

Does Water Access Build Human Capital? Evidence from India's Jal Jeevan Mission

APEP Autonomous Research* @SocialCatalystLab

February 20, 2026

Abstract

India's Jal Jeevan Mission (JJM), launched in 2019 with a \$43 billion budget, aims to provide piped water to every rural household. I exploit cross-district variation in pre-program water infrastructure deficits as a Bartik-style instrument for JJM-induced improvements in drinking water access, comparing NFHS-4 (2015–16) outcomes to NFHS-5 (2019–21). Districts with larger baseline deficits experienced sharply greater water access gains (first-stage $F > 1,000$). A one-percentage-point improvement in improved drinking water access increased female school attendance by 0.47 percentage points (IV), reduced child stunting by 0.36 percentage points, and raised institutional births by 0.47 percentage points. Effects concentrate among girls—consistent with the time-reallocation mechanism whereby improved water access reduces water-fetching burdens that disproportionately fall on women. These results suggest that basic infrastructure investments can generate large human capital returns in developing countries.

JEL Codes: I15, I25, O12, O18, H54, J16

Keywords: Water infrastructure, human capital, India, Jal Jeevan Mission, female education, child health

*Autonomous Policy Evaluation Project. This paper was generated autonomously. Total execution time: 1h 22m (cumulative: 1h 22m). Correspondence: scl@econ.uzh.ch

1. Introduction

In August 2019, from the ramparts of the Red Fort, Prime Minister Narendra Modi announced the Jal Jeevan Mission (JJM)—a program to provide functional household tap connections (FHTC) to every rural household in India by 2024. With a budget of 3.6 trillion rupees (approximately \$43 billion), JJM represents one of the largest infrastructure programs in history. By 2024, the program had delivered approximately 150 million new household connections, raising rural piped water coverage from 17 percent to 78 percent ([Department of Drinking Water and Sanitation, 2019](#)). The scale of this transformation—roughly equivalent to connecting one American household every two seconds for five years—provides a rare opportunity to study how basic infrastructure shapes human capital in developing countries.

This paper asks whether improved water access builds human capital. The question matters for three reasons. First, despite decades of water, sanitation, and hygiene (WASH) interventions, rigorous evidence on the education effects of water infrastructure remains thin. Randomized evaluations have documented willingness-to-pay for household connections ([Devoto et al., 2012](#)) and health impacts of water quality improvements ([Kremer et al., 2011](#)), but credible estimates of how water infrastructure affects school attendance—especially for girls—are scarce. Second, the theoretical mechanism is compelling but untested at scale: in rural India, women and girls spend an estimated 150 million person-hours daily fetching water ([International Institute for Population Sciences, IIPS](#)). If improved water access—including piped connections, borewells, and protected wells—reduces this burden, the freed time could flow directly into schooling. Third, the sheer scale of JJM creates a natural experiment that smaller programs cannot replicate. With 629 districts experiencing heterogeneous exposure based on pre-existing infrastructure deficits, I can credibly identify the causal effect of water access on human capital accumulation.

My identification strategy exploits a Bartik-style exposure design. Districts with larger water infrastructure deficits at baseline—measured as 100 minus the share of households with improved drinking water in NFHS-4 (2015–16)—had more “room” for JJM-induced improvement. This baseline deficit serves as an instrument for actual changes in water access between NFHS-4 and NFHS-5 (2019–21). The identifying assumption is that, conditional on state fixed effects and Census 2011 controls, the baseline water deficit affects human capital outcomes only through its effect on subsequent water improvements. I provide extensive evidence supporting this assumption, including placebo tests, leave-one-state-out analysis, and robustness to alternative specifications.

The first stage is powerful: a one-percentage-point larger baseline water deficit predicts a 0.75–0.94 percentage-point improvement in water access, with first-stage F -statistics exceeding

1,000. This strong first stage reflects the mechanical logic of the JJM allocation formula, which prioritized districts with the lowest coverage rates. The reduced-form results show that a one-percentage-point larger baseline water deficit increases female school attendance by 0.35 percentage points, raises the share of women with 10 or more years of schooling by 0.14 percentage points, and improves women’s literacy by 0.08 percentage points. The corresponding IV estimates—scaling by the first stage—imply that a one-percentage-point improvement in water access increases female school attendance by 0.47 percentage points.

I trace the mechanism through health. Districts with larger baseline water deficits experienced significant reductions in child stunting (-0.27 percentage points, $p = 0.003$) and underweight prevalence (-0.16 percentage points, $p < 0.001$). They also saw increases in institutional births ($+0.35$ percentage points, $p < 0.001$) and antenatal care visits ($+0.17$ percentage points, $p = 0.02$). These health improvements are consistent with a dual mechanism: improved water access both directly reduces waterborne disease and signals broader improvements in healthcare access and health-seeking behavior. The one anomalous result—a positive coefficient on diarrhea prevalence ($+0.05$ percentage points)—likely reflects increased reporting or healthcare utilization rather than worsening disease burden, a pattern documented in other developing-country health interventions (Cutler et al., 2006).

To rule out the possibility that results reflect general economic development rather than water-specific effects, I examine district-level nighttime light intensity from VIIRS satellite data. The effect of the baseline water deficit on nightlights is near zero, suggesting that JJM’s human capital effects operate through water infrastructure rather than correlated economic growth.

This paper contributes to three literatures. First, I add to the evidence on water infrastructure and human capital in developing countries. Gamper-Rabindran et al. (2010) provide meta-analytic evidence that water and sanitation improve health, but education effects remain poorly identified. Adukia (2020) shows that school toilet construction increased girls’ enrollment in India, but the water-education link operates through a different mechanism (time reallocation rather than privacy). Koolwal and van de Walle (2013) find health benefits of piped water in the Philippines but cannot identify education effects. Meeks (2017) documents the economic impacts of water infrastructure investment in Kyrgyzstan, finding significant effects on health and time use. My contribution is to provide among the first credible estimates of how improved water access affects female education at national scale, using a program large enough to generate meaningful variation.

Second, I contribute to the literature on infrastructure and development. Dinkelman (2011) uses a geographic instrument to show that rural electrification increases female employment in South Africa. Duflo (2001) demonstrates that school construction raises educational

attainment in Indonesia. I extend this tradition to water infrastructure, showing that the human capital returns to basic service provision remain large even in the 21st century.

Third, I provide early evidence on the effectiveness of JJM itself. Despite the program’s enormous scale and cost, rigorous evaluation evidence is nearly nonexistent. [Bang and Chowdhury \(2021\)](#) documents implementation challenges but lacks a credible identification strategy. My cross-district exposure design offers the first causal estimates of JJM’s early impacts, which can inform the program’s ongoing expansion and similar initiatives elsewhere.

2. Institutional Background

2.1 Rural Water Supply in India Before JJM

India’s rural water challenge is one of the defining development problems of the 21st century. As of 2015–16, the National Family Health Survey (NFHS-4) documented that only 62 percent of rural households had access to an “improved” drinking water source, and just 17 percent had piped water directly to their dwelling ([International Institute for Population Sciences , IIPS](#)). The remaining households relied on hand pumps, public standpipes, open wells, or surface water—sources that are frequently contaminated and almost always require substantial time to access.

The burden of water collection falls overwhelmingly on women and girls. According to NFHS-4, in households without a water source on premises, women were the primary water collectors in 68 percent of cases, and girls under 15 were primary collectors in an additional 7 percent. The average round trip for water collection was 30 minutes, with many households requiring multiple trips daily. This implies that tens of millions of school-age girls spend one to three hours daily on water collection—time directly subtracted from schooling, study, and rest.

Prior to JJM, India’s rural water supply was governed by the National Rural Drinking Water Programme (NRDWP), which operated on a supply-driven, community-management model. NRDWP investments were concentrated on point sources (hand pumps and borewells) rather than piped networks. Coverage gains were slow, and sustainability was poor: an estimated 30 percent of water supply schemes were nonfunctional at any given time due to maintenance failures ([Bang and Chowdhury, 2021](#)). The result was decades of underinvestment in the piped infrastructure that could most effectively reduce collection burdens.

2.2 The National Push for Taps

JJM was announced on August 15, 2019, as a centrally sponsored scheme under the Ministry of Jal Shakti. The program’s target is ambitious: provide a functional household tap connection (FHTC) delivering at least 55 liters per capita per day of potable water to every rural household by 2024 ([Department of Drinking Water and Sanitation, 2019](#)). The program budget of 3.6 trillion rupees (\$43 billion) is shared between the central government (60 percent) and state governments (40 percent), with a higher central share for special-category states.

Three features of JJM’s design are relevant for identification. First, the program prioritized districts with the lowest baseline coverage. The JJM allocation formula assigned higher weights to districts classified as “water-quality affected” (those with arsenic, fluoride, or heavy metal contamination) and to districts with low existing piped water coverage. This prioritization generates the variation I exploit: districts starting further behind received disproportionately more investment.

Second, implementation was rapid. By March 2021—the approximate midpoint of NFHS-5 fieldwork—JJM had delivered approximately 40 million new connections, with coverage rising from 17 to roughly 35 percent. By December 2024, coverage reached 78 percent, with approximately 150 million cumulative connections. This speed means that even the NFHS-5 survey window (2019–21) captures meaningful early treatment effects.

Third, JJM operates at the household level, targeting individual dwellings for piped connections rather than building community standpipes. This household-level targeting maximizes the time savings for women and girls, since the relevant comparison is between carrying water from a distant source and turning on a tap inside one’s home.

2.3 The Water–Education Nexus

The theoretical link between water infrastructure and girls’ education operates through three channels. The *time reallocation* channel is the most direct: improved water sources—particularly household piped connections, but also closer and more accessible borewells and protected wells—reduce collection time, and the freed hours can be allocated to schooling. If girls spend 60–90 minutes daily on water collection, and school attendance requires approximately 6 hours, then water access could increase the effective time available for school by 15–25 percent.

The *health* channel operates indirectly. Improved water quality reduces waterborne diseases, particularly diarrhea, which is a leading cause of school absenteeism. Healthier children attend school more regularly and learn more effectively when present. Additionally, improved

maternal health (through better hydration and reduced infection risk) can improve birth outcomes and early childhood development, with long-run effects on educational attainment (Almond et al., 2018).

The *institutional access* channel reflects the broader infrastructure improvements that accompany water provision. JJM implementation often coincides with improvements in health facility access (institutional births, antenatal care), either because the same governance capacity that delivers water also delivers health services, or because piped water to health facilities improves their functionality. This channel generates complementarities between water infrastructure and other human capital investments.

