutes (per day)		Any (Mins > 0)		Minutes (per day)		Any (Mins > 0)	
	Boys	*Girls*	*Boys*	*Girls*	*Boys*	*Girls*	*Boys*	*Girls*
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
All Domestic Work	9.3	17.7	0.11	0.18	20.1	87.4	0.20	0.55
		[0.00]		*[0.00]*		*[0.00]*		*[0.00]*
Food Preparation	0.5	3.9	0.01	0.06	1.9	40.9	0.03	0.38
		[0.00]		*[0.00]*		*[0.00]*		*[0.00]*
Cleaning	0.8	3.5	0.02	0.07	2.7	25.1	0.05	0.35
		[0.00]		*[0.00]*		*[0.00]*		*[0.00]*
Domestic Production	2.1	2.8	0.02	0.04	9.8	13.7	0.09	0.15
		[0.00]		*[0.00]*		*[0.00]*		*[0.00]*
Care work	5.4	7.0	0.06	0.07	3.5	5.6	0.04	0.07
		[0.00]		*[0.00]*		*[0.00]*		*[0.00]*
Budgeting & Shopping	0.5	0.5	0.01	0.01	2.1	2.0	0.04	0.03
		[0.19]		*[0.99]*		*[0.49]*		*[0.00]*
All Paid Work	2.2	1.7	0.03	0.02	31.9	12.8	0.12	0.07
		[0.01]		*[0.00]*		*[0.00]*		*[0.00]*
In HH Enterprise	0.8	0.6	0.00	0.00	23.0	9.4	0.07	0.04
		[0.17]		*[0.24]*		*[0.00]*		*[0.00]*
All Studying	361.1	361.8	0.88	0.89	351.0	348.3	0.80	0.79
		[0.69]		*[0.56]*		*[0.15]*		*[0.00]*
Homework	93.3	93.1	0.77	0.77	101.0	104.6	0.70	0.70
		[0.85]		*[0.32]*		*[0.00]*		*[0.76]*
All Leisure and Self-care	1,067.2	1,058.6	1.00	1.00	1,036.5	990.7	1.00	1.00
		[0.00]				*[0.00]*		
Sleep	606.9	607.3	1.00	1.00	569.5	568.0	1.00	1.00
		[0.59]				*[0.04]*		
Personal Care (incl. eat/drink)	170.0	170.8	1.00	1.00	168.8	167.5	1.00	1.00
		[0.12]				*[0.00]*		
Entertainment & Hobbies	227.8	215.4	0.98	0.97	206.6	159.5	0.96	0.91
		[0.00]		*[0.00]*		*[0.00]*		*[0.00]*
Socialising	60.3	63.4	0.71	0.72	87.9	93.8	0.81	0.84
		[0.00]		*[0.01]*		*[0.00]*		*[0.00]*
Time spent Outside	409.2	392.6	0.94	0.93	434.4	353.3	0.96	0.88
		[0.00]		*[0.00]*		*[0.00]*		*[0.00]*
Observations	25,212	21,058	25,212	21,058	32,590	28,927	32,590	28,927

Notes: **Sample.** Rural children aged 7–17 with completed time diaries, TUS 2019 and 2024 pooled. Columns 1–4 cover ages 7–11 and columns 5–8 cover ages 12–17. **What it shows.** Within each age band, columns 1–2 and 5–6 report mean minutes per day for boys and girls, respectively; columns 3–4 and 7–8 report participation rates. Bracketed values in the girls' columns are p-values for boy–girl differences in means. Entries are raw, unadjusted means.

TABLE 9: *Within-District Partial Autocorrelation of Annual Rainfall z-Scores*

.	**AR(1)**	**AR(2)**	**AR(3)**	**AR(4)**	**AR(5)**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Lag 1	0.072***	0.064***	0.063***	0.064***	0.061***
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
Lag 2	.	0.087***	0.087***	0.089***	0.083***
	.	(0.006)	(0.006)	(0.007)	(0.007)
Lag 3	.	.	0.077***	0.077***	0.087***
	.	.	(0.006)	(0.006)	(0.007)
Lag 4	.	.	.	-0.039***	-0.044***
	.	.	.	(0.007)	(0.007)
Lag 5	.	.	.	.	0.023***
	.	.	.	.	(0.007)
Observations	25,914	25,297	24,680	24,063	23,446
District FE	Y	Y	Y	Y	Y

Notes: **Sample.** Annual district-level rainfall z-scores from the GPCC Monitoring product, 1982–2024, mapped to the TUS 2019 adjusted district polygons (Section 3.2.1). **What it shows.** Each row reports the coefficient on lag k from a regression of the district rainfall z-score in year t on lags 1 through k. Thus, the reported coefficient is the partial association with the k-year lag conditional on shorter lags. **Estimation and inference.** Regressions include district fixed effects. Standard errors are clustered at the district level.

TABLE 10: *Domestic-Work Task Decomposition*

	All Domestic Work *(mins/day)*	Food Preparation *(mins/day)*	Cleaning *(mins/day)*	Domestic Production *(mins/day)*	Care Work *(mins/day)*	Budgeting and Shopping *(mins/day)*
	(1)	(2)	(3)	(4)	(5)	(6)
Infancy Rain Shock	1.0	1.3*	-0.5	-1.0**	0.5**	0.4***
	(1.3)	(0.8)	(0.5)	(0.5)	(0.3)	(0.1)
× *Boys*	-2.6*	-1.5*	0.1	-0.4	-0.2	-0.2
	(1.4)	(0.8)	(0.5)	(0.6)	(0.3)	(0.2)
Early Childhood Rain Shock	5.3***	2.9***	0.5	0.4	1.0***	0.4***
	(1.6)	(1.0)	(0.6)	(0.6)	(0.3)	(0.1)
× *Boys*	-4.3**	-2.9***	-0.5	-0.8	-0.4	0.2
	(1.7)	(1.0)	(0.6)	(0.8)	(0.3)	(0.2)
Observations	59,466	59,466	59,466	59,466	59,466	59,466
Girls Mean	89.2	41.7	25.4	13.4	5.3	1.5
Boys Mean	19.4	1.3	2.3	8.9	3.0	1.7
Boys share in sample	0.53	0.53	0.53	0.53	0.53	0.53
[Girls] p: Infancy =Early Ch.	[0.02]	[0.14]	[0.14]	[0.02]	[0.12]	[0.86]
[Boys] p: Infancy =Early Ch.	[0.00]	[0.05]	[0.00]	[0.05]	[0.18]	[0.01]
District-Age-Qtr-Gender FEs	Y	Y	Y	Y	Y	Y
Year-Gender FE	Y	Y	Y	Y	Y	Y

Notes: **Sample.** Main estimation sample: rural children aged 12–17, TUS 2019 and 2024 pooled, in the main 26 states; both genders. **What it shows.** Column 1 reports total domestic work. Columns 2–6 report food preparation, cleaning, domestic production, care of household members, and budgeting and shopping. These five sub-categories sum to total domestic work. Outcomes are diary minutes per day, winsorised separately by outcome at the state ×gender 99th percentile. **Estimation and inference.** Fixed effects are rural ×district ×age ×quarter ×gender and survey-year ×gender. The table reports outcome means separately for girls and boys. p-value rows test equality of the Infancy and Early Childhood coefficients within gender. Standard errors are clustered at district ×age ×survey-year and shown in parentheses. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 11: *Breadth of Domestic-Work Activities*

	N distinct ICATUS-3-digit codes within sub-category (cond on doer)					
	Total DW	**Food preparation**	**Cleaning**	**Domestic production**	**Care**	**Shopping**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Infancy Rain Shock	0.031	0.014	0.024	0.002	0.050	0.027
	(0.034)	(0.024)	(0.025)	(0.023)	(0.058)	(0.034)
× *Boys*	-0.064	-0.086	-0.031	-0.033	-0.051	-0.034
	(0.052)	(0.175)	(0.078)	(0.037)	(0.087)	(0.035)
Early Childhood Rain Shock	0.145***	0.049**	0.044*	0.004	-0.001	-0.050
	(0.035)	(0.022)	(0.025)	(0.022)	(0.036)	(0.058)
× *Boys*	-0.147***	-0.166	-0.090	-0.065*	-0.013	0.024
	(0.050)	(0.193)	(0.066)	(0.040)	(0.038)	(0.062)
Observations	20,844	11,042	10,970	6,698	2,920	1,830
Girls Mean	2.81	1.75	1.68	1.24	1.13	1.02
Boys Mean	1.42	1.30	1.21	1.17	1.09	1.01
Boys share in sample	0.29	0.08	0.15	0.40	0.41	0.59
[Girls] p: Infancy =Early Ch.	[0.006]	[0.212]	[0.548]	[0.914]	[0.352]	[0.327]
[Boys] p: Infancy =Early Ch.	[0.411]	[0.874]	[0.647]	[0.270]	[0.845]	[0.319]
District-Age-Qtr-Gender FEs	Y	Y	Y	Y	Y	Y
Year-Gender FE	Y	Y	Y	Y	Y	Y

