Discussion Paper Series

IZA DP No. 18771

July 2026

Private Schooling and Fertility Decline in India

Daniel Fernández Kranz
IE University and IZA@LISER

Trong-Anh Trinh
Monash University

IZA@LISER NETWORK
ISSN: 2365-9793

IZA DP No. 18771 July 2026

Private Schooling and Fertility Decline in India

Abstract

India reached below-replacement fertility (TFR = 2.0) despite low income, persistent traditional gender norms, and declining female labour force participation. We study how the expansion of a decentralized, fee-based education market reshaped fertility decisions by increasing households' exposure to the cost of child quality. Because sons are more likely than daughters to be enrolled in private schools, this cost channel interacts with son preference to generate gender-differentiated fertility stopping behavior. We develop a dynamic model of parity progression that yields testable predictions on the joint role of child gender composition and local schooling markets. We test these predictions using survey data from NFHS (2015–21) merged with administrative district-level UDISE data on 37,251 women in 525 districts at the parity-2-to-3 margin. We find that households with only daughters are less responsive to local private-school prevalence than otherwise similar households with sons. This result is stable across subsamples, the exclusion of recent migrants, controls for modernization proxies, and controls for differential baseline-fertility trends. At the district level, private-school expansion is associated with lower parity progression, but this effect is sensitive to district-level confounding.

JEL classification

J13, I25, O15, D13

Keywords

fertility, private schooling, household decision-making, developing economies

Corresponding author

Daniel Fernández Kranz

Daniel.Fernandez@ie.edu

1 Introduction

India recently reached below-replacement fertility (TFR = 2.0) despite low income, persistent traditional gender norms, and declining female labour force participation[1]—a pattern that sits uneasily with the standard demographic transition, which emphasizes rising income, female labor market opportunities, and the opportunity cost of time as central drivers (Becker, 1960). Over the same period, India experienced a rapid expansion of private schooling that transformed the education sector from a largely uniform, publicly provided system into a vertically differentiated market in which higher-quality education is increasingly available at a positive price. The expansion of the private schooling sector happened through decentralized entrepreneurial entry, supported by a low-cost production model based on locally hired teachers (Kingdon, 2020; Tooley, 2009). The result has been a large, fragmented market characterized by localized competition, substantial heterogeneity in price and quality, and limited scale economies. The conjunction of these two developments—a surprising fertility decline alongside a rapidly spreading market for fee-based education—has begun to attract attention well beyond academia: recent commentary on India's "surprise baby bust" points directly to an *educational arms race*, in which the spread of fee-paying schooling (now attended by close to two-fifths of children) and rising per-child education spending raise the effective cost of children and push fertility below replacement (The Economist, 2026). This paper formalises and tests precisely this channel, with particular emphasis on its interaction with son preference.

This expansion occurred in a context of well-documented deficiencies in the public education system. High rates of teacher absenteeism, weak infrastructure, and low learning outcomes in government schools (PROBE Team, 1999; Kremer et al., 2005) generated strong entry opportunities for alternatives. Private schools emerged as a competitive response, offering higher perceived quality and greater accountability, often at relatively low cost (Muralidharan and Kremer, 2008; Muralidharan and Sundararaman, 2015; Azam et al., 2016). As a result, households increasingly faced a menu of schooling options that differed in both price and quality, transforming education from a largely free public service into a differentiated market good.

In this paper, we argue that this change in the structure of the education market has implications for fertility decisions. In a setting with son preference, where boys are more likely than girls to be enrolled in private schools (Kaul, 2018; Nandi et al., 2023; Datta and Kingdon, 2019), and where private-school growth at the district level has been shown to widen—rather than narrow—the gender gap in enrolment (Kingdon and Pal, 2010), the expansion of fee-based schooling creates a gender-differentiated cost of child quality. Even with uniform posted fees, the effective cost of investing in children varies by gender

[1] Female labor force participation in India declined from approximately 34–35 percent in the mid-2000s to around 26 percent by 2020, before partially recovering to about 32–33 percent in recent years, remaining below its earlier peak (World Bank, 2024).

due to differential enrollment choices. This interaction between market structure and household preferences generates systematic differences in fertility stopping behavior.