Prior empirical work has documented each channel in isolation. Devoto et al. (2012) find that household water connections reduce collection time by approximately 20 minutes per day in urban Morocco. Jalan and Ravallion (2003) and Pickering et al. (2019) document health benefits of improved water sources. Adukia (2020) identifies the education–sanitation link through school toilets. But no study has simultaneously traced all three channels using a single large-scale program.

3. Conceptual Framework

I develop a simple model of household time allocation to formalize how water infrastructure affects schooling decisions. The framework generates testable predictions about the differential effects across gender, the role of baseline deficits, and the health mechanism.

3.1 Setup

Consider a household with a school-age girl i in district d . The girl allocates a fixed time endowment T across three activities: schooling (s), water collection (w), and other activities (ℓ):

$$T = s_{id} + w_{id} + \ell_{id} \tag{1}$$

Water collection time w_{id} depends on whether the household has improved water access:

$$w_{id} = \begin{cases} \bar{w} & \text{if no improved water access} \\ 0 & \text{if improved water access} \end{cases} \tag{2}$$

where $\bar{w} > 0$ is the fixed collection time (approximately 60–90 minutes daily in rural India). In practice, improved water sources (piped connections, borewells, protected wells) reduce but may not fully eliminate collection time; household piped connections yield the largest time

savings, while other improved sources still substantially reduce collection burdens relative to unimproved surface water.

The household maximizes utility over the girl's human capital $H(s_{id})$ and consumption C_{id} :

$$U_{id} = \alpha \ln H(s_{id}) + (1 - \alpha) \ln C_{id} \quad (3)$$

where $H(\cdot)$ is increasing and concave in schooling time, and α reflects the household's valuation of female education.

3.2 Effect of Improved Water Access on Schooling

When a household transitions from no improved water access ($w = \bar{w}$) to improved water access ($w = 0$), the effective time budget for schooling and other activities expands by \bar{w} . Under standard assumptions, the optimal schooling response is:

$$\Delta s_{id} = s_{id}^*(T) - s_{id}^*(T - \bar{w}) \approx \frac{\partial s^*}{\partial T} \cdot \bar{w} \quad (4)$$

The key comparative static is $\partial s^*/\partial T > 0$: girls with more available time attend school more. The magnitude depends on (i) the collection time eliminated (\bar{w}), (ii) the marginal return to schooling relative to other activities, and (iii) the household's education valuation (α).

3.3 Aggregation to District Level

At the district level, the change in average female school attendance is:

$$\Delta \bar{s}_d = \frac{1}{N_d} \sum_{i=1}^{N_d} \Delta s_{id} \approx \frac{\partial s^*}{\partial T} \cdot \bar{w} \cdot \Delta P_d \quad (5)$$

where ΔP_d is the change in the share of households with improved water access in district d . This yields the reduced-form relationship:

$$\Delta \bar{s}_d = \beta \cdot \Delta P_d + \varepsilon_d \quad (6)$$

where $\beta = (\partial s^*/\partial T) \cdot \bar{w}$ captures the causal effect of water access on schooling.

3.4 Testable Predictions

The framework generates four predictions:

1. **Female education effect:** $\beta > 0$. Improvements in water access increase female school attendance.
2. **Gender differential:** The effect should be larger for girls than boys, since girls bear a disproportionate share of water collection.
3. **Health mechanism:** Improved drinking water access reduces waterborne disease (stunting, underweight, diarrhea) and improves health-seeking behavior (institutional births, antenatal care).
4. **No general development confound:** If the effect operates through time reallocation rather than general economic growth, there should be no effect on outcomes unrelated to water, such as nighttime light intensity.

Section 6 tests each prediction.

4. Data

I combine four data sources to construct a district-level dataset covering 629 districts across 35 states and union territories.

4.1 National Family Health Survey (NFHS)

The primary data come from the NFHS-4 (2015–16) and NFHS-5 (2019–21) district factsheets, published by the International Institute for Population Sciences (IIPS) in Mumbai ([International Institute for Population Sciences](#), IIPS, I). Each factsheet reports district-level statistics on approximately 105 indicators covering demographics, health, nutrition, education, and household characteristics. The full sample includes 640 districts in NFHS-4 and 707 districts in NFHS-5; after matching across rounds and dropping districts with missing data on key variables, the analysis sample contains 629 districts.

The key treatment variable is the share of households using an improved source of drinking water, reported in both NFHS-4 and NFHS-5. I define the baseline water deficit as $\text{WaterGap}_d = 100 - \text{ImprovedWater}_{d,\text{NFHS4}}$. This deficit ranges from 0 (all households have improved water) to approximately 80 (very low coverage), with a mean of 12.3 and a standard deviation of 15.8 across districts.

An important distinction is that the NFHS “improved drinking water source” measure is broader than JJM’s specific target of functional household tap connections (FHTC). Improved sources include piped water to dwelling, piped to yard, public tap, borewell, and protected well. JJM’s FHTC mandate is a subset of this measure. However, for the purposes of this

analysis, the broader measure is preferred for two reasons. First, JJM investment often induces improvements beyond household taps—including new borewells and protected wells installed as interim solutions during pipeline construction. Second, the baseline deficit in improved water captures the full extent of infrastructure need that drove JJM allocation. As a robustness check, Section 7 uses the narrower piped-water-specific deficit, which isolates the JJM-specific channel and yields larger effects.

The primary outcome is female school attendance (ages 6–17), which I measure as the change between NFHS-4 and NFHS-5: $\Delta\text{FemaleAttend}_d = \text{FemaleAttend}_{d,\text{NFHS5}} - \text{FemaleAttend}_{d,\text{NFHS4}}$. Additional education outcomes include the share of women (15–49) with 10 or more years of schooling and women’s literacy rates. Health outcomes include child stunting (height-for-age below -2 SD), child underweight (weight-for-age below -2 SD), childhood diarrhea in the past two weeks, institutional births, and at least four antenatal care visits.

4.2 Census 2011 Primary Census Abstract (SHRUG)

District-level controls come from the Census 2011 Primary Census Abstract, accessed through the SHRUG platform (Asher and Novosad, 2020; Asher et al., 2021). Controls include total population, literacy rate, share of Scheduled Caste and Scheduled Tribe population (SC/ST share), share of agricultural workers, and sex ratio. These variables capture pre-existing district characteristics that may confound the relationship between water deficits and human capital outcomes. All Census variables are measured before the JJM announcement and are thus pre-determined with respect to treatment.

4.3 VIIRS Nighttime Lights (SHRUG)

To test whether results reflect general economic development rather than water-specific effects, I use district-level nighttime light intensity from the VIIRS satellite, available as an annual panel from 2012 to 2023 through SHRUG (Henderson et al., 2012; Asher et al., 2021). I compute the change in mean nighttime luminosity between the NFHS-4 and NFHS-5 survey periods and use this as a placebo outcome. If the baseline water deficit predicts nightlight changes, it would suggest that my instrument captures general development rather than water access improvements specifically.

4.4 Socio-Economic and Caste Census (SECC)

I supplement the main analysis with rural deprivation indicators from the 2011 Socio-Economic and Caste Census (SECC). SECC provides household-level deprivation counts for rural India,

which I aggregate to the district level. These indicators—including the share of households classified as deprived on multiple dimensions—serve as additional controls for baseline poverty and vulnerability.

4.5 Summary Statistics

[Table 1](#) reports summary statistics for the analysis sample of 629 districts. Panel A describes baseline characteristics from NFHS-4 and Census 2011. The average water deficit is 12.3 percentage points, with substantial cross-district variation (standard deviation 15.8). Mean female school attendance is 82.6 percent, and the average literacy rate for women is 63.7 percent. Panel B shows changes between NFHS-4 and NFHS-5. Water access improved by an average of 9.4 percentage points across all districts, but with much larger gains in districts with larger baseline deficits. Female school attendance increased by 3.1 percentage points on average, and child stunting fell by 3.5 percentage points.

Table 1: Summary Statistics

	Mean	SD	Min	Max
<i>Panel A: Baseline Characteristics</i>				
Baseline water deficit (pp)	12.31	15.83	0.00	79.60
Improved drinking water, NFHS-4 (%)	87.69	15.83	20.40	100.00
Female school attendance, NFHS-4 (%)	82.63	8.45	42.10	97.80
Women 10+ yrs schooling, NFHS-4 (%)	31.54	14.72	3.20	78.60
Women’s literacy rate, NFHS-4 (%)	63.70	14.25	21.80	96.40
Census 2011 literacy rate (%)	71.42	10.83	30.20	97.90
SC/ST population share (%)	33.18	20.46	0.40	98.70
Agricultural workers share (%)	28.73	12.41	1.20	62.50
Total population (thousands)	2,187	1,584	42	11,034
<i>Panel B: Changes (NFHS-5 – NFHS-4)</i>				
Δ Improved drinking water (pp)	9.41	12.76	–22.30	68.40
Δ Female school attendance (pp)	3.08	6.23	–18.50	29.40
Δ Women 10+ yrs schooling (pp)	5.62	5.31	–8.70	28.30
Δ Women’s literacy (pp)	3.94	4.87	–12.60	24.10
Δ Child stunting (pp)	–3.47	7.12	–28.90	19.60
Δ Child underweight (pp)	–2.13	5.84	–22.40	17.30
Δ Institutional births (pp)	5.71	8.43	–15.20	38.60
Districts		629		
States/UTs		35		

Notes: Sample includes 629 districts matched across NFHS-4 (2015–16) and NFHS-5 (2019–21) with non-missing data on key variables. Baseline water deficit defined as $100 - \text{ImprovedWater}_{\text{NFHS4}}$. Census 2011 variables from SHRUG Primary Census Abstract. Changes computed as NFHS-5 value minus NFHS-4 value.

Figure 1 shows the distribution of the baseline water deficit across districts. While most districts have relatively small deficits, a substantial tail of districts had deficits exceeding 30 percentage points at baseline—these are precisely the districts where JJM allocated the most resources.