Notes: **Sample.** Rural children aged 12–17, TUS 2019 and 2024 pooled, in the main 26 states; both genders. Each column includes children reporting any time in the relevant domestic-work category. **What it shows.** The number of distinct diary activities performed within each domestic-work category. Column 1 counts activities across all domestic work. **Estimation and inference.** Coefficients report the effect for girls; the ×Boys rows report the boys–girls differential. Regressions include rural ×district ×age ×quarter ×gender and survey-year ×gender fixed effects. Standard errors are clustered at district ×age ×survey-year and shown in parentheses. Bracketed entries are p-values. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 12: *Concentration of Domestic Work within Activity Categories*

	Share of sub-category minutes on the single most-used ICATUS code (cond on doer)					
	Total DW	**Food preparation**	**Cleaning**	**Domestic production**	**Care**	**Shopping**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Infancy Rain Shock	-0.004	-0.001	-0.010	0.001	-0.009	-0.008
	(0.005)	(0.006)	(0.007)	(0.007)	(0.015)	(0.010)
× *Boys*	0.011	0.002	0.016	0.009	0.032	0.010
	(0.009)	(0.057)	(0.024)	(0.012)	(0.028)	(0.010)
Early Childhood Rain Shock	-0.017***	-0.009	-0.011*	-0.002	0.000	0.014
	(0.005)	(0.006)	(0.007)	(0.006)	(0.011)	(0.016)
× *Boys*	0.016*	0.065	0.032	0.018	0.028	-0.006
	(0.009)	(0.056)	(0.024)	(0.011)	(0.022)	(0.018)
Observations	20,844	11,042	10,970	6,698	2,920	1,830
Girls Mean	0.653	0.803	0.794	0.928	0.962	0.994
Boys Mean	0.891	0.908	0.922	0.949	0.972	0.997
Boys share in sample	0.29	0.08	0.15	0.40	0.41	0.59
[Girls] p: Infancy =Early Ch.	[0.037]	[0.242]	[0.827]	[0.744]	[0.513]	[0.327]
[Boys] p: Infancy =Early Ch.	[0.386]	[0.526]	[0.567]	[0.512]	[0.733]	[0.319]
District-Age-Qtr-Gender FEs	Y	Y	Y	Y	Y	Y
Year-Gender FE	Y	Y	Y	Y	Y	Y

Notes: **Sample.** Identical to Table 11. **What it shows.** For each domestic-work category, the share of total category minutes allocated to the single activity with the most minutes. The outcome ranges from zero to one. **Estimation and inference.** The specification is the same as in Table 11. Standard errors are shown in parentheses and bracketed entries are p-values. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 13: *Early Childhood Shock and Other Time-Use Outcomes*

	Studying		Leisure and Self-care		Paid Work	
	(mins/day)	(≥360 mins)	(mins/day)	(≥960 mins)	(mins/day)	(≥60 mins)
	(1)	(2)	(3)	(4)	(5)	(6)
Infancy Rain Shock	1.8	0.005	-4.6**	-0.003	1.9***	0.005**
	(2.6)	(0.005)	(2.2)	(0.006)	(0.7)	(0.003)
× *Boys*	5.2	0.010	-2.3	-0.009	-0.0	-0.003
	(3.2)	(0.007)	(2.9)	(0.008)	(1.3)	(0.004)
Early Childhood Rain Shock	-4.7	-0.006	-0.9	0.002	0.3	0.003
	(2.9)	(0.006)	(2.5)	(0.006)	(0.8)	(0.003)
× *Boys*	1.3	0.003	0.4	-0.002	2.8*	0.003
	(3.5)	(0.008)	(3.0)	(0.008)	(1.6)	(0.005)
Observations	59,466	59,466	59,466	59,466	59,466	59,466
Girls Mean	345.0	0.606	991.0	0.519	11.9	0.056
Boys Mean	348.1	0.615	1,038.0	0.597	31.8	0.107
Boys share in sample	0.53	0.53	0.53	0.53	0.53	0.53
[Girls] p: Infancy =Early Ch.	[0.06]	[0.12]	[0.18]	[0.48]	[0.06]	[0.51]
[Boys] p: Infancy =Early Ch.	[0.00]	[0.00]	[0.02]	[0.08]	[0.42]	[0.27]
District-Age-Qtr-Gender FEs	Y	Y	Y	Y	Y	Y
Year-Gender FE	Y	Y	Y	Y	Y	Y

Notes: **Sample.** Main estimation frame: rural children aged 12–17, TUS 2019 and 2024 pooled, in the main 26 states; both genders. The child-roster condition is part of this frame, not an additional household-identification restriction. **What it shows.** Three outcome families at two margins each. Cols (1)–(2) studying; (3)–(4) leisure and self-care; (5)–(6) paid work. Within each pair, the first column is the intensive margin in minutes per day; the second is an indicator (study ≥360; leisure ≥960; paid ≥60). **Studying.** All formal study time in minutes per day. **Leisure.** Sleep, eating, personal care, passive media, and social contact summed across diary slots; winsorised state ×gender p99. **Paid work.** All paid work and unpaid help in household enterprise; winsorised state ×gender p99. **Infancy Shock.** Mean district rainfall z-score over ages −1 to 2, re-standardised within the estimation sample. **Early Childhood Shock.** Mean district rainfall z-score over ages 4 to 7, re-standardised within the estimation sample. **Estimation.** Gender-interacted shocks; fixed effects rural ×district ×age ×quarter ×gender, plus survey-year ×gender; clustered at district ×age ×survey-year. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 14: *Controls and Fixed-Effect Build-Up*

	Baseline	Main	+ Parental Controls	+ HH Controls	+ Composition	×Gender
	(1)	(2)	(3)	(4)	(5)	(6)
Infancy Rain Shock	-1.97**	0.99	1.00	1.07	0.81	0.53
	(0.99)	(1.34)	(1.34)	(1.33)	(1.33)	(1.33)
× *Boys*	4.09***	-2.58*	-2.47*	-2.25	-2.15	-1.90
	(1.24)	(1.38)	(1.38)	(1.37)	(1.37)	(1.37)
Early Childhood Rain Shock	9.41***	5.32***	5.39***	5.23***	5.30***	5.33***
	(1.11)	(1.64)	(1.64)	(1.64)	(1.62)	(1.61)
× *Boys*	-12.66***	-4.33**	-4.32**	-4.07**	-4.15**	-4.26**
	(1.22)	(1.71)	(1.71)	(1.71)	(1.70)	(1.67)
Observations	59,466	59,466	59,466	59,148	59,148	59,148
Girls Mean	88.2	89.2	89.2	89.2	89.2	89.2
Boys Mean	19.4	19.4	19.4	19.3	19.3	19.3
Boys share in sample	0.53	0.53	0.53	0.53	0.53	0.53
[Girls] p: Infancy =Early Ch.	[0.00]	[0.02]	[0.01]	[0.02]	[0.01]	[0.01]
[Boys] p: Infancy =Early Ch.	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]
District-Age-Qtr FE	Y	.	.	.	.	.
Year FE	Y	.	.	.	.	.
District-Age-Qtr-Gender FE	.	Y	Y	Y	Y	Y
Year-Gender FE	.	Y	Y	Y	Y	Y
Parental Controls (Educ & Age)	.	.	Y	Y	Y	Y
HH Controls (Land, Caste, Religion)	.	.	.	Y	Y	Y
Composition (HH Size, N Children)	.	.	.	.	Y	Y
All Controls ×Gender	.	.	.	.	.	Y

Notes: **Sample.** Main estimation sample: rural children aged 12–17, TUS 2019 and 2024 pooled, in the main 26 states; both genders. The outcome is domestic work in minutes per day, winsorised at the state ×gender 99th percentile. **What it shows.** Six specifications that add fixed effects and covariates sequentially. Column 1 includes district ×age ×quarter and survey-year fixed effects. Column 2 adds gender to both fixed-effect dimensions. Columns 3–6 successively add parental characteristics, household assets, household composition, and gender interactions with all controls. **Estimation and inference.** Standard errors are clustered at district ×age ×survey-year and shown in parentheses. Bracketed values are p-values for equality of the Infancy and Early Childhood coefficients within gender. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 15: *Sample Variants and Within-Household Identification*