To formalize this mechanism, we develop a dynamic model of fertility in which households make sequential parity[2] progression decisions under uncertainty about child gender and future schooling costs. The model incorporates three key elements: (i) diminishing marginal utility of additional children, (ii) son preference, and (iii) endogenous investment in child quality through schooling. The presence of a private schooling market increases the marginal cost of children for households that invest in quality, but this effect is heterogeneous across households depending on their existing child composition.

The model yields two main predictions. First, an increase in the prevalence of private schooling reduces fertility at higher parities by raising the cost of additional children. Second, this effect interacts with gender composition: households with only daughters are less deterred from continuing fertility than households with sons. The mechanism is not that the expected cost of the *next* child differs by family type—it does not, since the next child is equally likely to be a boy or a girl. Rather, the mechanism operates through the *existing* cost burden: families with sons already enrolled in private schools face higher total schooling expenditures, which tighten the household budget and raise the effective (shadow) cost of any additional child. Families with only daughters, whose children attend cheaper government schools, have more budgetary slack and are therefore less sensitive to the expansion of private schooling.

We test these predictions using data from the National Family Health Surveys (NFHS-4 and NFHS-5), combined with district-level information on school composition from the Unified District Information System for Education (UDISE). Our empirical strategy focuses on contemporaneous parity progression decisions, defined as whether a woman has an additional birth within a fixed time window prior to the survey, conditional on being at risk at the beginning of that window. This approach aligns fertility decisions with the contemporaneous schooling environment and mitigates concerns related to long-term retrospective measures of completed fertility.

We exploit variation in private school prevalence across districts and over time, controlling for district fixed effects, survey-wave effects, and age profiles. While we do not claim that variation in private school prevalence is fully exogenous, the institutional context suggests that much of this variation is driven by supply-side factors, including local labor market conditions for teachers (Muralidharan and Kremer, 2008) and heterogeneous regulatory environments following the Right to Education Act (Shah and Steinberg, 2019). Importantly, our identification strategy relies on the interaction between within-district changes in private school penetration and household-level gender composition,

[2] Parity progression refers to the probability that a woman with n children has an additional birth (i.e., transitions from parity n to $n+1$), conditional on having reached parity n. It is a standard measure used to study fertility behavior at different margins.

which is less likely to be confounded by district-level omitted variables. Furthermore, we show that this interaction-based evidence is stable across subsamples, alternative observation windows, and the exclusion of recent migrants, alleviating selective migration concerns. Finally, balance tests indicate that within-state private-school expansion is not systematically predicted by baseline fertility preferences or gender norms, and a supply-side consistency check based on predetermined within-district changes in teacher supply produces consistent results, suggestive of the relevance of supply-side factors.

Our results provide consistent support for our model's proposed mechanisms, with the strongest and most robust evidence coming from within-district variation across household types. At the parity 2 to 3 margin, households with only daughters are significantly more likely to continue childbearing than otherwise-similar households with sons, and this gap is amplified where private schooling is more prevalent—precisely the gender-differentiated response the model predicts when fee-charging schooling raises the cost of sons more than that of daughters. Because this interaction is identified from comparisons among households in the same district and survey wave, it is insulated from district-level confounders, and it proves stable across subsamples, observation windows, the exclusion of recent migrants, the inclusion of controls for differential baseline-fertility trends, and a supply-side consistency check.

At the district level, higher private school prevalence is associated with significantly lower probabilities of an additional birth, and the relationship is monotone in the size of the local private-school expansion. Taken at face value, the estimated cost-channel coefficient implies that the observed 1.3-percentage-point increase in district private-school share between waves reduced the 36-month probability of a third birth by about 0.4 percentage points—on the order of one-sixth of the observed decline in parity-2-to-3 progression over this period. We caution, however, that this aggregate magnitude is sensitive to district-level confounding: districts that expanded private schooling started from higher fertility, and once we allow districts on different baseline-fertility trajectories to trend differentially, the district-level coefficient attenuates substantially and loses significance (Section 7.8). We therefore treat the aggregate decomposition as illustrative of the channel's potential scale rather than as a precise causal estimate, and rest our causal interpretation on the within-district interaction.