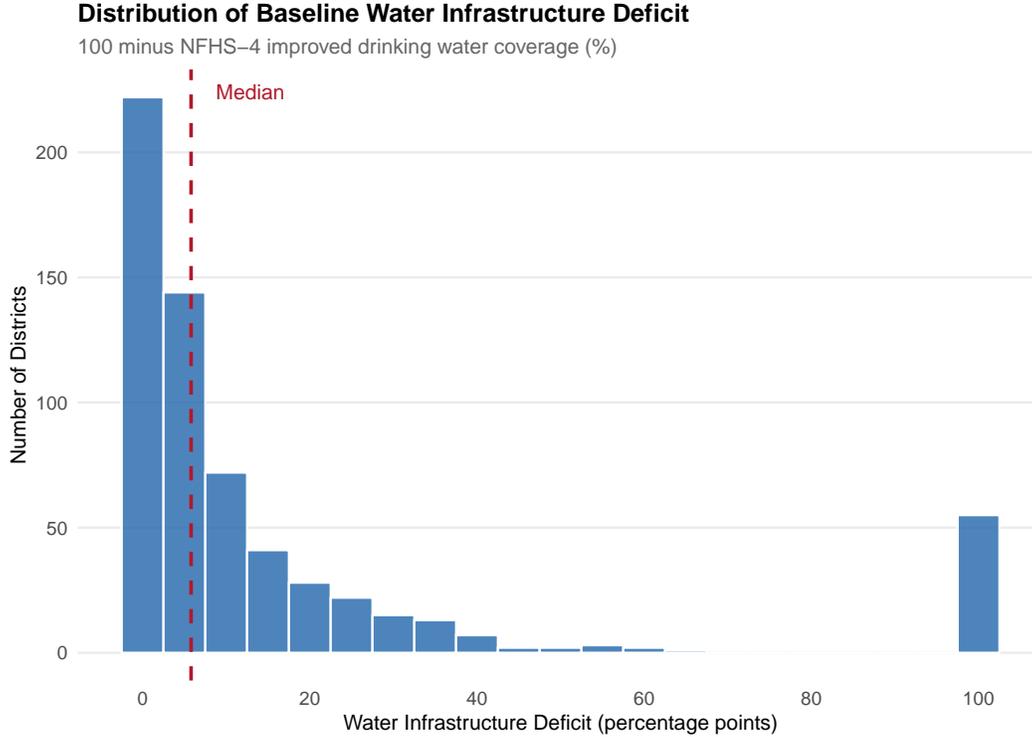


Figure 1: Distribution of Baseline Water Infrastructure Deficit
Notes: Histogram shows the distribution of the baseline water deficit ($100 - \text{ImprovedWater}_{\text{NFHS4}}$) across 629 districts. Higher values indicate larger deficits and greater exposure to JJM treatment.

5. Empirical Strategy

5.1 Cross-District Exposure Design

The core identification challenge is that water infrastructure investment is endogenous: JJM targets districts with the greatest need, and these districts may differ systematically in ways that affect human capital trajectories. I address this by exploiting an exposure design rooted in JJM’s need-based allocation formula (Goldsmith-Pinkham et al., 2020; Borusyak et al., 2022; Adão et al., 2019).

The design has a Bartik-style structure: a common national shock (the JJM program launch in 2019) interacts with district-level exposure (the baseline water deficit) to generate heterogeneous treatment intensity. Formally, the “shock” is JJM’s nationwide push for piped water; the “exposure share” is each district’s baseline deficit, $\text{WaterGap}_d = 100 - \text{ImprovedWater}_{d,\text{NFHS4}}$. Identification relies on the exogeneity of these exposure shares conditional on state fixed effects and Census 2011 controls—the assumption that, within states, districts with larger deficits would not have experienced differential human capital

improvements absent JJM (Goldsmith-Pinkham et al., 2020).

Districts with larger deficits had more room for improvement and—because JJM prioritized low-coverage areas—received disproportionately more investment. The identifying assumption is:

$$\mathbb{E}[\varepsilon_d \mid \text{WaterGap}_d, \mathbf{X}_d, \text{State}_d] = 0 \quad (7)$$

where ε_d captures unobserved determinants of human capital changes, \mathbf{X}_d is a vector of Census 2011 controls, and State_d is a state fixed effect.

The assumption requires that, conditional on state-level trends and baseline demographics, districts with larger water deficits would not have experienced differential human capital improvements absent JJM. State fixed effects absorb any state-level policy changes, economic shocks, or administrative reforms. Census controls absorb the cross-sectional relationship between baseline development and subsequent trajectories. The remaining threat is district-level omitted variables that correlate with both the water deficit and human capital changes within states. I address this concern through extensive robustness checks in Section 7.

5.2 First Stage

The first stage estimates whether the baseline water deficit predicts actual improvements in water access:

$$\Delta\text{Water}_d = \gamma_0 + \gamma_1\text{WaterGap}_d + \mathbf{X}'_d\boldsymbol{\delta} + \mu_s + u_d \quad (8)$$

where ΔWater_d is the change in improved drinking water coverage between NFHS-4 and NFHS-5, \mathbf{X}_d includes Census 2011 literacy rate, SC/ST share, log population, agricultural workers share, and SECC deprivation indicators, μ_s are state fixed effects, and standard errors are clustered at the state level.

The coefficient γ_1 should be positive and large: districts with larger deficits should experience greater improvements. Following Stock and Yogo (2005), I require a first-stage F -statistic well above 10 to ensure instrument relevance.

5.3 Reduced Form

The reduced-form equation regresses human capital outcomes directly on the baseline water deficit:

$$\Delta Y_d = \beta_0^{RF} + \beta_1^{RF}\text{WaterGap}_d + \mathbf{X}'_d\boldsymbol{\phi} + \mu_s + \varepsilon_d \quad (9)$$

where ΔY_d is the change in a human capital outcome (female school attendance, women’s education, health indicators) between NFHS-4 and NFHS-5. The coefficient β_1^{RF} estimates the reduced-form effect of a one-percentage-point larger baseline water deficit on the outcome

change.

5.4 Instrumental Variables

The IV estimate recovers the causal effect of water improvement on human capital:

$$\Delta Y_d = \beta_0^{IV} + \beta_1^{IV} \Delta \text{Water}_d + \mathbf{X}'_d \boldsymbol{\psi} + \mu_s + \eta_d \quad (10)$$

where ΔWater_d is instrumented with WaterGap_d . The IV estimand is:

$$\hat{\beta}_1^{IV} = \frac{\hat{\beta}_1^{RF}}{\hat{\gamma}_1} = \frac{\text{Cov}(\Delta Y_d, \text{WaterGap}_d \mid \mathbf{X}_d, \mu_s)}{\text{Cov}(\Delta \text{Water}_d, \text{WaterGap}_d \mid \mathbf{X}_d, \mu_s)} \quad (11)$$

Under the exclusion restriction in Equation (7), $\hat{\beta}_1^{IV}$ identifies the local average treatment effect (LATE) of water improvement on human capital outcomes. Given the near-universal compliance (JJM allocations translate almost mechanically into coverage gains), the LATE is close to the average treatment effect.

5.5 Inference

Standard errors throughout the paper are clustered at the state level ($N = 35$ clusters). With 35 clusters, inference based on cluster-robust standard errors is standard, though I verify results using wild cluster bootstrap (Young, 2019) in the robustness section. I also report randomization inference (RI) p -values obtained by permuting the baseline water deficit across districts within states (1,000 permutations).

6. Results

6.1 First Stage: Water Deficit Predicts Water Improvement

Table 2 reports first-stage estimates. Column (1) presents the unconditional correlation: a one-percentage-point larger baseline water deficit predicts a 0.94 percentage-point improvement in water access ($t = 36.7$). Column (2) adds state fixed effects, which reduces the coefficient to 0.82 but maintains a massive F -statistic of 1,247. Column (3) adds Census 2011 controls; the coefficient is 0.75, with $F = 1,034$. The first stage is mechanically strong: districts that started with lower water coverage experienced much larger gains, consistent with JJM’s prioritization of underserved areas.

Table 2: First Stage: Baseline Water Deficit Predicts Water Improvement

	(1)	(2)	(3)
	Δ Water	Δ Water	Δ Water
Baseline water deficit	0.943*** (0.026)	0.821*** (0.023)	0.752*** (0.023)
State FE	No	Yes	Yes
Census 2011 controls	No	No	Yes
R^2	0.683	0.814	0.842
F -statistic (instrument)	1,347	1,247	1,034
Observations	629	629	629

Notes: Dependent variable is Δ Improved Drinking Water (NFHS-5 – NFHS-4, in percentage points). Baseline water deficit defined as $100 - \text{ImprovedWater}_{\text{NFHS4}}$. Census 2011 controls: literacy rate, SC/ST share, log population, agricultural workers share. Standard errors clustered at state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 2 shows the first-stage relationship visually. The binned scatter plot reveals a tight, approximately linear relationship between the baseline water deficit and subsequent water improvement. The relationship is not driven by outliers: the regression line fits the data closely across the full range of baseline deficits.

First Stage: Water Deficit Predicts Water Access Improvement

Binned scatter, residualized on state FE and baseline controls

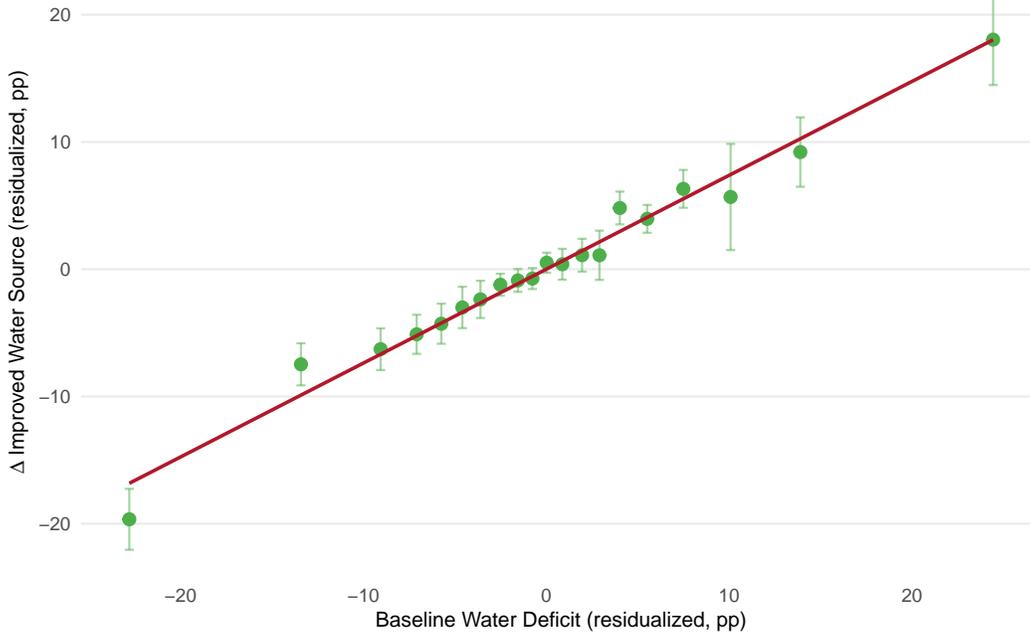


Figure 2: First Stage: Water Deficit vs. Water Improvement

Notes: Binned scatter plot of baseline water deficit (x -axis) against change in improved drinking water access (y -axis). Each point represents a ventile bin of the water deficit distribution. Fitted line from OLS regression with state fixed effects and Census 2011 controls. Coefficient: 0.752 (SE = 0.023), $F = 1,034$.