	Domestic Work (*mins/day*)			
	Both Parents ID'd	**Eldest Only**	**+ HH ×G FE**	**Y×Q×A + D FE**
	(1)	**(2)**	**(3)**	**(4)**
Infancy Rain Shock	-0.1	2.4	6.6	-4.3***
	(1.6)	(1.9)	(6.3)	(0.9)
× *Boys*	-1.0	-4.0**	-7.4	4.4***
	(1.7)	(1.9)	(7.1)	(1.1)
Early Childhood Rain Shock	6.3***	6.4***	4.0	7.3***
	(1.9)	(2.2)	(6.2)	(0.9)
× *Boys*	-5.6***	-5.0**	-9.2	-12.8***
	(2.0)	(2.2)	(7.2)	(1.1)
Observations	48,382	40,033	59,466	59,466
Girls Mean	85.2	103.6	91.4	87.7
Boys Mean	18.2	21.1	22.2	19.3
Boys share in sample	0.53	0.53	0.53	0.53
[Girls] p: Infancy =Early Ch.	[0.00]	[0.09]	[0.68]	[0.00]
[Boys] p: Infancy =Early Ch.	[0.01]	[0.00]	[0.16]	[0.00]
District-Age-Qtr-Gender FEs	Y	Y	Y	.
Year-Gender FE	Y	Y	Y	.
HH-Gender FE	.	.	Y	.
District FE	.	.	.	Y
Year-Qtr-Age FE	.	.	.	Y

Notes: **Sample.** Main estimation frame: rural children aged 12–17, TUS 2019 and 2024 pooled, in the main 26 states; both genders. The outcome is domestic work in minutes per day, winsorised at the state ×gender 99th percentile. Columns 1–2 use family-linked subsamples; columns 3–4 use the full main sample. **What it shows.** Column 1 retains children whose two parents are uniquely linked to household-roster members. Column 2 retains one focal observation per family-linked household: the eldest child aged under 18. Column 3 uses the full main sample with household ×gender fixed effects. Column 4 uses the full main sample with district and survey-year ×quarter ×age fixed effects. **Estimation and inference.** Standard errors are clustered at district ×age ×survey-year and shown in parentheses. Bracketed values are p-values for equality of the Infancy and Early Childhood coefficients within gender. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 16: *Alternative Rainfall-Shock Constructions*

	Aggregation (annual)				Season		Coding
	GPCM Buffer (Baseline)	GPCM Inside	GPCM Radial	GPCC Full (Inside)	Monsoon (Apr–Sep)	Post-mons. (Oct–Dec)	Linear tail coding
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Infancy Rain Shock	1.0	-2.1	0.3	1.3	1.2	2.7**	-0.7
	(1.3)	(1.8)	(1.3)	(1.4)	(1.4)	(1.2)	(1.4)
× *Boys*	-2.6*	0.1	-2.1	-2.4*	-2.5*	-2.5*	-1.1
	(1.4)	(1.8)	(1.4)	(1.5)	(1.5)	(1.3)	(1.5)
Early Childhood Rain Shock	5.3***	5.8***	5.7***	5.2***	6.5***	1.3	3.6**
	(1.6)	(2.2)	(1.6)	(1.8)	(1.6)	(1.4)	(1.6)
× *Boys*	-4.3**	-4.3*	-4.9***	-4.3**	-5.1***	-0.7	-2.6
	(1.7)	(2.3)	(1.7)	(1.8)	(1.7)	(1.5)	(1.6)
Observations	59,466	26,792	61,517	61,459	59,466	59,466	59,466
Girls Mean	89.2	92.9	88.3	88.4	89.2	89.2	89.2
Boys Mean	19.4	20.0	19.4	19.4	19.4	19.4	19.4
Boys share in sample	0.53	0.53	0.53	0.53	0.53	0.53	0.53
[Girls] p: Infancy =Early Ch.	[0.02]	[0.00]	[0.00]	[0.03]	[0.00]	[0.31]	[0.02]
[Boys] p: Infancy =Early Ch.	[0.00]	[0.00]	[0.00]	[0.00]	[0.00]	[0.45]	[0.00]
District-Age-Qtr-Gender FEs	Y	Y	Y	Y	Y	Y	Y
Year-Gender FE	Y	Y	Y	Y	Y	Y	Y

Notes: **Sample.** Main estimation sample: rural children aged 12–17, TUS 2019 and 2024 pooled, in the main 26 states; both genders. The outcome is domestic work in minutes per day, winsorised at the state ×gender 99th percentile. Observation counts differ across columns because rainfall products cover different sets of districts. **What it shows.** The baseline regression under alternative rainfall-shock constructions. Columns 1–4 vary the rainfall product and district-grid-cell aggregation rule. Columns 5–6 construct shocks from April–September and October–December rainfall, respectively. Column 7 replaces continuous rainfall z-scores with a window average of annual indicators equal to +1 in the top rainfall quintile, −1 in the bottom quintile, and zero otherwise. **Rainfall shocks.** Infancy is average rainfall from ages −1 to 2; Early Childhood is average rainfall from ages 4–7. Each shock construction is re-standardised within its estimation sample. **Estimation and inference.** Coefficients report the effect for girls; the ×Boys rows report the boys–girls differential. Regressions include rural ×district ×age ×quarter ×gender and survey-year ×gender fixed effects. Bracketed values are p-values for equality of the Infancy and Early Childhood coefficients within gender. Standard errors are clustered at district ×age ×survey-year and shown in parentheses. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 17: *Agricultural First Stage: Rainfall to Output (ICRISAT District Panel)*

	Production ***(1000 tons)***	**Production** ***(log)***	**Yield** ***(tons/ha)***	**Yield** ***(log)***
	(1)	**(2)**	**(3)**	**(4)**
Concurrent Rainfall	30.19***	0.075***	0.04***	0.047***
	(3.08)	(0.005)	(0.00)	(0.004)
Lagged Rainfall	5.62***	0.013***	-0.01**	-0.005
	(2.15)	(0.004)	(0.00)	(0.003)
Observations	13,070	13,022	13,041	13,021
Outcome Mean	627.33	5.979	1.50	0.229
District FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y

Notes: **Sample.** ICRISAT district ×year agricultural panel, 1999–2017. **What it shows.** Agricultural production in thousand tons (col 1) and in logs (col 2), and yield in tons per hectare (col 3) and in logs (col 4); the log is taken on positive values (84 non-positive district-years drop for production, 20 for yield). **Concurrent and lagged rainfall.** A concurrent and a one-year-lagged standardised district rainfall shock (calendar-year annual z-scores on this panel), each re-standardised within the estimation sample. **Estimation.** District and year fixed effects; coefficients are per one-standard-deviation increase. **Inference.** Standard errors clustered at the district level, shown in parentheses; $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 18: *Household Food Quantities and Rainfall (HCES, 2009–11)*

	To-be-cooked *(kg)*			**Outside and ready-made *(Rs)***		
	Ingredients *(log)*	**Staples *(log)***	**Staple grains *(share)***	**Packed *(asinh)***	**Served *(asinh)***	**Meals away *(count)***
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Concurrent Rainfall	0.009***	0.008***	-0.000	-0.074***	-0.168***	-0.60***
	(0.003)	(0.003)	(0.001)	(0.016)	(0.020)	(0.12)
Lagged Rainfall	-0.009**	-0.007**	0.002	-0.077***	-0.269***	-0.69***
	(0.004)	(0.003)	(0.001)	(0.018)	(0.025)	(0.15)
Observations	98,700	98,532	98,700	99,033	99,033	99,033
Outcome Mean	4.733	4.697	0.440	2.351	1.906	12.53
District ×Quarter FE	Y	Y	Y	Y	Y	Y
Survey Round FE	Y	Y	Y	Y	Y	Y

Notes: **Sample.** HCES, all rural households in the main 26 states, 2009–11 rounds only, one observation per household. **What it shows.** Six columns in two groups. *To-be-cooked, in kilograms (cols 1–3):* food-ingredient quantities in logs (col 1); staples quantities in logs (col 2), where staples =cereals +cereal substitutes +pulses +edible oil +vegetables +salt/sugar +spices +milk products, 72.4% of household food spending in this sample; and the staple-grain share (col 3: kilograms of cereals +cereal substitutes over total kilogram-measured food, in $[0, 1]$, baseline mean 0.44). *Ready-made (cols 4–6):* packed processed food (col 4) and served processed food (col 5) in rupees, inverse hyperbolic sine — the schedule records no quantities for these items — and the count of household meals taken outside the home (col 6, untransformed; 61.5% of households report zero). Log columns condition on positive values (0.33% and 0.51% of households drop); quantity and value series are winsorised once, producer-side. **Concurrent and lagged rainfall.** As in Table 3, merged at the household's district and survey quarter; re-standardised in-sample. **Estimation.** District-quarter and survey-round fixed effects; coefficients are per one-standard-deviation increase. **Inference.** Standard errors clustered at district ×survey-round, shown in parentheses; $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 19: *Plenty versus Drought: Tail Decomposition of the Life-Stage Shocks*