We also document that the effect of private school prevalence is concentrated at higher parities and is absent at earlier fertility margins, consistent with the idea that schooling costs become salient only once children reach school age. Additional heterogeneity analyses show that the results are stronger among Hindu households—where son–preference is more acute—and broadly consistent across age groups. We also find suggestive evidence that the effects are larger for lower-income households, although these estimates are imprecisely measured.

Taken together, our findings highlight a new channel through which market expansion in a service sector can generate equilibrium spillovers in household behavior. The growth of private schooling in India did not affect only educational outcomes; it also altered the economic tradeoffs underlying fertility decisions by introducing a priced dimension to child quality that interacts with gender norms and household constraints.

This paper contributes to several strands of the literature.

First, we contribute to the large literature on private school expansion in India. Evidence on learning outcomes is mixed: Muralidharan and Sundararaman (2015) found that voucher recipients in Andhra Pradesh scored 0.13 SD higher at one-quarter of government per-pupil cost; Singh (2015) found positive effects on English but not mathematics; and Chudgar and Quin (2012) found no significant private-school advantage after propensity score matching. Evidence from Pakistan reinforces the complexity of these markets: Andrabi et al. (2008) documented the possibilities and limits of low-cost private schooling; Andrabi et al. (2020) showed that grants to private schools improved scores by 0.3 SD; Andrabi et al. (2025) demonstrated that the private premium is concentrated among higher-quality providers; and Andrabi et al. (2024) found that public spending can crowd *in* private quality. On competition and stratification, Bagde et al. (2022) showed that private school entry induced sorting by gender, income, and caste without harming public school students, a pattern confirmed by Reddy et al. (2026) for urban caste-based segregation. On equity, Srivastava and Noronha (2016) showed that hidden costs under the RTE Act disproportionately burdened disadvantaged households, while Srivastava (2013) found that low-fee private schools mitigate but do not eliminate inequity. On teacher labour markets, Chudgar and Sakamoto (2021) documented that private school teachers earn substantially less than government counterparts despite similar work, enabling the low-cost model that drives expansion (Kingdon, 2020). What remains largely unexplored is the effect of private school expansion on outcomes *beyond* the education sector. Apart from work on gendered household expenditure patterns (Azam et al., 2016; Datta and Kingdon, 2019), this literature has focused on learning, equity, and competition within schooling markets. Our paper fills this gap by showing that the expansion of fee-based education affects household fertility decisions through the interaction of gendered schooling costs with son preference and budget constraints.

Second, we contribute to the rather limited literature linking education costs and fertility decisions in developing countries. Burlando and Bbaale (2022) provide causal evidence that education costs influence reproductive behavior. After Uganda's Universal Primary Education reform abolished fees for up to 4 children per household, the paper finds that families exceeding the 4-child cap had lower fertility (4.2 pp less likely to have another child). That is, higher marginal schooling costs reduced fertility, consistent with the quality–quantity tradeoff. Similarly, Duflo et al. (2015) find that lowering schooling costs in Kenya delays childbearing and marriage among adolescent girls. More recently, Mahler et al. (2025) develop a theoretical framework in which education expenditures are shaped by social comparisons, generating an "educational arms race" that raises the effective cost of children. In the most closely related work, Meng et al. (2025) study the mirror image of our mechanism: China's 2021 ban on competition-driven private tutoring—which reduced educational competition and costs—increased expected total fertility by 7–8%, with the primary channel being a perceived reduction in educational competition. Our paper studies the opposite direction: how the *expansion* of fee-based schooling raises the cost of child quality and lowers fertility. Together, the two papers provide complementary evidence—from opposite policy directions and different countries—that education market competition is a first-order determinant of fertility decisions.