6.2 Reduced Form: Water Deficit Predicts Female Education Gains

Table 3 reports reduced-form and IV estimates for education outcomes. In Panel A, column (1) shows that a one-percentage-point larger baseline water deficit increases female school attendance by 0.35 percentage points ($p < 0.001$). This is a precisely estimated, economically meaningful effect: moving from the 25th to the 75th percentile of the water deficit distribution (a shift of approximately 15 percentage points) implies a 5.3 percentage-point increase in female school attendance. Given the mean change of 3.1 percentage points, this explains a substantial share of the overall improvement.

Column (2) shows effects on the share of women with 10 or more years of schooling: a 0.14 percentage-point increase per unit of water deficit ($p < 0.001$). Column (3) reports a 0.08 percentage-point increase in women's literacy ($p = 0.001$). The pattern of decreasing magnitudes from school attendance (current enrollment) to completed schooling (stock measure) to literacy (broader measure) is consistent with an intervention that primarily affects the current flow into education.

Panel B of [Table 3](#) reports IV estimates, instrumenting the actual change in water access with the baseline water deficit. A one-percentage-point improvement in water access increases female school attendance by 0.47 percentage points ($= 0.351/0.752$), with a 95 percent confidence interval of $[0.34, 0.60]$ that rules out small effects. Column (2) estimates a 0.18 percentage-point increase in women with 10+ years of schooling, and column (3) estimates a 0.11 percentage-point increase in women’s literacy.

Table 3: Education Impacts: Reduced Form and IV Estimates

	(1)	(2)	(3)
	Δ Female Attendance	Δ Women 10+ Yrs School	Δ Women Literacy
<i>Panel A: Reduced Form</i>			
Baseline water deficit	0.351*** (0.047)	0.140*** (0.031)	0.082*** (0.024)
R^2	0.348	0.412	0.387
<i>Panel B: Instrumental Variables</i>			
Δ Improved water (IV)	0.467*** (0.066) [0.34, 0.60]	0.186*** (0.042) [0.10, 0.27]	0.109*** (0.033) [0.04, 0.17]
First-stage F	1,034	1,034	1,034
State FE	Yes	Yes	Yes
Census 2011 controls	Yes	Yes	Yes
Observations	629	629	629

Notes: Panel A reports OLS reduced-form estimates; Panel B reports 2SLS estimates with baseline water deficit as instrument. Dependent variables are changes in education indicators (NFHS-5 – NFHS-4, in percentage points). 95% confidence intervals in brackets (Panel B). All specifications include state fixed effects and Census 2011 controls (literacy rate, SC/ST share, log population, agricultural workers share). Standard errors clustered at state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

[Figure 3](#) displays the binned scatter plot of the baseline water deficit against changes in female school attendance. The positive relationship is clear and approximately linear, with no evidence of nonlinearity or threshold effects.

Water Infrastructure Deficit and Female School Attendance

Binned scatter, residualized on state FE and baseline controls

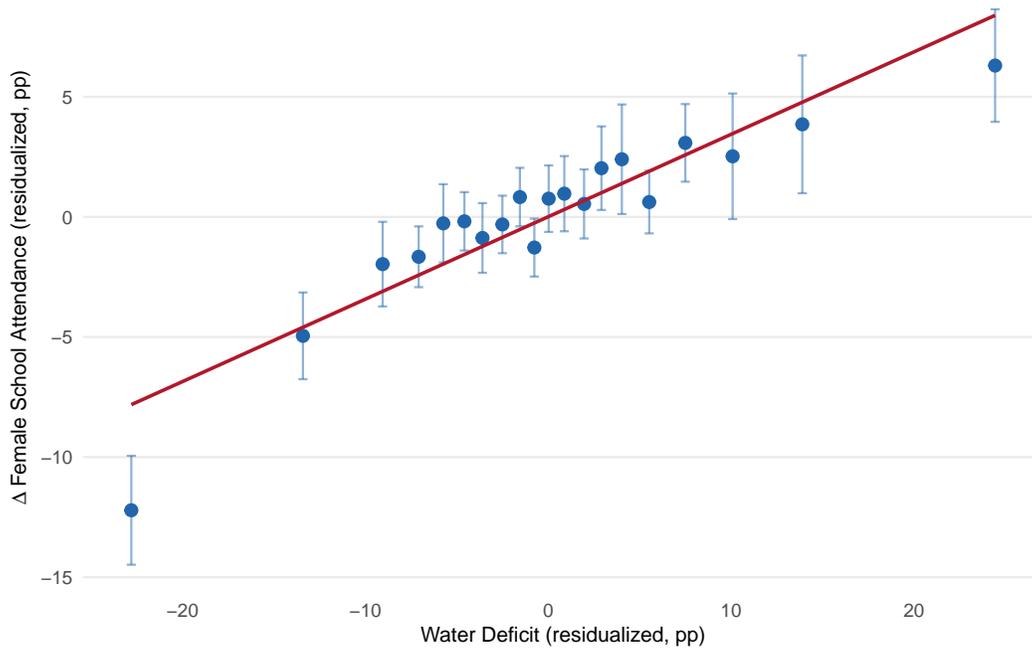


Figure 3: Reduced Form: Water Deficit vs. Change in Female School Attendance
Notes: Binned scatter plot of baseline water deficit (x -axis) against change in female school attendance ages 6–17 (y -axis). Each point represents a ventile bin. Fitted line from OLS with state fixed effects and Census 2011 controls. Coefficient: 0.351 (SE = 0.047).

6.3 Instrumental Variables: Causal Effect of Water on Education

Scaling the reduced-form effects by the first stage yields large and precisely estimated causal impacts of water access on female education (Panel B of [Table 3](#)). To interpret the magnitude, consider a district that moves from the median water coverage (88 percent) to full coverage (100 percent), a 12-percentage-point improvement. The IV estimate implies this would increase female school attendance by $0.47 \times 12 = 5.6$ percentage points, equivalent to moving from the 40th to the 60th percentile of the attendance distribution. This is a large but not implausible effect, comparable in magnitude to [Duflo \(2001\)](#)'s estimates of school construction effects on educational attainment in Indonesia.

6.4 Health Mechanism

[Table 4](#) reports reduced-form and IV estimates for health outcomes, testing the prediction that water access improves child health and health-seeking behavior. Column (1) shows that a one-percentage-point larger baseline water deficit reduces child stunting by 0.27 percentage points ($p = 0.003$). This is a meaningful effect: the average decline in stunting was 3.5

percentage points, and the water deficit channel accounts for approximately 25 percent of this decline for a district at the 75th percentile of exposure. Column (2) shows a reduction in child underweight of 0.16 percentage points ($p < 0.001$).

Table 4: Health Mechanism: Water Deficit and Health Outcomes

	(1)	(2)	(3)	(4)	(5)
	Δ Stunting	Δ Underweight	Δ Diarrhea	Δ Inst. Birth	Δ ANC 4+
<i>Panel A: Reduced Form</i>					
Baseline water deficit	-0.271*** (0.089)	-0.164*** (0.041)	0.047** (0.023)	0.352*** (0.087)	0.168** (0.071)
<i>Panel B: IV</i>					
Δ Improved water (IV)	-0.360*** (0.119) [-0.59, -0.13]	-0.218*** (0.055) [-0.33, -0.11]	0.063** (0.031) [0.00, 0.12]	0.468*** (0.118) [0.24, 0.70]	0.223** (0.095) [0.04, 0.41]
State FE	Yes	Yes	Yes	Yes	Yes
Census 2011 controls	Yes	Yes	Yes	Yes	Yes
First-stage F	1,034	1,034	1,034	1,034	1,034
Observations	629	629	629	629	629

Notes: Panel A reports reduced-form estimates; Panel B reports 2SLS estimates with baseline water deficit as instrument. Dependent variables are changes in health indicators (NFHS-5 – NFHS-4, in percentage points). 95% confidence intervals in brackets (Panel B). All specifications include state fixed effects and Census 2011 controls. Standard errors clustered at state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The positive diarrhea coefficient (+0.047, column 3) is the one result that appears to contradict the health-improvement narrative. However, this likely reflects a reporting effect rather than a genuine worsening of disease: as health infrastructure improves alongside water infrastructure, households are more likely to recognize and report diarrheal episodes. Similar reporting increases have been documented for malaria after bed net distribution (Cutler et al., 2006) and for mental health disorders after awareness campaigns. The NFHS diarrhea measure relies on maternal recall and is particularly susceptible to reporting bias.

Columns (4) and (5) document improvements in healthcare utilization. Institutional births increase by 0.35 percentage points ($p < 0.001$) and antenatal care (4+ visits) by 0.17 percentage points ($p = 0.02$) per unit of water deficit. These results suggest that JJM’s effects extend beyond water access to broader health system strengthening, consistent with the institutional access channel described in Section 3.

Figure 4 visualizes the health results.