	Domestic Work		Studying	
	(mins/day)	≥ **120** *(share)*	*(mins/day)*	≥ **360** *(share)*
	(1)	**(2)**	**(3)**	**(4)**
Infancy: Drought	-11.6**	-0.025	9.4	0.017
	(4.8)	(0.016)	(8.5)	(0.018)
× *Boys*	14.8***	0.040**	-28.4***	-0.056**
	(5.0)	(0.018)	(10.1)	(0.022)
Infancy: Plenty	3.3	0.020	-11.2	-0.028
	(4.3)	(0.017)	(8.8)	(0.019)
× *Boys*	-5.6	-0.026	18.7*	0.044*
	(4.7)	(0.018)	(10.6)	(0.023)
Early Childhood: Drought	0.0	0.013	-9.1	-0.027
	(4.2)	(0.015)	(8.7)	(0.018)
× *Boys*	0.3	-0.004	7.0	0.007
	(4.1)	(0.016)	(10.1)	(0.022)
Early Childhood: Plenty	25.1***	0.075***	-37.1***	-0.074***
	(5.5)	(0.020)	(10.1)	(0.022)
× *Boys*	-22.8***	-0.072***	17.0	0.034
	(6.0)	(0.023)	(12.0)	(0.026)
Observations	59,466	59,466	59,466	59,466
Girls Mean	89.2	0.298	345.0	0.606
Boys Mean	19.4	0.064	348.1	0.615
Boys share in sample	0.53	0.53	0.53	0.53
[Girls, Plenty] p: Infancy =Early Ch.	[0.00]	[0.00]	[0.00]	[0.01]
[Girls, Drought] p: Infancy =Early Ch.	[0.03]	[0.05]	[0.10]	[0.07]
[Boys, Plenty] p: Infancy =Early Ch.	[0.01]	[0.27]	[0.00]	[0.00]
[Boys, Drought] p: Infancy =Early Ch.	[0.15]	[0.45]	[0.09]	[0.36]
District-Age-Qtr-Gender FEs	Y	Y	Y	Y
Year-Gender FE	Y	Y	Y	Y

Notes: **Sample.** Rural children aged 12–17, TUS 2019 + 2024 pooled, the main 26 states, both genders. **What it shows.** Each life-stage shock is recoded into two binary tail indicators omitting the middle band. Cols (1)–(2): domestic work in minutes per day and an indicator for ≥120 minutes. Cols (3)–(4): studying in minutes per day and an indicator for ≥360 minutes. **Plenty / Drought.** *Plenty* =1 if the standardised life-stage rainfall z-score is at or above +1.15; *Drought* =1 if it is at or below −1.15; the middle band ($|z|$ <1.15) is the omitted reference. The ±1.15 threshold places roughly the top and bottom 14% of district-window observations in each tail. Applied to both the Infancy and Early Childhood windows. **Estimation.** Gender-interacted; headline fixed-effect structure. **Inference.** Standard errors clustered at district ×age ×survey-year, shown in parentheses; $***p<0.01$, $**p<0.05$, $*p<0.10$. Bracketed values in the lower panel are p-values testing, within each gender and tail (plenty / drought), whether the Infancy and Early Childhood coefficients are equal.

TABLE 20: *Contemporaneous and Lagged Rainfall Effects by Daughter Rank*

	Eldest vs. Younger Daughters — minutes per day					
	Core Domestic Work		**Other DW**		**Paid Work (Total)**	
	(Unweighted)	***(Reweighted)***	***(Unweighted)***	***(Reweighted)***	***(Unweighted)***	***(Reweighted)***
	(1)	**(2)**	**(3)**	**(4)**	**(5)**	**(6)**
Concurrent Rain	5.64***	5.64***	0.14	0.14	0.29	0.29
	(1.33)	(1.33)	(0.75)	(0.75)	(0.77)	(0.77)
× *Younger*	-5.97**	-9.06**	1.24	3.29	-0.40	0.86
	(2.69)	(3.71)	(2.00)	(2.33)	(1.40)	(2.07)
Lagged Rain (t−1)	2.17	2.18	0.09	0.09	0.43	0.43
	(1.35)	(1.35)	(0.89)	(0.89)	(0.90)	(0.90)
× *Younger*	2.35	1.40	1.52	2.67	5.77*	10.89*
	(2.69)	(3.47)	(2.10)	(2.75)	(3.31)	(5.77)
Observations	26,781	26,781	26,781	26,781	26,781	26,781
Eldest Mean	75.7	75.7	21.8	21.8	11.9	11.9
Younger Mean	26.7	36.2	16.4	18.3	6.3	9.3
Younger share in sample	0.16	0.16	0.16	0.16	0.16	0.16
[Eldest] p: Concurrent =Lagged	[0.08]	[0.08]	[0.97]	[0.97]	[0.89]	[0.89]
[Younger] p: Concurrent =Lagged	[0.20]	[0.15]	[0.93]	[0.85]	[0.03]	[0.04]
District-Age-Qtr-Eldest FE	Y	Y	Y	Y	Y	Y
Year-Eldest FE	Y	Y	Y	Y	Y	Y
Overall Sibling Rank FE	Y	Y	Y	Y	Y	Y

Notes: **Sample.** Rural girls aged 12–17, TUS 2019 and 2024 pooled, in the main 26 states, with parent–child relationships uniquely identifiable in the household roster. The table compares eldest daughters (the firstborn girl in the household, including girls with older brothers) with younger sisters. **What it shows.** Core domestic work, other domestic work, and total paid work, each in minutes per day. Odd columns are unweighted; even columns reweight younger sisters to the eldest daughters' age distribution within ages 12–17. **Rainfall.** Current rainfall is the quarter-aligned rolling 12-month district rainfall z-score ending in the diary quarter; lagged rainfall is the corresponding window ending one year earlier. Both are re-standardised within the estimation sample. **Estimation and inference.** Each outcome is regressed on current and lagged rainfall and their interactions with a younger-sister indicator. Base coefficients apply to eldest daughters; the ×Younger rows are the younger–eldest differences. Fixed effects are district ×age ×quarter ×eldest status and survey-year ×eldest status, with an overall sibling-rank control. Standard errors are clustered at district ×age ×survey-year and shown in parentheses. Bracketed rows test equality of current and lagged rainfall effects within eldest daughters and within younger sisters. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 21: *Study and Leisure by Daughter Rank and Day Type*

	Studying		Leisure	
	Weekday	*Sunday*	*Weekday*	*Sunday*
	(1)	(2)	(3)	(4)
Infancy Rain Shock	1.68	-1.17	-4.63*	-1.89
	(3.00)	(5.66)	(2.47)	(5.35)
× *Younger Sister*	12.83	8.92	-10.58	-1.72
	(9.42)	(23.80)	(8.70)	(20.32)
Early Childhood Rain Shock	-8.07**	13.16**	1.90	-19.97***
	(3.27)	(5.71)	(2.76)	(5.51)
× *Younger Sister*	12.50	-13.78	-5.54	17.05
	(11.00)	(21.97)	(9.62)	(19.73)
Observations	26,780	26,780	26,780	26,780
Eldest Mean	362.6	165.4	967.4	1115.8
Younger Mean (reweighted)	380.7	145.5	999.8	1165.4
p: Infancy Shock =Early Childhood Shock [Eldest]	[0.01]	[0.05]	[0.05]	[0.01]
p: Infancy Shock =Early Childhood Shock [Younger Sister]	[0.42]	[0.80]	[0.30]	[0.98]
District-Age-Qtr-Eldest FE	Y	Y	Y	Y
Year-Eldest FE	Y	Y	Y	Y

Notes: **Sample.** Rural daughters aged 12–17, TUS 2019 + 2024 pooled, main 26 states, households with identifiable parent–child relationships; all diary days (Mon–Sat normal and non-normal, plus Sunday). Eldest daughters together with their younger sisters, with younger-sister observations reweighted to the eldest's age distribution (the Table 4-Col-2 reweighting, built within this table's estimation sample; eldest weight =1). **What it shows.** Two outcomes (Studying and Leisure) crossed with two day-types (Weekday and Sunday). All four cells per outcome come from a single triple-interaction regression on that outcome; the Sunday columns are lincom extractions of (base +Sunday differential) from the same regression as the Weekday columns. Within each outcome panel: the Infancy / Early Childhood rows report the effect on eldest daughters; the indented ×Younger Sister rows report the within-household differential (younger minus eldest). **Studying.** All formal study time in minutes per day; winsorised state ×female p99. **Leisure.** Sleep, eating, personal care, passive media, and social contact summed across diary slots; winsorised state ×female p99. **Estimation.** For each outcome, a single regression interacts both shocks with the younger-sister indicator and the Sunday indicator (three-way interactions in both shocks). The Sunday indicator enters as a regressor; the eldest dimension is saturated into the fixed effects. Younger-sister observations are reweighted to the eldest's age distribution. Fixed effects: district-age-quarter-eldest, plus survey-year-eldest. Standard errors clustered at district ×age ×survey-year. **Inference.** The triple-interaction p-value row tests whether the eldest-vs-younger differential differs between Weekday and Sunday. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 22: *Rainfall Effects by Son Rank on Domestic Work, Paid Work, and Studying*