Third, we contribute to the standard economic framework of fertility decline, originating in Becker (1960) and extended by Becker et al. (1990), which emphasizes rising income, increasing returns to human capital, and the opportunity cost of time—particularly for women—as key drivers of the shift from quantity to quality. While this framework has been highly influential, India's experience departs from its core predictions. Chatterjee and Vogl (2018) document that fertility decline in developing countries typically accompanies economic growth, yet India's transition has followed an alternative pathway, reaching below-replacement fertility without the standard preconditions of delayed marriage, high female labor force participation, or rising income (Park et al., 2023). Earlier work by Drèze and Murthi (2001) shows that within India, fertility variation is more strongly associated with female literacy and child mortality than with income or urbanization. More recently, Park et al. (2023) highlight that India's transition differs not only in timing but also in mechanism. Using DHS data, they show that fertility decline has occurred primarily through early stopping of childbearing rather than delayed marriage or delayed first birth, a mechanism that is central to our analysis.

Fourth, our analysis builds on the quality–quantity framework of Becker and Lewis (1973), in which increases in the price of child quality raise the marginal cost of additional children. A large empirical literature tests this mechanism. Early work by Rosenzweig and Wolpin (1980) found evidence of tradeoffs in India, while studies using twin births and other sources of exogenous variation have produced null or opposite results, such as Black et al. (2005) and Angrist et al. (2010) for Norway and Israel, respectively. More recently, Bhalotra and Clarke (2020) show that OLS estimates and twin-IV estimates bound the true effect from opposite directions. For India, Kugler and Kumar (2017) provide evidence that larger family size reduces children's educational attainment, particularly among poorer and rural households. What remains less understood is what makes the quality–quantity tradeoff binding in a given context. We propose that the expansion of private schooling—which transformed education from a free public good into a priced, differentiated service—raised the effective cost of quality and activated the tradeoff for a large share of households.

Finally, we contribute to the large literature that documents the central role of son preference in shaping fertility decisions in India. Jayachandran (2017) shows that fertility decline can exacerbate sex imbalance by increasing the likelihood that families without sons resort to sex selection. Jayachandran and Pande (2017) show that eldest-son preference generates a birth-order gradient in child stunting that explains much of the India–Africa height gap, while Jayachandran and Kuziemko (2011) document that mothers breastfeed daughters less than sons in order to conceive again sooner—a mechanism through which son-driven fertility behavior directly harms girls' health. More recently, Jayachandran (2023) shows that son preference in India increasingly operates through fertility stopping rules, while Anukriti (2018) demonstrates that policies reducing fertility without addressing son preference can worsen sex ratios. Our contribution is to show that the expansion of private schooling introduces a new channel through which son preference affects fertility. While gender gaps in educational attainment have long been documented (Kingdon, 2002), the channel of gender bias has shifted from enrolment to school choice as private schooling expanded (Datta and Kingdon, 2019). By creating gender-differentiated costs of child quality, this shift generates systematic interactions between household composition and local market conditions that shape parity progression decisions. While prior work on private schooling in India has focused on school choice and learning outcomes (Kingdon, 2020; Muralidharan and Sundararaman, 2015), our analysis shows that the growth of a differentiated, fee-based education sector has consequences well beyond the education sector itself, altering the economic environment in which fertility decisions are made.

The rest of the paper is organized as follows. Section 2 describes the institutional setting and characterizes the expansion of the private schooling sector as a fragmented, differentiated market. Section 3 describes the data sources and the construction of the main variables. Section 4 presents the theoretical framework and derives testable predictions on the interaction between schooling costs, gender composition, and fertility decisions. Section 5 outlines the empirical strategy. Section 6 shows the main results and Section 7 presents the robustness checks. Section 8 concludes.

2 Market Structure and Expansion of the Private School Sector

This section characterizes the supply-side features of the private school sector in India, emphasizing market structure, pricing behavior, regulatory constraints, and competitive dynamics. We interpret the rapid expansion of private schooling not merely as a response to household demand, but as the outcome of a large-scale transformation in a decentralized service industry whose growth reshaped the economic environment facing households.