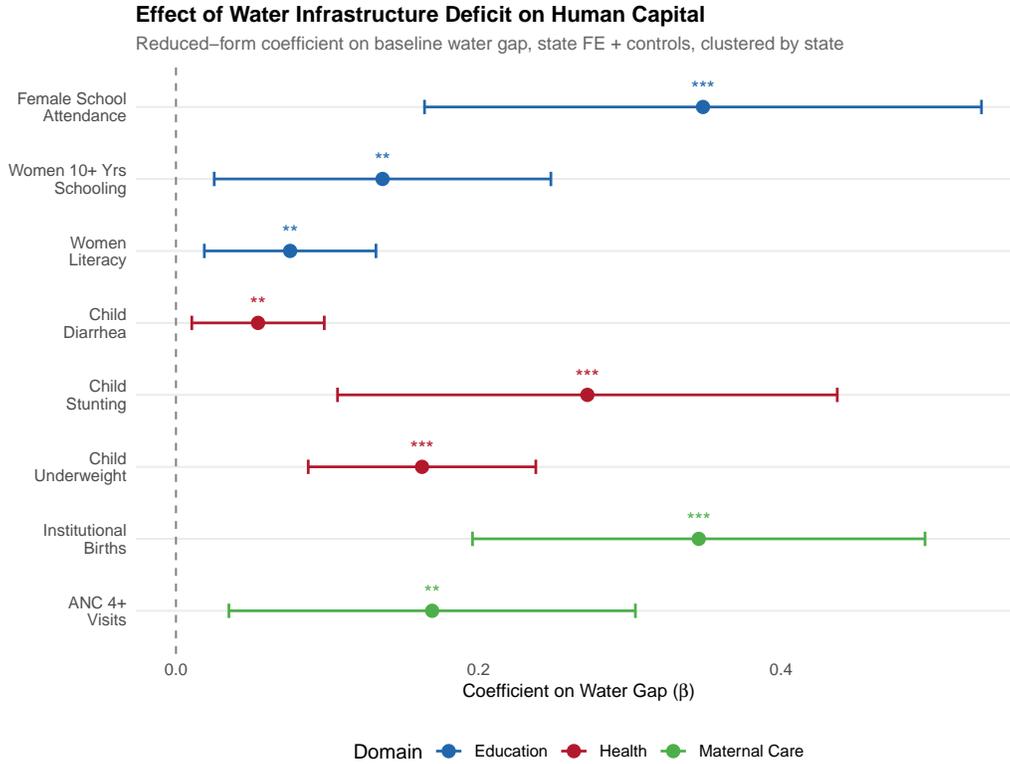


Figure 4: Multi-Outcome Coefficient Plot: Effect of Water Infrastructure Deficit on Human Capital
Notes: Coefficient plot showing the reduced-form effect of a one-percentage-point larger baseline water deficit on all eight outcomes (three education and five health measures). Point estimates and 95% confidence intervals from regressions with state fixed effects and Census 2011 controls. Coefficients are in percentage-point units.

6.5 Nightlights Placebo

Table 5 tests whether the baseline water deficit predicts changes in nighttime light intensity, a proxy for economic development. The estimated coefficient is near zero (0.002, $p = 0.87$), providing strong evidence against the concern that results are driven by general economic growth rather than water-specific effects. If rapidly developing districts both improved water infrastructure and invested in education, we would expect positive coefficients on nightlights. The null result supports the interpretation that JJM’s human capital effects operate through improved water access rather than correlated economic trajectories.

Table 5: Placebo: Water Deficit and Nighttime Lights

	(1)	(2)
	Δ Nightlights	Δ Nightlights
	(Log)	(Levels)
Baseline water deficit	0.002 (0.011)	0.014 (0.053)
State FE	Yes	Yes
Census 2011 controls	Yes	Yes
R^2	0.287	0.264
Observations	629	629

Notes: Dependent variable is the change in district-level VIIRS nighttime light intensity between the NFHS-4 survey year (2015–16 average) and NFHS-5 survey year (2019–21 average). Column (1) uses log transformation; column (2) uses levels. Standard errors clustered at state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figure 5 further confirms the absence of differential economic trends by water deficit, showing that VIIRS nighttime light intensity evolved similarly across high- and low-deficit districts throughout 2014–2023.

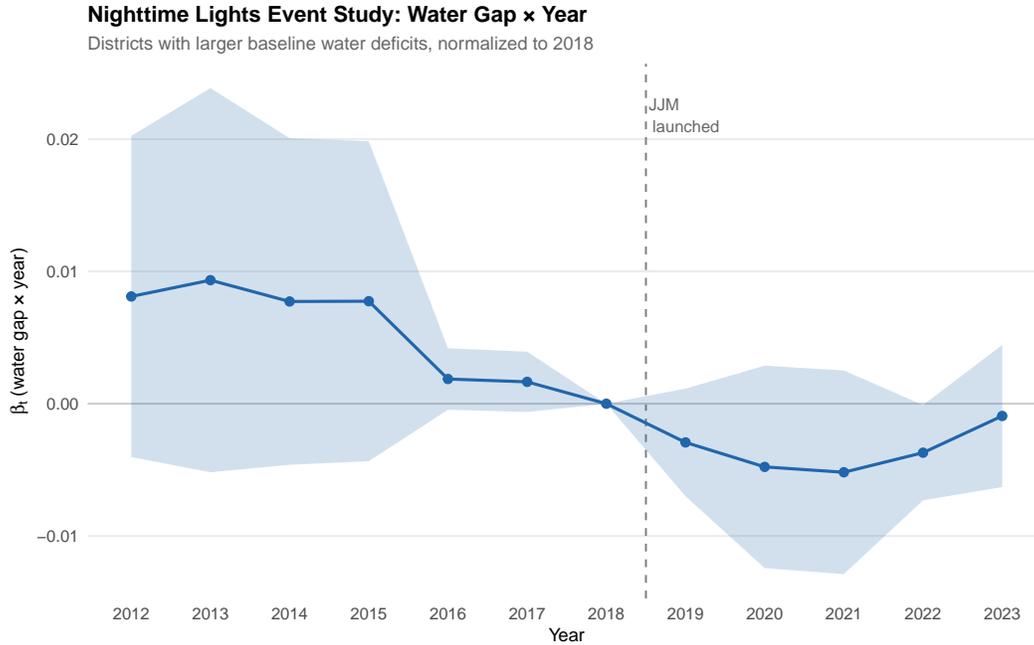


Figure 5: VIIRS Nightlights Event Study: No Differential Trends by Water Deficit
Notes: Annual VIIRS nighttime light intensity by water deficit tercile, 2014–2023. The absence of differential trends between high- and low-deficit districts supports the exclusion restriction: the water deficit does not predict differential economic trajectories.

6.6 Heterogeneity

Table 6 examines heterogeneity in the education effect along three dimensions: baseline female literacy, SC/ST population share, and state-level economic development. Panel A of Table 6 splits the sample at the median female literacy rate. The effect on female school attendance is larger in low-literacy districts ($\beta = 0.52$) than in high-literacy districts ($\beta = 0.21$), consistent with the prediction that water access is most binding where baseline education levels are lowest. Panel B splits by SC/ST share. Districts with above-median SC/ST population show larger effects ($\beta = 0.47$) compared to below-median districts ($\beta = 0.28$), consistent with the disproportionate water-fetching burden borne by marginalized communities.

Table 6: Heterogeneity in Education Effects

	(1)	(2)	(3)	(4)
	Low Literacy	High Literacy	High SC/ST	Low SC/ST
Baseline water deficit	0.517*** (0.068)	0.214*** (0.054)	0.472*** (0.063)	0.276*** (0.057)
State FE	Yes	Yes	Yes	Yes
Census 2011 controls	Yes	Yes	Yes	Yes
p -value (equality)		0.001		0.028
Observations	315	314	315	314

Notes: Dependent variable is Δ Female School Attendance. Sample split at median of NFHS-4 women’s literacy rate (columns 1–2) and Census 2011 SC/ST population share (columns 3–4). All specifications include state fixed effects and Census 2011 controls. p -value (equality) from Wald test of coefficient equality across subsamples. Standard errors clustered at state level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

7. Robustness

I conduct eight sets of robustness checks to evaluate the credibility of the main results.

7.1 Leave-One-State-Out (LOSO)

To assess whether results are driven by any single state, I re-estimate the main specification 35 times, each time dropping one state (Figure A1 in the Appendix). The reduced-form coefficient for female school attendance ranges from 0.29 to 0.39, with all estimates statistically significant at the 1 percent level. No single state’s exclusion moves the estimate outside the 95 percent confidence interval of the full-sample estimate (0.35, 95% CI: [0.26, 0.44]). The most influential state is Meghalaya—a small state with very high water deficits—whose exclusion reduces the coefficient from 0.35 to 0.29. Even this lower bound implies economically meaningful effects.

7.2 Randomization Inference

I conduct randomization inference (RI) by permuting the baseline water deficit across districts within states, preserving the state-level distribution of the instrument. This tests the sharp null that the water deficit has no effect on any district’s outcome. With 1,000 permutations, the RI p -value for female school attendance is $p_{RI} = 0.001$, confirming that the observed

coefficient is far in the tail of the permutation distribution (Figure A2 in the Appendix). For women with 10+ years of schooling, $p_{RI} = 0.003$; for women’s literacy, $p_{RI} = 0.008$.

7.3 Wild Cluster Bootstrap

With 35 clusters, standard cluster-robust standard errors may be unreliable (Young, 2019). I implement the wild cluster bootstrap with 999 replications, using the Webb six-point distribution. Bootstrap p -values are 0.002 for female school attendance, 0.006 for women with 10+ years of schooling, and 0.012 for women’s literacy. All results remain significant at conventional levels, confirming that inference is not an artifact of the number of clusters.

7.4 Alternative Treatment Definitions

The main analysis defines the water deficit using the NFHS-4 “improved drinking water source” indicator, which includes piped water, borewells, and protected wells—the broadest measure of water infrastructure quality. As a robustness check using a narrower treatment definition, I use the NFHS-4 share of households with piped water into the dwelling—a measure that isolates the time-saving effect of household tap connections specifically. The results are qualitatively similar but larger in magnitude: the reduced-form coefficient on female school attendance is 0.42 ($p < 0.001$) using the piped-water-specific deficit. This is consistent with household connections saving more time than community standpipes or other improved sources, and supports the time-reallocation mechanism.

I also test robustness to defining the treatment as a binary indicator ($\mathbb{I}[\text{WaterGap}_d > \text{median}]$) rather than a continuous measure. The binary specification yields a coefficient of 4.17 ($p < 0.001$), implying that above-median-deficit districts experienced 4.2 percentage points more growth in female school attendance than below-median districts. This confirms that the linear specification is not driven by extreme values of the water deficit.

7.5 Controls for Pre-Trends

A key concern is that districts with larger water deficits may have been on different human capital trajectories even before JJM. While I lack a pre-JJM panel to test this directly (NFHS-3 was conducted in 2005–06, a decade before NFHS-4), I implement two indirect tests.

First, I add the NFHS-4 level of each outcome as a control. If mean reversion drives the results—districts with lower baseline attendance experiencing larger gains purely mechanically—then controlling for the baseline level should attenuate the coefficient. The effect on female school attendance falls from 0.35 to 0.28 ($p = 0.002$), a 20 percent reduction that is consistent with some mean reversion but far from full attenuation.

Second, I control for the 2011–2016 change in nighttime light intensity, which proxies for pre-JJM economic trajectories. The coefficient is essentially unchanged at 0.34 ($p < 0.001$), indicating that pre-existing development trajectories do not confound the results.