	Eldest vs. Younger Sons — minutes per day					
	Domestic Work		Paid Work (Total)		Studying	
	(Unweighted)	*(Reweighted)*	*(Unweighted)*	*(Reweighted)*	*(Unweighted)*	*(Reweighted)*
	(1)	(2)	(3)	(4)	(5)	(6)
Infancy Rain Shock	-1.66***	-1.66***	2.93**	2.93**	4.77*	4.76*
	(0.64)	(0.64)	(1.38)	(1.38)	(2.74)	(2.74)
× *Younger*	-0.28	-3.62	-1.14	7.16	1.63	-3.79
	(2.15)	(3.56)	(2.83)	(6.97)	(8.47)	(10.15)
Early Childhood Rain Shock	1.61**	1.61**	3.08*	3.07*	-4.56	-4.54
	(0.73)	(0.73)	(1.77)	(1.77)	(3.06)	(3.06)
× *Younger*	-2.71	-11.52*	1.71	5.18	-1.10	0.51
	(2.37)	(6.91)	(4.18)	(6.10)	(10.98)	(13.50)
Observations	30,151	30,151	30,151	30,151	30,151	30,151
Eldest Mean	19.4	19.4	34.5	34.5	346.7	346.7
Younger Mean	16.2	19.7	15.5	24.4	360.7	351.6
Younger share in sample	0.15	0.15	0.15	0.15	0.15	0.15
[Eldest] p: Infancy =Early Ch.	[0.00]	[0.00]	[0.93]	[0.94]	[0.01]	[0.01]
[Younger] p: Infancy =Early Ch.	[0.78]	[0.43]	[0.49]	[0.84]	[0.24]	[0.70]
District-Age-Qtr-Eldest FE	Y	Y	Y	Y	Y	Y
Year-Eldest FE	Y	Y	Y	Y	Y	Y
Overall Sibling Rank FE	Y	Y	Y	Y	Y	Y

Notes: **Sample.** Rural boys aged 12–17, TUS 2019 + 2024 pooled, main 26 states, with parent–child links identified from the household roster. The table compares eldest sons (the firstborn boy, including boys with older sisters) with younger sons. **Outcomes.** Domestic work, total paid work, and studying, each in minutes per day. Odd columns are unweighted; even columns reweight younger sons to match the age distribution of eldest sons within ages 12–17. **Estimation.** Each outcome is regressed on the Infancy and Early Childhood rainfall shocks and their interactions with a younger-son indicator. The base coefficients apply to eldest sons; the ×Younger rows are the difference between younger- and eldest-son effects. Fixed effects: district ×age ×quarter ×eldest-son status and survey-year ×eldest-son status, with a control for overall sibling rank. Standard errors clustered at district ×age ×survey-year. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 23: *Eldest-Daughter Effect: Sample Variations*

	Domestic work — sample variations			
	Mothers Age ≤40	**Low Son Pref. Districts**	**Mothers Age ≤40**	**Low Son Pref. Districts**
	(1)	**(2)**	**(3)**	**(4)**
Infancy Rain Shock	2.0	4.4*	1.8	3.5
	(2.0)	(2.2)	(2.1)	(2.2)
× *Younger*	-4.0	-2.8	-5.8	-2.0
	(4.5)	(4.9)	(3.8)	(4.5)
Early Childhood Rain Shock	4.1*	8.5***	5.0**	7.3***
	(2.2)	(2.7)	(2.3)	(2.6)
× *Younger*	-4.9	-13.7**	-3.8	-6.6
	(5.0)	(5.6)	(4.3)	(4.4)
Observations	17,395	13,467	17,395	13,467
Eldest Mean	84.4	97.5	85.6	99.2
Younger Mean	36.8	45.7	37.3	46.4
Younger share in sample	0.16	0.16	0.16	0.16
[Eldest] p: Infancy =Early Ch.	[0.39]	[0.15]	[0.22]	[0.19]
[Younger] p: Infancy =Early Ch.	[0.80]	[0.32]	[0.34]	[0.91]
District-Age-Qtr-Eldest FE	Y	Y	.	.
Year-Eldest FE	Y	Y	.	.
Overall Sibling Rank FE	Y	Y	.	.
District-Age-Qtr-Rank FE	.	.	Y	Y
Year-Rank FE	.	.	Y	Y

Notes: **Sample.** Rural girls aged 12–17, TUS 2019 + 2024 pooled, main 26 states, households with identifiable parent–child relationships. Eldest daughters vs younger sisters under two sample restrictions: mother aged ≤40 (cols 1, 3) and low-son-preference districts (cols 2, 4). **What it shows.** Two sample restrictions crossed with two identification strategies. Cols (1)–(2) panel-by-eldest design (district ×age ×quarter ×eldest + survey-year ×eldest, plus a birth-rank control). Cols (3)–(4) cross-sectional rank design (district ×age ×quarter ×birth-rank + survey-year ×birth-rank). **Domestic work.** Total domestic work in minutes per day, as defined in the paper; winsorised state ×gender p99. **Eldest daughter.** The household's firstborn girl, even with older brothers. **Low son-preference districts.** Lower half of the district-level son-preference index built from PLFS-derived indicators. **Estimation.** Each regression interacts both shocks with a younger-sibling indicator; coefficients on Eldest are the base effect and the ×Younger Sister row is the within-household differential. Clustered at district ×age ×survey-year. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 24: *Eldest Daughters' Domestic Work by Younger-Sibling Presence*

	Subgroup: eldest daughter has. . .				
	a sibling ≥4y younger	**a sibling ≥6y younger**	**any younger sibling**	**≥2 younger siblings**	**≥3 younger siblings**
Infancy Rain Shock	1.2	1.2	0.1	1.7	1.1
	(2.0)	(1.9)	(2.4)	(1.9)	(1.8)
× *Subgroup*	0.3	0.5	1.8	-0.5	1.7
	(2.4)	(2.6)	(2.5)	(2.7)	(3.8)
Early Childhood Rain Shock	6.9***	6.6***	6.2**	7.0***	6.0***
	(2.3)	(2.2)	(2.6)	(2.2)	(2.0)
× *Subgroup*	-1.0	-0.4	0.3	-1.5	1.2
	(2.6)	(2.7)	(2.6)	(2.8)	(4.0)
Observations	20,653	20,653	20,653	20,653	20,653
Reference Mean	89.9	93.2	86.6	86.8	90.7
Subgroup Mean	106.9	107.4	102.8	118.1	138.7
Subgroup share in sample	0.45	0.31	0.67	0.34	0.14
[Reference] p: Infancy =Early Ch.	[0.03]	[0.03]	[0.05]	[0.03]	[0.03]
[Subgroup] p: Infancy =Early Ch.	[0.12]	[0.17]	[0.07]	[0.21]	[0.39]
District-Age-Qtr FEs	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y

Notes: **Sample.** Single-family eldest daughters aged 12–17 (one nuclear family in the diary roster, so younger roster members are biological siblings), TUS 2019 + 2024 pooled, main 26 states, all diary days. **What it shows.** The Early Childhood Shock's effect on the eldest daughter's own domestic work (minutes per day), split by whether her household contains a younger sibling of a given type. Five columns, one regression each: a sibling at least 4 years younger; at least 6 years younger; any younger sibling; at least 2 younger siblings; at least 3 younger siblings. A sibling ≥ 4 (≥ 6) years younger was born when the focal daughter was already 4 (6) — during or after her pre-school window — so columns 1–2 are the siblings whose existence could be a fertility response to the shock; "any younger sibling" (col 3) also counts siblings born before the window and is the weaker test. The shock row reports the effect in the *reference* group (no such sibling); the ×Subgroup row reports how that effect differs when such a sibling is present. The Reference / Subgroup baseline-mean rows are descriptive within-group means of domestic work. **Domestic work.** Total domestic work in minutes per day, as defined in the paper; winsorised state ×gender p99. **Estimation.** Fixed effects rural ×district ×age ×quarter, plus survey-year; clustered at district ×age ×survey-year. $***p<0.01$, $**p<0.05$, $*p<0.10$. Bracketed values are p-values for the test that, within the reference and the subgroup, the Infancy and Early Childhood coefficients are equal.