2.1 A Rapidly Expanding, Fragmented Service Market

Between 2000 and 2020, India experienced a dramatic expansion of private schooling with no close parallel in other developing countries. Private school enrollment in rural India increased from approximately 19 percent in 2006 to over 30 percent by 2014 (Kingdon, 2020), and by the time of NFHS-5 (2019–21), private schools accounted for more than 10 percent of all schools in the median district, with substantially higher shares in the upper tail of the distribution.

This expansion reflects the emergence of a large, highly fragmented market for low-cost education. The sector is dominated by small, locally operated schools rather than large chains or centralized providers (FSG, 2013; Tooley, 2009). As a result, competition is primarily local, with schools operating in geographically bounded markets and facing limited consumer mobility.

By the mid-2010s, the scale of the sector was substantial: an estimated 300,000 to 400,000 private schools served around 90 million children (Kingdon, 2020). In some urban areas, private provision dominates the market entirely, with private schools accounting for the majority of institutions.

2.2 Entry, Competition, and the Role of Public School Quality

The expansion of private schooling is closely linked to the performance of incumbent public providers. A large body of evidence documents severe inefficiencies in India's government school system, including high teacher absenteeism and weak learning outcomes. Kremer et al. (2005) report that 25 percent of teachers in government schools are absent on any given day, with substantial cross-state variation. Earlier evidence from the PROBE report highlights severe infrastructure deficits and low instructional quality (PROBE Team, 1999), while ASER data document persistent learning gaps among enrolled students.

In this environment, private schools emerged as a competitive response to unmet demand for quality. Observational evidence shows that private school students perform substantially better on standardized tests and attend more regularly (Muralidharan and Kremer, 2008). Experimental evidence further indicates that private schools can deliver comparable-to-better outcomes at a fraction of the cost (Muralidharan and Sundararaman, 2015).

From a market perspective, these patterns suggest that private school entry is responsive to local demand conditions and the quality of incumbent providers. Consistent with this interpretation, private schools are more likely to operate in areas with larger populations and weaker public school performance (Muralidharan and Kremer, 2008). More direct evidence on entry comes from Bagde et al. (2022), who show that the introduction of private schools induces sorting across sectors, with higher-achieving and more advantaged students disproportionately selecting into private providers.

2.3 Production Technology and Cost-Based Pricing

A defining feature of the private school sector is its low-cost production model. Private schools achieve cost advantages primarily through differences in labor contracts: teachers are locally recruited, often with lower formal qualifications, and are paid substantially less than government school teachers (Kingdon, 2020; Chudgar and Sakamoto, 2021). Salaries in private schools can be as low as one-fifth to one-tenth of public-sector wages.

This cost structure enables private schools to offer schooling at relatively low fees, supporting widespread entry even in low-income areas. At the same time, the sector is not homogeneous: there is substantial variation in both fees and quality across schools. Evidence from Andhra Pradesh shows that private secondary schooling significantly improves cognitive skills (Azam et al., 2016), consistent with a vertically differentiated market in which providers offer distinct quality–price bundles. The availability of low-cost labor also acts as a key supply-side constraint. Andrabi et al. (2013) show that in Pakistan the local availability of secondary-educated women plays an important role in determining where private schools can operate, as these women constitute the primary labor pool for private school teachers. This creates geographic variation in school supply that is driven by historical human capital accumulation rather than contemporaneous household demand.

2.4 Pricing Behavior, Geographic and Gender Differences

Cost differential between government and private schooling. The magnitude of the cost differential between government and private schooling is central to our argument. Data from the National Sample Survey (75th Round, 2017–18) reveal that average annual household expenditure per student at the primary level was Rs 1,253 in government schools but Rs 14,485 in private unaided schools—a ratio of nearly 12 to 1 (National Statistical Office, 2019). The gap narrows at higher levels but remains large: Rs 2,181 versus Rs 17,360 at the upper-primary level (8:1) and Rs 4,078 versus Rs 20,804 at secondary (5:1). These figures include tuition, examination fees, and all compulsory payments. Even acknowledging that government schools impose hidden costs (uniforms, textbooks, transport), the private-school premium is substantial. In Uttar Pradesh, Kingdon and Muzammil (2015) estimate that government per-pupil expenditure substantially exceeds the fees charged by low-cost private schools, implying that the private sector delivers education at a fraction of the public-sector cost per student—yet even these low fees represent a significant outlay for poor households. Nationally, the Central Square Foundation estimates that 70% of private-school parents pay less than Rs 1,000 per month, yet even this "low-cost" segment represents a large multiple of government-school expenditure (Central Square Foundation, 2020).