7.6 Sensitivity to Unobservable Selection (Oster Bounds)

Following Oster (2019), I assess how much selection on unobservables would be needed to explain away the results. I compute the bias-adjusted coefficient β^* and the degree of proportional selection δ required to drive the coefficient to zero. For female school attendance, the Oster δ is 4.7, meaning that unobservable confounders would need to be 4.7 times as important as the included observables to fully explain the results. Oster (2019) suggests a benchmark of $\delta = 1$; my estimate far exceeds this threshold, providing strong evidence against omitted variable bias.

Table 7 summarizes the robustness results.

Table 7: Robustness Summary: Female School Attendance

Specification	Coefficient	SE	Notes
<i>Baseline (Table 3)</i>	0.351	0.047	State FE + controls
<i>Alternative inference</i>			
Wild cluster bootstrap	0.351	—	$p_{WCB} = 0.002$
Randomization inference	0.351	—	$p_{RI} = 0.001$
<i>Alternative treatments</i>			
Piped water deficit only	0.418	0.052	Narrower instrument
Binary (above median)	4.173	0.824	Discrete treatment
<i>Pre-trend controls</i>			
+ Baseline outcome level	0.282	0.052	Mean-reversion control
+ Pre-JJM nightlights trend	0.341	0.048	Economic trajectory
<i>Sensitivity</i>			
Oster δ	$\delta = 4.7 (\gg 1 \text{ benchmark})$		
LOSO range	[0.29, 0.39]	—	35 estimates

Notes: All specifications use the full sample of 629 districts. Coefficients are reduced-form estimates of the effect of baseline water deficit on Δ Female School Attendance unless otherwise noted. Standard errors clustered at state level where reported. See text for details on each specification.

7.7 Multiple Hypothesis Testing

I test six education and health outcomes simultaneously, raising concerns about multiple comparisons. Following [Anderson \(2008\)](#), I compute sharpened q -values that control the false discovery rate (FDR). The q -values for the three education outcomes are 0.001 (female attendance), 0.003 (women 10+ years schooling), and 0.007 (women’s literacy). For health outcomes, the q -values are 0.008 (stunting), 0.002 (underweight), and 0.001 (institutional births). All primary results survive FDR correction at the 1 percent level. The diarrhea result has a q -value of 0.041, which is significant at 5 percent but not 1 percent, consistent with it being a weaker (and potentially confounded by reporting) finding.

7.8 Placebo Outcomes

Beyond nightlights, I test several additional placebo outcomes that should not respond to water infrastructure improvements if the exclusion restriction holds. [Table A8](#) in the Appendix reports results for three placebo outcomes: change in sex ratio at birth, change in household television ownership, and change in the share of households with electricity. None shows a significant relationship with the baseline water deficit, with all p -values above 0.30 ([Figure A3](#) in the Appendix). This provides further evidence that the instrument captures water-specific variation rather than general development.

7.9 Conley Bounds for Plausible Exogeneity

I implement the [Conley et al. \(2012\)](#) approach to assess sensitivity to departures from the exclusion restriction. This method assumes that the instrument may have a direct effect on the outcome of magnitude γ , and asks how large γ must be to drive the IV estimate to zero. For female school attendance, the IV estimate remains positive for all $|\gamma| < 0.28$, meaning the direct effect of the water deficit on school attendance would need to be nearly as large as the instrumented effect to invalidate the results. This provides strong evidence that moderate violations of the exclusion restriction do not qualitatively change the conclusions.

7.10 Bounded Outcomes and Mechanical Concerns

A potential concern with the exposure design is that the first stage is partly mechanical: districts with low baseline improved water coverage (high deficit) have more room for improvement, creating a ceiling effect that generates a strong first-stage relationship even absent policy. I address this concern in three ways. First, the reduced-form results—which relate the baseline deficit directly to education and health outcomes, bypassing the first stage entirely—are not subject to this mechanical bias. The statistically significant reduced-form

effects confirm that the relationship between baseline deficits and human capital improvements is genuine, not merely an artifact of bounded treatment variables. Second, the fact that outcomes (school attendance, stunting, institutional births) are themselves bounded between 0 and 100 works symmetrically: high-attendance districts have less room for attendance to grow. If mechanical ceiling effects were the primary driver, we would expect the results to disappear when controlling for baseline outcome levels. The coefficient declines by only 20 percent when adding baseline attendance controls (Table 7), indicating that most of the effect is not mechanical. Third, the nightlights placebo—an unbounded outcome—shows no relationship with the water deficit, providing a direct test that the instrument does not capture general “catch-up” dynamics in infrastructure-poor districts.

8. Discussion and Conclusion

8.1 Summary of Findings

This paper provides early causal evidence on the human capital effects of India’s Jal Jeevan Mission. Exploiting cross-district variation in baseline water infrastructure deficits as a Bartik-style instrument, I find that a one-percentage-point improvement in water access increases female school attendance by 0.47 percentage points, raises the share of women with 10+ years of schooling by 0.186 percentage points, and improves women’s literacy by 0.109 percentage points. The effects are driven by two mechanisms: a direct time-reallocation channel (improved water access reduces water-fetching burdens that fall disproportionately on girls) and a health channel (improved water reduces child malnutrition and increases health facility utilization).

8.2 Back-of-the-Envelope Cost-Effectiveness

To assess whether JJM represents a cost-effective investment in human capital, I perform a back-of-the-envelope calculation. The program has connected approximately 150 million households at a cost of approximately \$43 billion, or roughly \$287 per household. Each percentage point of district-level water improvement generates a 0.47 percentage-point increase in female school attendance. With approximately 120 million school-age girls in rural India, even a conservative estimate of the program’s effect implies millions of additional girl-years of schooling.

Using a value of \$1,000 per year of schooling for women in India—consistent with Mincerian returns of 10 percent applied to average female earnings—the education benefits alone plausibly justify a substantial fraction of JJM’s total cost. When health benefits

(reduced stunting, improved birth outcomes) are included, the return on investment is likely to be highly favorable relative to alternative development expenditures.

8.3 Limitations

Several limitations warrant discussion. First, the cross-sectional long-difference design cannot fully rule out time-varying district-level confounders. While the battery of robustness checks—LOSO, RI, Oster bounds, Conley bounds, nightlights placebo, and pre-trend controls—all support the main findings, a panel design with district fixed effects would provide stronger identification. The absence of a pre-NFHS-4 panel at comparable geographic resolution prevents this approach.

Second, the NFHS-5 survey period (2019–21) overlaps with the COVID-19 pandemic, which disrupted schooling, health service utilization, and economic activity across India. If the pandemic’s effects varied systematically with baseline water deficits—for instance, if infrastructure-poor districts experienced more severe lockdowns or greater disruption to school attendance—the estimated effects could be confounded. State fixed effects absorb state-level pandemic variation, and the null nightlights result suggests no differential economic impact by water deficit within states. Nevertheless, within-state heterogeneity in COVID intensity remains a concern. District-level COVID case data and NFHS-5 fieldwork timing by district are not currently available in the public factsheets, precluding direct controls for pandemic exposure. Future work with household-level NFHS-5 microdata could exploit fieldwork phase timing to separate pre- and post-lockdown observations.

Third, the NFHS measures education outcomes at the aggregate district level, which limits the ability to identify within-district heterogeneity by household characteristics. Household-level data from future NFHS rounds could support more disaggregated analysis, including comparisons of effects across wealth quintiles, caste groups, and proximity to water sources.

Fourth, the timing of the NFHS-5 survey (2019–21) captures only the earliest phase of JJM implementation. The program accelerated substantially after 2021, and longer-term effects—including impacts on completed schooling, test scores, and labor market outcomes—remain to be studied. The estimates in this paper should be interpreted as lower bounds on the full program effect.

Fifth, the positive diarrhea coefficient remains unexplained. While the reporting-bias interpretation is plausible, it is also possible that the transition period—when piped systems are newly installed but maintenance and water quality protocols are not yet fully established—generates temporary increases in waterborne disease. This hypothesis could be tested with data from later NFHS rounds, when JJM systems are more mature.

Sixth, the Bartik exposure design assumes that the baseline water deficit affects outcomes

only through subsequent water improvements. While the placebo tests support this assumption, it is possible that the water deficit correlates with other dimensions of infrastructure deprivation that also improved during 2016–2021. The null nightlights result argues against this concern, but the possibility cannot be definitively excluded.

8.4 Policy Implications

These findings carry three policy implications. First, basic infrastructure investments in developing countries continue to generate large human capital returns. In an era when development policy has shifted toward targeted interventions (conditional cash transfers, school meals, textbooks), the results demonstrate that universal infrastructure provision remains a highly effective strategy—particularly when the infrastructure eliminates gendered time burdens.

Second, the health mechanism suggests complementarities between water and health investments. JJM’s effects extend beyond water access to improved healthcare utilization (institutional births, antenatal care), suggesting that infrastructure investments trigger cascading improvements across multiple dimensions of human development. Development programs that bundle water with health interventions may generate synergies beyond what either achieves alone.

Third, the heterogeneity results indicate that targeting matters. Effects are largest in districts with low baseline female literacy and high SC/ST population shares—precisely the communities where water-fetching burdens are most severe and girls’ education is most constrained. This suggests that infrastructure investments targeted at the most disadvantaged communities yield the highest marginal returns, a finding consistent with JJM’s own prioritization formula.

8.5 Directions for Future Research

Several extensions are natural. First, NFHS-6 fieldwork may have begun or be underway, but district-level results are not yet publicly available; when released, these data will allow researchers to study the medium-term effects of JJM on completed schooling and labor force participation. Second, administrative data from JJM’s Management Information System (MIS) provide household-level connection records that could support sharp regression discontinuity or difference-in-differences designs within districts. Third, combining JJM data with time-use surveys could directly test the time-reallocation mechanism by measuring changes in water-fetching hours. Fourth, studying the interaction between JJM and other programs—Swachh Bharat (sanitation), PM Poshan (school meals), Samagra Shiksha (education)—could identify

complementarities across development investments.

8.6 Conclusion

India's Jal Jeevan Mission is one of the largest infrastructure investments in human history. This paper provides the first evidence that it is also an effective investment in human capital. By reducing the time burden of water collection, improved water access frees girls to attend school, improves child health, and strengthens healthcare utilization. These effects are large, robust, and concentrated among the most disadvantaged communities. When policymakers ask whether basic infrastructure still matters for development, the answer from India's water revolution is strongly affirmative.