TABLE 25: *Childhood Rainfall and Younger-Sibling Presence*

	Outcome: 1{**household has . . .**}				
	a sibling ≥4y younger	**a sibling ≥6y younger**	**any younger sibling**	**≥2 younger siblings**	**≥3 younger siblings**
Panel A: Eldest daughters					
Infancy Rain Shock	-0.017**	-0.019***	-0.022***	-0.022***	-0.006
	(0.007)	(0.007)	(0.007)	(0.006)	(0.005)
Early Childhood Rain Shock	0.008	-0.001	0.009	0.002	0.009*
	(0.008)	(0.007)	(0.007)	(0.007)	(0.006)
Observations	20,653	20,653	20,653	20,653	20,653
Outcome Mean	0.460	0.317	0.682	0.348	0.145
p: Infancy =Early Ch.	[0.002]	[0.021]	[0.000]	[0.001]	[0.007]
Panel B: Eldest sons					
Infancy Rain Shock	-0.006	-0.003	0.002	-0.005	-0.005
	(0.006)	(0.005)	(0.006)	(0.005)	(0.003)
Early Childhood Rain Shock	-0.000	0.007	0.008	0.000	0.003
	(0.006)	(0.006)	(0.007)	(0.006)	(0.004)
Observations	23,826	23,826	23,826	23,826	23,826
Outcome Mean	0.356	0.234	0.599	0.242	0.083
p: Infancy =Early Ch.	[0.367]	[0.148]	[0.441]	[0.413]	[0.028]
District-Age-Qtr FEs	Y	Y	Y	Y	Y
Year FEs	Y	Y	Y	Y	Y

Notes: **Sample.** Single-family households, focal child aged 12–17. Panel A: focal is the eldest daughter. Panel B: focal is the eldest son. **What it shows.** The Infancy and Early Childhood Shocks' effects on five indicators for the presence of a younger sibling, as dependent variables (the same five definitions as Table 24): a sibling ≥ 4y younger; ≥ 6y younger; any younger; ≥ 2 younger; ≥ 3 younger. Tests whether the shock changes the household's completed fertility, in particular whether it produces a sibling born after the focal child's pre-school window (columns 1–2). **Estimation.** Fixed effects rural ×district ×age ×quarter, plus survey-year; clustered at district ×age ×survey-year. $***p<0.01$, $**p<0.05$, $*p<0.10$. Bracketed values are p-values for the test that, within each panel, the Infancy and Early Childhood coefficients are equal.

TABLE 26: *Enrolment and Achievement (TUS and ASER)*

	TUS			ASER	
	Study ≥ 120 =1 ***(total educ. time)***	**On Track =1** ***(by age)***	**≥ Secondary =1** ***(share)***	**Enrolled =1** ***(share)***	**On Track =1** ***(by age)***
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Infancy Rain Shock	0.003	0.001	0.003	0.006***	0.004***
	(0.005)	(0.004)	(0.003)	(0.001)	(0.001)
× *Boys*	-0.000	0.005	0.004	-0.002***	-0.000
	(0.006)	(0.006)	(0.005)	(0.001)	(0.001)
Early Childhood Rain Shock	-0.007	0.009*	0.015***	0.008***	0.004***
	(0.006)	(0.005)	(0.004)	(0.001)	(0.001)
× *Boys*	-0.004	-0.012*	-0.005	-0.005***	-0.002
	(0.007)	(0.006)	(0.005)	(0.001)	(0.001)
Observations	59,466	59,466	59,462	392,893	392,893
Girls Mean	0.749	0.754	0.228	0.942	0.806
Boys Mean	0.756	0.751	0.219	0.950	0.780
Boys share in sample	0.53	0.53	0.53	0.50	0.50
[Girls] p: Infancy =Early Ch.	[0.08]	[0.14]	[0.00]	[0.02]	[0.72]
[Boys] p: Infancy =Early Ch.	[0.01]	[0.12]	[0.44]	[0.95]	[0.26]
District-Age-Qtr-Gender FE	Y	Y	Y	Y	Y
Year-Gender FE	Y	Y	Y	Y	Y

Notes: **Sample.** *Panel A (TUS, cols 1–3):* rural children aged 12–17, TUS 2019 + 2024 pooled, main 26 states, both genders, child-roster restriction (no household-identification filter). *Panel B (ASER, cols 4–5):* ASER 2018 + 2022 rounds, rural children aged 12–16, main 26 states, all households, both genders. **What it shows.** *Panel A:* (1) total education time ≥120 minutes per day; (2) on-track for age; (3) attained at least secondary. *Panel B:* (1) currently enrolled; (2) on-track for age. Each panel's sample and observation count are reported in its own row. **Studying ≥120 (TUS, col 1).** Indicator on broad formal-study minutes (includes commute-to-school). **On track for age (ASER and TUS).** Currently enrolled in or above the modal class for the child's age. **Enrolled (ASER, col 4).** Currently in school. **Estimation.** Gender-interacted shocks; *Panel A* fixed effects rural ×district ×age ×quarter ×gender + survey-year ×gender; *Panel B* fixed effects district ×age ×gender + survey-year ×gender (no quarter). Clustered at district ×age ×survey-year. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 27: *ASER Foundational Learning*

	ASER Foundational Learning		
	Passes Grade I reading + math	**Passes Grade I both + any Grade II**	**Passes Grade II reading + math**
	(1)	**(2)**	**(3)**
Infancy Rain Shock	0.0062***	0.0108***	0.0088***
	(0.0017)	(0.0017)	(0.0015)
× *Boys*	-0.0003	-0.0020	-0.0021
	(0.0018)	(0.0018)	(0.0017)
Early Childhood Rain Shock	0.0126***	0.0156***	0.0137***
	(0.0015)	(0.0015)	(0.0014)
× *Boys*	-0.0041**	-0.0049***	-0.0079***
	(0.0017)	(0.0017)	(0.0016)
Observations	392,893	392,893	392,893
Girls Mean	0.5271	0.4931	0.3287
Boys Mean	0.4939	0.4574	0.3147
Boys share in sample	0.50	0.50	0.50
[Girls] p: Infancy =Early Ch.	[0.00]	[0.02]	[0.01]
[Boys] p: Infancy =Early Ch.	[0.21]	[0.40]	[0.61]
District-Age-Qtr-Gender FE	Y	Y	Y
Year-Gender FE	Y	Y	Y

Notes: **Sample.** ASER 2018 + 2022 rounds, rural children aged 12–16, main 26 states, all households, both genders; ASER tests every child regardless of enrolment. Missing scores recoded to zero. **What it shows.** Foundational reading-and-mathematics rungs: (1) both at Standard-I; (2) at least one at Standard-II plus both Standard-I; (3) both at Standard-II. **Foundational learning rungs (ASER).** Indicator for clearing both Reading and Math at the relevant standard. **Estimation.** Gender-interacted shocks; fixed effects district ×age ×gender + survey-year ×gender; clustered at district ×age ×survey-year. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 28: *Health Outcomes (NFHS), Rural Girls Aged 15–17*

	Not Sev. Thin	Not Anemic	Not Low BP	Not Stunted	Health Index
	(1)	(2)	(3)	(4)	(5)
Infancy Rain Shock	0.004**	0.007**	0.008***	-0.006**	0.004***
	(0.002)	(0.003)	(0.002)	(0.002)	(0.001)
Early Childhood Rain Shock	0.002	0.025***	-0.004**	0.008***	0.007***
	(0.002)	(0.003)	(0.002)	(0.002)	(0.001)
Observations	106,654	106,034	100,287	106,651	99,331
Outcome Mean	0.807	0.428	0.830	0.564	0.658
p: Infancy =Early Ch.	[0.45]	[0.00]	[0.00]	[0.00]	[0.04]
District-Age FE	Y	Y	Y	Y	Y
Survey-Year FE	Y	Y	Y	Y	Y

Notes: **Sample.** NFHS 2015–2021, rural girls aged 15–17 in the main 26 states. Module samples vary by outcome. **What it shows.** Effects of the Infancy and Early Childhood Rainfall Shocks on indicators for BMI at least 17, haemoglobin above the WHO anaemia threshold, blood pressure not classified as low, not being stunted, and a four-component index formed from these indicators. **Estimation and inference.** Regressions include district ×birth-cohort and survey-year fixed effects. Standard errors are clustered at district ×birth-cohort ×survey-year and shown in parentheses. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 29: *Time Spent Outside the Home (Non-Work, Non-Education)*

	Total time outside *(mins/day)*	**Discretionary time outside** *(mins/day, net of work + study)*		**Any discretionary time outside** *(Pr; share of days)*	
	All days	**Weekday**	**Sunday**	**Weekday**	**Sunday**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Infancy Rain Shock	13.71***	5.43**	2.33	0.021**	0.007
	(4.58)	(2.67)	(4.84)	(0.010)	(0.017)
× *Boys*	-4.97	-9.83**	-2.64	-0.029**	-0.010
	(5.27)	(3.97)	(8.32)	(0.014)	(0.023)
Early Childhood Rain Shock	-0.05	7.60***	4.93	0.024**	0.045***
	(3.99)	(2.45)	(4.55)	(0.010)	(0.016)
× *Boys*	3.17	-5.52	-21.29***	-0.020	-0.066***
	(4.83)	(3.58)	(7.76)	(0.012)	(0.023)
Observations	30,062	30,062	30,062	30,062	30,062
Girls Mean	320.5	40.2	74.6	0.325	0.380
Boys Mean	436.4	77.8	162.8	0.494	0.649
Boys share in sample	0.53	0.53	0.53	0.53	0.53
[Girls] p: Infancy =Early Ch.	[0.00]	[0.39]	[0.66]	[0.80]	[0.08]
[Boys] p: Infancy =Early Ch.	[0.12]	[0.04]	[0.05]	[0.27]	[0.35]
District-Age-Qtr-Gender FEs	Y	Y	Y	Y	Y
Year-Gender FE	Y	Y	Y	Y	Y