Geographical variation in private school fees. Private school fees exhibit substantial variation across states and between rural and urban areas (Kingdon, 2020), and even within the low-cost segment, schools operate at different points along the quality–price spectrum (FSG, 2013). This dispersion is consistent with a market in which providers differentiate along both price and quality dimensions. Moreover, emerging evidence suggests that pricing is responsive to policy and competitive conditions. Private schools adjust fees and quality in response to changes in effective demand, such as those induced by voucher programs (Sahai, 2023). Similarly, policy-induced changes in the regulatory environment can affect both entry and fee levels across districts (Vinod, 2024). Using unit-level data from the NSS 71st Round (2014), Kingdon (2020) documents that median monthly private-school fees range from Rs 117 in rural Uttar Pradesh to Rs 692 in rural Punjab (a sixfold difference) and from Rs 250 in urban Uttar Pradesh to Rs 1,800 in urban Delhi (a sevenfold difference). The rural–urban gap is equally large: average per-student expenditure across all school types was Rs 5,240 in rural areas versus Rs 16,308 in urban areas, a ratio of 3.1 to 1. At the state level, per-student education expenditure ranges from Rs 5,656 in Bihar to Rs 25,720 in Haryana—a nearly fivefold difference (National Statistical Office, 2025). This cross-state and rural–urban heterogeneity in the price of quality schooling is the source of the identifying variation in our empirical analysis: households in high-P_d districts face a steeper quality–cost gradient than those in low-P_d districts.

Economic burden of private schooling. Private-school fees represent a non-trivial burden even for middle-income households. Kingdon (2020) estimates that the mean private-school fee amounts to approximately 9% of state per capita income and roughly 10% of the annual minimum wage. For lower-income families with multiple children, education expenditure can absorb 20–30% of total household income (Endow, 2019). The "budget" private school segment—schools charging less than Rs 15,000 per year, serving an estimated 60–90 million children in 200,000–400,000 schools nationwide—illustrates the binding nature of this cost: even in the cheapest segment, shifting a child from government to private school requires a discrete increase in household expenditure that is proportional to P_d and amplified by the number of children (Tooley, 2009; Central Square Foundation, 2020). It is precisely this cost structure—a per-child schooling premium that cumulates with family size and interacts with gendered enrolment norms—that our theoretical model formalises in Section 4.