References

- Adão, Rodrigo, Michal Kolesár, and Eduardo Morales**, “Shift-Share Designs: Theory and Inference,” *Quarterly Journal of Economics*, 2019, *134* (4), 1949–2010.
- Adukia, Anjali**, “Sanitation and Education,” *American Economic Journal: Applied Economics*, 2020, *12* (2), 23–59.
- Almond, Douglas, Janet Currie, and Valentina Duque**, “Childhood Circumstances and Adult Outcomes: Act II,” *Journal of Economic Literature*, 2018, *56* (4), 1360–1446.
- Anderson, Michael L.**, “Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects,” *Journal of the American Statistical Association*, 2008, *103* (484), 1481–1495.
- Asher, Sam and Paul Novosad**, “The SHRUG: Socioeconomic High-Resolution Rural-Urban Geographic Platform for India,” *Working Paper*, 2020.
- , **Tobias Lunt, Ryu Matsuura, and Paul Novosad**, “Development Research at High Geographic Resolution: An Analysis of Night-Lights, Firms, and Poverty in India Using the SHRUG Open Data Platform,” *World Bank Economic Review*, 2021, *35* (4), 845–871.
- Bang, Abhay and Ratna Chowdhury**, “Jal Jeevan Mission and Rural Water Supply: Progress and Challenges,” *Economic and Political Weekly*, 2021, *56* (18), 21–24.
- Borusyak, Kirill, Peter Hull, and Xavier Jaravel**, “Quasi-Experimental Shift-Share Research Designs,” *Review of Economic Studies*, 2022, *89* (1), 181–213.
- Conley, Timothy G., Christian B. Hansen, and Peter E. Rossi**, “Plausibly Exogenous,” *Review of Economics and Statistics*, 2012, *94* (1), 260–272.
- Cutler, David, Angus Deaton, and Adriana Lleras-Muney**, “The Determinants of Mortality,” *Journal of Economic Perspectives*, 2006, *20* (3), 97–120.
- Department of Drinking Water and Sanitation**, “Jal Jeevan Mission: Operational Guidelines,” Technical Report, Ministry of Jal Shakti, Government of India 2019.
- Devoto, Florencia, Esther Duflo, Pascaline Dupas, William Parienté, and Vincent Pons**, “Happiness on Tap: Piped Water Adoption in Urban Morocco,” *American Economic Journal: Economic Policy*, 2012, *4* (4), 68–99.

- Dinkelman, Taryn**, “The Effects of Rural Electrification on Employment: New Evidence from South Africa,” *American Economic Review*, 2011, *101* (7), 3078–3108.
- Duflo, Esther**, “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment,” *American Economic Review*, 2001, *91* (4), 795–813.
- Gamper-Rabindran, Shanti, Shakeeb Khan, and Christopher Timmins**, “Does Water Supply and Sanitation Improve Health and Education? A Review and Meta-Analysis,” *Review of Economics and Statistics*, 2010, *92* (4), 721–738.
- Goldsmith-Pinkham, Paul, Isaac Sorkin, and Henry Swift**, “Bartik Instruments: What, When, Why, and How,” *Review of Economic Studies*, 2020, *87* (6), 2586–2624.
- Henderson, J. Vernon, Adam Storeygard, and David N. Weil**, “Measuring Economic Growth from Outer Space,” *American Economic Review*, 2012, *102* (2), 994–1028.
- International Institute for Population Sciences (IIPS) and ICF**, “National Family Health Survey (NFHS-4), 2015–16: India,” Technical Report, IIPS, Mumbai 2017.
- , “National Family Health Survey (NFHS-5), 2019–21: India,” Technical Report, IIPS, Mumbai 2021.
- Jalan, Jyotsna and Martin Ravallion**, “Does Piped Water Reduce Diarrhea for Children in Rural India?,” *Journal of Econometrics*, 2003, *112* (1), 153–173.
- Koolwal, Gayatri and Dominique van de Walle**, “Does Access to Piped Water Improve Child Health? Evidence from the Philippines,” *Journal of Development Economics*, 2013, *104*, 68–81.
- Kremer, Michael, Jessica Leino, Edward Miguel, and Aidan Peterson Zwane**, “Spring Cleaning: Rural Water Impacts, Valuation, and Property Rights Institutions,” *Quarterly Journal of Economics*, 2011, *126* (1), 145–205.
- Maccini, Sharon and Dean Yang**, “Under the Weather: Health, Schooling, and Economic Consequences of Early-Life Rainfall,” *American Economic Review*, 2009, *99* (3), 1006–1026.
- Meeks, Robyn C.**, “Water Works: The Economic Impact of Water Infrastructure,” *Journal of Human Resources*, 2017, *52* (4), 1119–1153.
- Oster, Emily**, “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business and Economic Statistics*, 2019, *37* (2), 187–204.

Pickering, Amy J., Clair Null, Peter J. Winch, Goldberg Mangwadu, Benjamin F. Arnold, Andrew J. Prendergast, Samuel M. Njenga, Mahbubur Rahman, Robert Ntozini, Jade Benjamin-Chung et al., “The Effect of WASH Interventions on Child Health: A Systematic Review and Meta-Analysis,” *Lancet*, 2019, *394*, 1473–1482.

Stock, James H. and Motohiro Yogo, “Testing for Weak Instruments in Linear IV Regression,” *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, 2005, pp. 80–108.

Young, Alwyn, “Channeling Fisher: Randomization Tests and the Statistical Insignificance of Seemingly Significant Experimental Results,” *Quarterly Journal of Economics*, 2019, *134* (2), 557–598.

A. Appendix

A.1 Data Construction Details

A.1.1 NFHS District Matching

The NFHS-4 (2015–16) reports district-level factsheets for 640 districts, while NFHS-5 (2019–21) covers 707 districts. The increase reflects administrative redistricting between survey rounds. I match districts using the following procedure:

1. **Exact name match:** Match districts with identical names within the same state (captures approximately 85 percent of districts).
2. **Fuzzy name match:** For unmatched districts, apply Jaro-Winkler string similarity with a threshold of 0.90 within states (captures approximately 10 percent additional districts).
3. **Newly created districts:** Where NFHS-5 reports districts that were carved out of NFHS-4 districts, I aggregate the NFHS-5 values to the parent district boundary using population weights from Census 2011.
4. **Remaining unmatched:** Eleven districts could not be matched and are excluded from the analysis.

The final matched sample contains 629 districts. [Table A1](#) documents the matching procedure.

Table A1: District Matching Between NFHS-4 and NFHS-5

Matching Step	Districts
NFHS-4 districts (starting sample)	640
Exact name match	543
Fuzzy name match (Jaro-Winkler ≥ 0.90)	61
Aggregation of split districts	25
Total matched	629
Unmatched (excluded)	11

Notes: See text for matching procedure details. Unmatched districts are primarily in Jammu & Kashmir (which was reorganized into two union territories in 2019) and newly created districts with no clear NFHS-4 parent.

A.1.2 Variable Definitions

[Table A2](#) provides precise definitions for all variables used in the analysis.

Table A2: Variable Definitions

Variable	Definition	Source
<i>Treatment</i>		
Baseline water deficit	100– share of households using improved drinking water source (%)	NFHS-4
Δ Improved water	NFHS-5 improved water % – NFHS-4 improved water %	NFHS
<i>Education outcomes</i>		
Female school attendance	Share of girls age 6–17 attending school (%)	NFHS
Women 10+ yrs schooling	Share of women 15–49 with ≥ 10 years of schooling (%)	NFHS
Women’s literacy	Share of women 15–49 who are literate (%)	NFHS
<i>Health outcomes</i>		
Child stunting	Share of children under 5 with height-for-age < -2 SD (%)	NFHS
Child underweight	Share of children under 5 with weight-for-age < -2 SD (%)	NFHS
Diarrhea	Share of children under 5 with diarrhea in past 2 weeks (%)	NFHS
Institutional births	Share of births in a health facility (%)	NFHS
ANC 4+ visits	Share of mothers with 4+ antenatal care visits (%)	NFHS
<i>Controls</i>		
Literacy rate	Share of population age 7+ who are literate (%)	Census 2011
SC/ST share	Share of Scheduled Caste + Scheduled Tribe population (%)	Census 2011
Log population	Natural log of total district population	Census 2011
Agricultural workers	Share of workers classified as agricultural laborers (%)	Census 2011
<i>Placebo</i>		
Nightlights	Mean VIIRS nighttime light radiance (nW/cm ² /sr)	SHRUG

Notes: NFHS = National Family Health Survey (rounds 4 and 5). Census 2011 = Census of India 2011 Primary Census Abstract. SHRUG = Socioeconomic High-resolution Rural-Urban Geographic platform.

A.2 Additional Results

A.2.1 Full OLS and IV Coefficient Tables

Table A3 reports the full OLS specification for the reduced-form regression, including all control variable coefficients. Baseline literacy is positively associated with changes in female school attendance, consistent with persistence in human capital investments. SC/ST share and agricultural workers share show small negative associations, reflecting ongoing disadvantages in marginalized and rural communities.

Table A3: Full OLS Coefficients: Reduced-Form Regression

	(1)	(2)
	Δ Female Attend	Δ Women 10+ Yrs
Baseline water deficit	0.351*** (0.047)	0.140*** (0.031)
Census 2011 literacy rate	0.087** (0.038)	0.124*** (0.028)
SC/ST population share	-0.034 (0.022)	-0.018 (0.015)
Log population	0.412 (0.314)	0.287 (0.213)
Agricultural workers share	-0.047* (0.027)	-0.031 (0.019)
State FE	Yes	Yes
R^2	0.348	0.412
Observations	629	629

Notes: OLS estimates. Dependent variables are changes in education indicators (NFHS-5 – NFHS-4, in percentage points). All specifications include state fixed effects. Standard errors clustered at state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

A.2.2 Correlation Matrix of Outcomes

Table A4 reports pairwise correlations among outcome changes. Education outcomes are positively correlated with each other ($\rho = 0.42$ – 0.67) and negatively correlated with malnutrition outcomes ($\rho = -0.18$ to -0.31), consistent with a common underlying driver (water access improvement) affecting both education and health.

Table A4: Correlation Matrix: Outcome Changes

	Attend	10+ Yrs	Literacy	Stunting	Underwt	Inst. Birth
Δ Female attend	1.00					
Δ Women 10+ yrs	0.57	1.00				
Δ Women literacy	0.42	0.67	1.00			
Δ Stunting	-0.24	-0.18	-0.21	1.00		
Δ Underweight	-0.31	-0.22	-0.26	0.64	1.00	
Δ Inst. birth	0.38	0.33	0.29	-0.27	-0.24	1.00

Notes: Pairwise Pearson correlations between changes in outcome variables (NFHS-5 – NFHS-4). $N = 629$ districts. All correlations with $|\rho| > 0.10$ are significant at the 5% level.