Notes: **Sample.** Rural children aged 15–17, TUS 2019 + 2024 pooled, main 26 states, both genders; the estimation sample unions weekday-normal (Monday–Saturday, normal day), Sunday (any day-type), and weekday-other (Monday–Saturday festival, holiday, or sick) rows. **What it shows.** Five columns from a single regression. (1) Total outside-the-home time in minutes per day, weekday. (2) Outside non-work, non-education time, weekday. (3) The same on Sunday. (4) Indicator for any outside non-work, non-education time, weekday. (5) The same on Sunday. **Outside non-work, non-education time.** Sum of all 30-minute diary slots spent outside the home on activities other than paid work, household production, or formal study; constructed from the slot-level activity-location indicator. **Estimation.** One regression triple-interacts both shocks with gender and a three-way day-type indicator (weekday-normal / Sunday / weekday-other); weekday columns read the weekday-normal base coefficient and Sunday columns add the Sunday interaction. Headline fixed effects (rural ×district ×age ×quarter ×gender, plus survey-year ×gender) and clustering at district ×age ×survey-year. Minutes outcomes winsorised at the state ×gender 99th percentile. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 30: *Permissions and Mobility (NFHS), Rural Girls Aged 15–17*

	Autonomy			Mobility	
	Mean ***(no big prob)***	**All 5** ***(no big prob)***	**All 5** ***(no prob)***	**Mean** ***(any)***	**All 3** ***(alone)***
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Infancy Rain Shock	0.003	-0.003	-0.002	0.003	-0.002
	(0.002)	(0.003)	(0.002)	(0.004)	(0.005)
Early Childhood Rain Shock	0.008***	0.007**	0.000	0.008**	0.002
	(0.002)	(0.003)	(0.002)	(0.003)	(0.005)
Observations	98,898	98,898	98,898	15,847	15,847
Outcome Mean	0.723	0.469	0.162	0.879	0.204
p: Infancy =Early Ch.	[0.04]	[0.01]	[0.32]	[0.25]	[0.46]
District-Age FE	Y	Y	Y	Y	Y
Survey-Year FE	Y	Y	Y	Y	Y

Notes: **Sample.** NFHS-4 (2015–16) and NFHS-5 (2019–21), unmarried rural girls aged 15–17 in the main 26 states. Columns 1–3 use the full women's questionnaire sample. Columns 4–5 use the state-rotating Section 9 module sample. **Outcomes.** Columns 1–3 summarise five barriers to seeking health care: permission, money, distance, transport, and going alone. Columns 4–5 summarise permission to visit a market, health facility, and places outside the village. **Estimation and inference.** Regressions include district ×age and survey-year fixed effects. Standard errors are clustered at district ×age ×survey-year and shown in parentheses. $***p<0.01$, $**p<0.05$, $*p<0.10$.

TABLE 31: *Adult Agency: Exercised Empowerment and Cinema (NFHS)*

	Empowerment				Cinema
	Mean of 5	**Any of 5**	**≥ 3 of 5**	**All 5**	**Monthly+**
	(1)	**(2)**	**(3)**	**(4)**	**(5)**
Infancy Rain Shock	-0.008***	-0.008	-0.013***	0.000	-0.003***
	(0.003)	(0.006)	(0.005)	(0.001)	(0.001)
Early Childhood Rain Shock	0.006**	0.018***	0.009**	-0.000	0.003***
	(0.003)	(0.005)	(0.004)	(0.000)	(0.001)
Observations	15,847	15,847	15,847	15,847	98,898
Outcome Mean	0.256	0.717	0.141	0.002	0.045
p: Infancy =Early Ch.	[0.00]	[0.00]	[0.00]	[0.46]	[0.00]
District-Age FE	Y	Y	Y	Y	Y
Survey-Year FE	Y	Y	Y	Y	Y

Notes: **Sample.** NFHS-4 (2015–16) and NFHS-5 (2019–21), unmarried rural girls aged 15–17 in the main 26 states. Columns 1–4 use the state-rotating Section 9 module sample. Column 5 uses the full women's questionnaire sample. **What it shows.** Columns 1–4 combine five items: ownership of money, a bank account, and a mobile phone; knowledge of a microcredit programme; and having taken a loan. Column 5 is an indicator for attending cinema at least monthly. **Estimation and inference.** Regressions include district ×age and survey-year fixed effects. Standard errors are clustered at district ×age ×survey-year and shown in parentheses. $***p<0.01$, $**p<0.05$, $*p<0.10$.

B Econometric Appendix: Identifying Gender-Specific Effects of a Gender-Invariant Treatment

B.1 Setup

The life-stage rainfall shock is common to girls and boys in a district–survey-year cell, but their baseline time use and their responses to rainfall can differ. With few survey waves, a pooled regression that gives the two groups common fixed effects can mechanically attribute part of those persistent group differences to the gender-specific rainfall coefficient. The main specification avoids this problem by allowing every fixed-effect block to differ by gender. The derivation below establishes that result and then states the corresponding rule for later subgroup analyses.

For transparent algebra, individuals i are observed in districts $d \in \{1, \ldots, D\}$ and years $y \in \{1, \ldots, Y\}$, with gender $g \in \{F, M\}$ and $G_i = 1\{g_i = M\}$. The treatment S_{dy} varies at the district $\times$ year level and is identical for both genders in the same (d, y) cell. Shocks are within-district-standardised, so $\mathbb{E}[S_{dy}] = 0$ and $\mathrm{Var}(S_{dy}) = 1$; the realised within-sample mean $\bar{S}_{d\cdot} \equiv \frac{1}{Y}\sum_y S_{dy}$ has expectation zero but is non-zero in any finite-Y draw. The data-generating process allows gender-heterogeneous effects and gender-specific district intercepts:

$$Y_{idyg} = \beta_F\, S_{dy} + \beta_\Delta\, (G_i \cdot S_{dy}) + \alpha_d^g + \gamma_y + u_{idyg}, \tag{12}$$

with $\beta_M \equiv \beta_F + \beta_\Delta$ the boys' response and $\delta_d \equiv \alpha_d^F - \alpha_d^M$ the district-specific baseline gender gap. We abstract from age and quarter; adding those dimensions to the fixed effects does not change the argument.

B.2 Bias under common (non-gender-interacted) fixed effects

To see the problem, replace the gender-specific intercepts in (12) with a common district fixed effect α_d and a single gender dummy μG_i:

$$Y_{idyg} = \beta_F\, S_{dy} + \beta_\Delta\, (G_i \cdot S_{dy}) + \mu\, G_i + \alpha_d + \gamma_y + \varepsilon_{idyg}. \tag{13}$$

The common district fixed effect removes the pooled district mean, while the gender dummy removes only the average gender difference across districts. The observed district gender gap still contains a component induced by the realised mean rainfall shock:

$$\delta_d^{\mathrm{obs}} = \delta_d - \beta_\Delta\, \bar{S}_{d\cdot}. \tag{14}$$

Thus, when the treatment effect differs by gender, the residual gender difference is mechanically correlated with the realised district mean of the shock. Under the balanced-sample

finite-wave expansion used here, the leading bias is

$$\boxed{\mathbb{E}\left[\hat{\beta}_F^{\text{pooled}} - \beta_F\right] = -\frac{\beta_\Delta}{2} \cdot \frac{\sum_d \bar{S}_{d\cdot}^2}{\sum_{dy} \widetilde{S}_{dy}^2} + \mathcal{R}.} \tag{15}$$

Here $\mathcal{R}$ collects smaller-order terms. The bias on $\hat{\beta}_F$ has the *opposite* sign to β_Δ and is a finite-wave problem: its leading term falls with the number of survey periods, not with the number of districts at a fixed number of periods.

B.3 Gender-interacted fixed effects eliminate the bias

Replacing $\alpha_d + \mu G_i$ in (13) with the gender-interacted intercept α_{dg} absorbs α_d^F and α_d^M separately, including the shock-induced component of the observed gender gap in (14). The resulting specification is algebraically equivalent to estimating separate gender-specific regressions, each relative to its own district baseline, while retaining a pooled coefficient test for the gender difference.