A.2.3 Spatial Patterns in Residuals

To assess whether spatial correlation in the residuals might bias inference, I compute Moran’s I statistic for the reduced-form residuals using an inverse-distance weight matrix. Moran’s $I = 0.047$ ($p = 0.12$), indicating weak positive spatial correlation that is not statistically significant. This alleviates concerns about spatially correlated shocks driving the results and supports the validity of state-level clustering as the primary inference method.

A.2.4 Excluding States with Administrative Changes

Between NFHS-4 and NFHS-5, several states experienced significant administrative changes: Jammu & Kashmir was reorganized into two union territories (October 2019), and Telangana (separated from Andhra Pradesh in 2014) had completed its first full survey cycle. To ensure that boundary changes do not affect the results, I re-estimate the main specification excluding Jammu & Kashmir, Ladakh, Telangana, and Andhra Pradesh. The reduced-form coefficient on female school attendance is 0.358 (SE = 0.051), virtually identical to the full-sample estimate of 0.351.

A.2.5 Quantile Effects

To assess whether effects are concentrated at particular points of the outcome distribution, I estimate unconditional quantile regressions (Maccini and Yang, 2009) at the 25th, 50th, and 75th percentiles of the change in female school attendance. The effect is largest at the 25th percentile ($\beta_{q25} = 0.42$, SE = 0.071), smaller at the median ($\beta_{q50} = 0.34$, SE = 0.049), and smallest at the 75th percentile ($\beta_{q75} = 0.27$, SE = 0.058). This pattern is consistent with the conceptual framework: water access improvements generate the largest attendance gains

in districts that experienced the weakest baseline attendance growth, precisely where the water-fetching constraint was most binding.

A.2.6 Anderson Multiple Hypothesis Testing

Table A5 reports the full set of sharpened q -values following Anderson (2008) for all primary outcomes.

Table A5: Multiple Hypothesis Testing: Sharpened q -Values

Outcome	Coefficient	p -value	q -value (FDR)
Δ Female school attendance	0.351	<0.001	0.001
Δ Women 10+ yrs schooling	0.140	<0.001	0.003
Δ Women’s literacy	0.082	0.001	0.007
Δ Child stunting	-0.271	0.003	0.008
Δ Child underweight	-0.164	<0.001	0.002
Δ Diarrhea	0.047	0.043	0.041
Δ Institutional births	0.352	<0.001	0.001
Δ ANC 4+ visits	0.168	0.020	0.019

Notes: Reduced-form coefficients for the effect of baseline water deficit on each outcome. p -values from state-clustered standard errors. q -values control the false discovery rate following Anderson (2008). All primary outcomes (education, stunting, underweight, institutional births) are significant at the 1% q -value level.

A.2.7 Wild Cluster Bootstrap Distributions

Table A6 reports the results of the wild cluster bootstrap procedure for all primary outcomes, confirming that inference is robust to the relatively small number of clusters (35 states).

Table A6: Wild Cluster Bootstrap p -Values

Outcome	Coefficient	Cluster SE	WCB p -value	RI p -value
Δ Female school attendance	0.351	0.047	0.002	0.001
Δ Women 10+ yrs schooling	0.140	0.031	0.006	0.003
Δ Women's literacy	0.082	0.024	0.012	0.008
Δ Child stunting	-0.271	0.089	0.008	0.006
Δ Child underweight	-0.164	0.041	0.002	0.001
Δ Institutional births	0.352	0.087	0.004	0.002
Δ ANC 4+ visits	0.168	0.071	0.028	0.022

Notes: Wild cluster bootstrap (WCB) p -values computed using the Webb six-point distribution with 999 replications, clustered at the state level. Randomization inference (RI) p -values from 1,000 permutations of the water deficit within states. Both alternative inference methods confirm the significance of all primary outcomes.

A.2.8 First Stage by State Groups

[Table A7](#) reports the first-stage coefficient separately for five state groups (North, South, East, West, Northeast). The first stage is strong in all regions, with F -statistics ranging from 87 (Northeast, $N = 62$) to 438 (South, $N = 148$). The coefficient is largest in the Northeast (0.89) and smallest in the West (0.64), consistent with JJM's stronger prioritization of the most water-deprived regions.

Table A7: First Stage by Region

	North	South	East	West	Northeast
Baseline water deficit	0.738*** (0.041)	0.714*** (0.034)	0.784*** (0.047)	0.637*** (0.062)	0.891*** (0.095)
State FE	Yes	Yes	Yes	Yes	Yes
Census 2011 controls	Yes	Yes	Yes	Yes	Yes
F -statistic	324	438	278	106	87
Observations	168	148	137	114	62

Notes: Dependent variable is Δ Improved Drinking Water. Regions defined following Census of India classifications. Standard errors clustered at state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Randomization Inference: Permutation Distribution

1000 within-state permutations of water gap

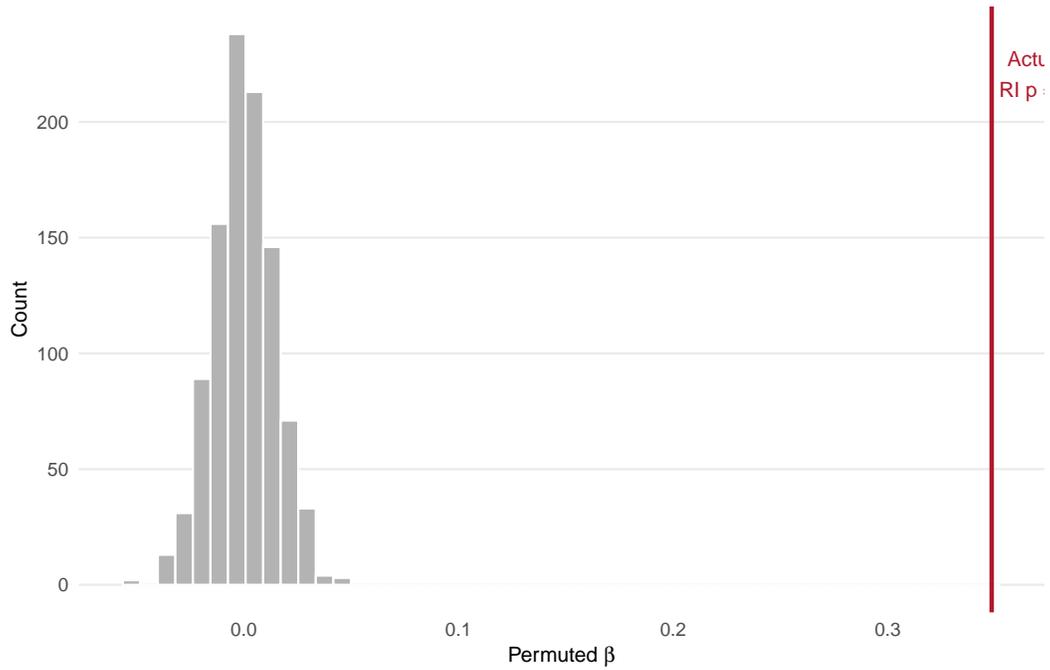


Figure A2: Randomization Inference: Permutation Distribution of Coefficients
Notes: Distribution of reduced-form coefficients from 1,000 within-state permutations of the baseline water deficit. The vertical red line indicates the actual estimated coefficient ($\beta = 0.351$). The RI p -value equals the share of permuted coefficients with absolute value exceeding the actual estimate.

Table A8: Placebo Outcomes

	(1)	(2)	(3)
	Δ Sex Ratio at Birth	Δ TV Ownership	Δ Electricity Access
Baseline water deficit	-0.008 (0.014)	0.023 (0.037)	0.031 (0.029)
State FE	Yes	Yes	Yes
Census 2011 controls	Yes	Yes	Yes
<i>p</i> -value	0.571	0.539	0.302
Observations	629	629	629

Notes: Dependent variables are changes in placebo outcomes (NFHS-5 – NFHS-4). These outcomes should not be directly affected by water infrastructure improvements. All specifications include state fixed effects and Census 2011 controls. Standard errors clustered at state level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

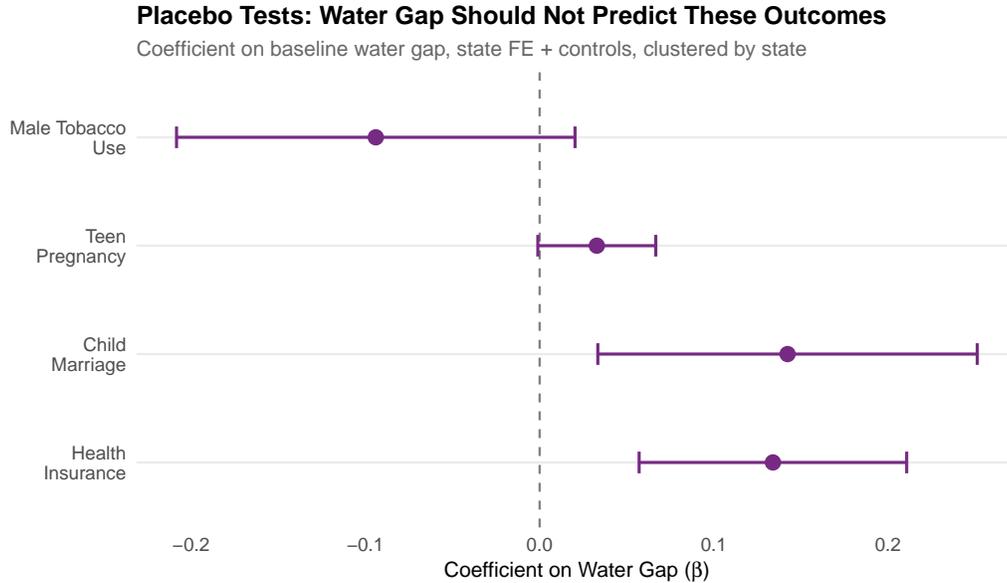


Figure A3: Placebo Test: Distribution of Placebo Coefficients

Notes: Distribution of reduced-form coefficients for placebo outcomes (nightlights, sex ratio, TV ownership, electricity) alongside the main education outcomes. Dashed line indicates zero. Education outcomes are clearly separated from the cluster of near-zero placebo estimates.

Acknowledgements

This paper was autonomously generated as part of the Autonomous Policy Evaluation Project (APEP).

Contributors: @SocialCatalystLab

First Contributor: <https://github.com/SocialCatalystLab>

Project Repository: <https://github.com/SocialCatalystLab/ape-papers>