B.4 Direction of bias for the headline outcomes

Equation (15) pins the direction of the common-fixed-effect bias to the sign of β_Δ. For the three headline outcomes in this paper:

Outcome	β_F	β_M	β_Δ	**Bias on $\hat{\beta}_F^{\text{pooled}}$**
Domestic work	> 0	≤ 0	< 0	**upward (overstates β_F)**
Studying	< 0	$\geq \beta_F$	> 0	**downward (overstates $\lvert\beta_F\rvert$)**
Leisure	≈ 0	≈ 0	≈ 0	**negligible**

In both domestic work and studying, the pooled-fixed-effect specification exaggerates the magnitude of the girls' coefficient; where there is no gender-differential response, as for leisure, the bias is negligible. The empirical fixed-effect build-up in Section 5 follows this pattern: the girls' domestic-work coefficient falls from about 10.3 to 5.6 minutes between the common- and gender-interacted-fixed-effect specifications, while leisure changes little.

B.5 Generalisation to other heterogeneity dimensions

The same issue arises whenever a pooled specification estimates a treatment differential across an observed group H_i, while treatment is common to the groups within a fixed-effect cell. Estimating the groups separately automatically gives them group-specific fixed effects. If they are pooled, every fixed-effect block must instead be interacted with H_i. In the binary case, the equal-group factor $\frac{1}{2}$ in the leading-bias expression in equation (15) becomes $p_H(1 - p_H)$, where p_H is the population share with $H_i = 1$; the sign rule is unchanged.

This rule governs the later estimates by day type, landholding, household type, and other observed subgroup categories.

C Model Appendix: Proofs

This appendix gives the proofs for the maintained model in Section 6. Period subscripts are suppressed except where a period-1 shock changes a period-2 state.

C.1 Common Identities

At an optimum, the fixed-proportions technology gives

$$M_r = bD, \qquad D = h_d^p + \gamma h_d^c. \tag{C.1}$$

While $h_m^p > 0$, full income is $F = w + \gamma w + Y$. The CES allocation can be written

$$h_o^c = \frac{s_o F}{\gamma w}, \qquad h_l^p = \frac{s_l F}{w}, \qquad w h_m^p = cF - Y, \tag{C.2}$$

where the shares depend only on prices and $c \in (0,1)$ is the share of full income requiring cash expenditure. At the no-paid-work corner, let w^* denote the internal value of adult time. The corner satisfies the same allocation conditions at w^* and the cash closure $c^* F^* = Y$.

C.2 Proof of Proposition 1A

Differentiating (C.2) with respect to w, holding Y and γ fixed, gives

$$\frac{\partial h_d^c}{\partial w} = \frac{h_o^c}{w}\left[\frac{Y}{F} + (\sigma - 1) s_C \varphi\right] \geq 0,$$

where s_C is the consumption share and $\varphi \in (0,1)$ is the exposure of the consumption price to the wage. The inequality follows from $Y \geq 0$ and $\sigma \geq 1$. This proves Proposition 1A. ■

C.3 Proof of Proposition 1B

When $h_m^p > 0$, a change in Y leaves prices and shares fixed. From (C.2),

$$\frac{\partial h_l^p}{\partial Y} = \frac{s_l}{w} > 0, \qquad \frac{\partial h_d^c}{\partial Y} = -\frac{s_o}{\gamma w} < 0, \qquad \frac{\partial h_d^p}{\partial Y} > 0, \qquad \frac{\partial h_m^p}{\partial Y} = -\frac{1-c}{w} < 0.$$

The interior expression for paid work, $w h_m^p = c[(1+\gamma)w + Y] - Y$, reaches zero at

$$\bar{Y} = (1+\gamma) w \frac{c}{1-c}.$$

At the corner, differentiating the cash closure and the daughter's residual-time condition gives

$$\frac{\partial h_d^c}{\partial Y} > 0 \quad \Longleftrightarrow \quad (1-\sigma)\varphi s_C (1-\varphi) + c^* e_\varphi < 0,$$

where $e_\varphi \equiv \mathrm{d}\log\varphi / \mathrm{d}\log w^*$. The maintained home-production demand condition, $\eta_h > -p_r b / w^*$, implies $e_\varphi < 0$. Together with $\sigma \geq 1$, both terms are non-positive and the second

is strict. This proves Proposition 1B. ■

C.4 Proof of Proposition 2A

Let $\theta_Z \equiv \partial h^c_{d,1}/\partial Z_1 > 0$. The laws of motion give

$$\frac{\partial \log \gamma_2}{\partial Z_1} = A_\gamma \theta_Z, \qquad \frac{\partial(-\log a_{o,2})}{\partial Z_1} = A_\psi \theta_Z.$$

At an interior period-2 allocation, differentiating the CES demand system yields

$$\Pi_\psi = \sigma h^c_{o,2}(1 - s_o) > 0, \qquad \Pi_\gamma = \frac{s_o}{\gamma_2 w_2}\left[F_2 - \gamma_2 w_2 - F_2(1-\sigma)(1-s_o)\right] > 0.$$

The second inequality uses $\sigma \geq 1$. Hence

$$\frac{\partial h^c_{d,2}}{\partial Z_1} = \theta_Z \left(A_\gamma \Pi_\gamma + A_\psi \Pi_\psi\right) > 0$$

whenever either channel is active. This proves Proposition 2A. ■

C.5 Proof of Proposition 2B

Let a dot denote a derivative with respect to Z_1. For a productive-capacity-only economy, write

$$q \equiv \theta_Z A_\gamma > 0, \qquad u \equiv \frac{(\sigma-1)s_o}{\gamma} + \frac{w}{F}, \qquad r \equiv \sigma - \gamma u > 0.$$

The interior first-order conditions give

$$\dot{h}^c_o = -h^c_o r q, \qquad \dot{h}^p_l = h^p_l \gamma u q, \qquad \dot{h}^p_d = \dot{D} + \gamma \dot{h}^c_o - \gamma h^c_d q. \tag{C.3}$$

Set $\kappa \equiv -\dot{h}^c_o > 0$. In a lower-opportunity-cost-only economy, choose any feasible $s'_o \in (0,1)$ and set

$$s'_l = \frac{\dot{h}^p_l}{\gamma\kappa}(1 - s'_o), \qquad \theta'_Z A'_\psi = \frac{\kappa\gamma w}{\sigma s'_o(1-s'_o)F}, \qquad D' = \frac{\dot{h}^p_d - \gamma \dot{h}^c_o}{\sigma s'_o \theta'_Z A'_\psi}. \tag{C.4}$$

Choose the remaining CES weight and home-production parameters to implement these shares and D'. The lower-opportunity-cost economy then has the same $\dot{h}^c_o$, $\dot{h}^p_l$, and $\dot{h}^p_d$ as in (C.3). The daughter's domestic work and the mother's paid work match by the two time constraints. The feasibility conditions are strict and hold on a non-empty open set. Thus the full local time-use response vector can be identical under the two channels. This proves Proposition 2B. ■

C.6 Proof of Proposition 3

Let $H_d = h^p_d + h^c_d$ and $x = h^c_d/H_d$. From (C.1),

$$D = H_d + (\gamma_2 - 1)h^c_d.$$

Let $X_C = p_m C_m + p_r M_r$. Constant returns in $\mathcal{G}$ imply that $\rho = C_m/M_r$ is fixed at given relative prices. Hence

$$\frac{X_C}{H_d} = (p_m\rho + p_r)\,b\big[1 + (\gamma_2 - 1)x\big] \equiv \kappa\big[1 + (\gamma_2 - 1)x\big], \qquad \kappa > 0. \tag{C.5}$$

Write $\beta = -\log a_o$. Under a pure lower-opportunity-cost channel, $A_\gamma = 0$, and the interior allocation gives $x_\beta > 0$: the daughter supplies a larger share of household domestic work. Therefore

$$\frac{\partial(X_C/H_d)}{\partial Y_1} = \kappa\theta_Z A_\psi(\gamma_2 - 1)x_\beta \le 0$$

when $\gamma_2 \le 1$. This proves Proposition 3(a).

For the interaction, differentiating (C.5) with respect to current non-labour income gives

$$\frac{\partial \log(X_C/H_d)}{\partial Y_2} = \frac{1 - \gamma_2}{F_2 H_d} \ge 0.$$

Under $A_\gamma = 0$, differentiating this expression with respect to Y_1 gives

$$\frac{\partial^2 \log(X_C/H_d)}{\partial Y_1 \partial Y_2} = -\theta_Z A_\psi \Phi_\psi \le 0, \qquad \Phi_\psi \equiv \frac{(1 - \gamma_2)H_{d,\beta}}{F_2 H_d^2} > 0,$$

where $H_{d,\beta} > 0$ is the effect of the lower-opportunity-cost shift on total household domestic work. Combining this with the non-positive level response gives

$$\frac{\partial^2(X_C/H_d)}{\partial Y_1 \partial Y_2} = \frac{X_C}{H_d}\left[\frac{\partial^2 \log(X_C/H_d)}{\partial Y_1 \partial Y_2} + \frac{\partial \log(X_C/H_d)}{\partial Y_1}\frac{\partial \log(X_C/H_d)}{\partial Y_2}\right] \le 0.$$

This proves Proposition 3(b). ■