Gendered cost of private education. Gender gaps in educational attainment in India have been extensively documented (Kingdon, 2002), and while enrolment gaps have narrowed, the margin of gender differentiation has shifted toward the *type* of schooling children receive. The choice of schooling *medium*—English-medium private versus vernacular government—is itself an economically driven and gender-differentiated household decision: Munshi and Rosenzweig (2006) show that investment in English-language schooling responds sharply to the rising returns to English and is mediated by caste networks and gender within the household, reinforcing that the operative gender margin is which *type* of school a child attends rather than whether the child attends at all. Experimental evidence from developing countries confirms that parental educational investment decisions are actively managed and responsive to perceived returns: Dizon-Ross (2019) shows that parents in Malawi update their schooling investments when provided with information about their children's academic ability, and that these information frictions disproportionately affect poorer households. In the Indian context, this active allocation process operates along gender lines. While schools typically post uniform fees for all students, there is strong evidence that households allocate children across schooling options in a *gender-differentiated* manner: boys are systematically more likely than girls to be enrolled in fee-charging private schools, with the surplus allocated to free or low-cost government institutions. Using the 2004–05 and 2011–12 rounds of the India Human Development Survey (IHDS), Maitra et al. (2016) estimate that private-school enrollment among children aged 6–14 rose from 27% to 36% for boys but only from 23% to 29% for girls, widening the gender gap from 4.1 to 6.7 percentage points—an increase in relative terms from 18% to 23% of the female enrolment rate. Among children aged 7–14, the gap reached 7.7 percentage points by 2012, equivalent to a 27% premium in boys' enrolment relative to girls'. Using household fixed effects to compare siblings within the same family, they find that the within-household female disadvantage in private-school enrolment is roughly twice as large as the corresponding gap in overall school enrolment, confirming that the relevant margin is *which type* of school children attend, not whether they attend at all. Azam and Kingdon (2013), using the same survey, identify the key mechanism: households channel pro-male educational investment by sending sons to fee-charging private schools while enrolling daughters in free government schools, with the bias concentrated at ages 10–14 and in rural areas. Importantly, Datta and Kingdon (2019) show that while gender bias in school *enrolment* has declined substantially between 1995 and 2014 (more girls are now enrolled), gender bias in *conditional* educational expenditure—how much families spend per enrolled child—has significantly *risen* over the same period. This shift in the channel of gender bias, from whether children attend school to which type of school they attend, is precisely the mechanism that our model formalises: as private schooling expanded, the relevant margin of gender discrimination moved from enrolment to school choice, generating a widening gap in the effective per-child cost of education between sons and daughters. At the district level, Kingdon and Pal (2010) provide complementary evidence: using panel data across Indian districts, they show that private school growth is associated with higher literacy overall but also with a *wider* gender gap in enrolment, particularly among children aged 10–14—precisely the age range where private-school fees become most salient. India's National Sample Survey (75th Round, 2017–18) confirms the pattern with independent data: at the primary level, 30% of boys attended private institutions compared with 27% of girls (a 3 percentage-point gap, or 11% of the female rate); at the higher-secondary level, the gap widens to 4 percentage points (28% versus 24%, a 17% relative premium for boys) (National Statistical Office, 2019). The most recent administrative data from the Unified District Information System for Education (UDISE+, 2021–22) reports national enrolment shares of 28.4% for boys and 23.2% for girls across all levels of schooling, a gap of 5.2 percentage points or 22% of the female rate (Department of School Education and Literacy, 2022). More recently, using NFHS-5 (2019–21) microdata[3]—the same survey underlying our fertility analysis—Nandi et al. (2023) document that girls aged 2–4 are 9–17% less likely than boys to be enrolled in private preschools; in a mother fixed-effects specification comparing siblings within the same household, girls are 19% less likely to attend private preschool than their brothers, and the gap widens in poorer and less-educated households.

Taken together, these patterns establish the model's key premise directly—at school age in IHDS-based studies that use within-household sibling comparisons (Maitra et al., 2016; Azam and Kingdon, 2013), and at the preschool margin in our own first stage (Section 6.1)—and, combined with the large private–government cost differential documented above, it follows mechanically that sons carry a higher expected schooling cost than daughters. What no single survey observes is the link from a household's *realised* schooling costs to its subsequent fertility. Rather than estimate that link from endogenous enrollment choices, we identify it through district-level exposure to private schooling (P_d): an intent-to-treat measure of the cost environment that is more exogenous than a household's own private-enrolment decision.

2.5 Regulation and Market Constraints

The regulatory environment governing private schools introduces an additional layer of heterogeneity across markets. The Right to Education Act (RTE, 2009) imposed na-

[3]The NFHS-5 household questionnaire records the type of institution (government, private, or other) only for *preschool* attendance (variable SH18, children aged 3–5). For school-age children (aged 6 and above), the survey records whether the child attends school and at what level, but does not distinguish between government and private institutions. Nandi et al.'s analysis is therefore restricted to the preschool margin. Our fertility analysis uses the women's questionnaire (Individual Recode), which contains complete birth histories but no child-level schooling variables; the district-level private-school prevalence measure P_d from UDISE serves as our proxy for household exposure to private-school costs.

tionwide requirements on private schools, including infrastructure standards (separate toilets, drinking water, boundary walls, playground), mandatory seat reservations for disadvantaged students, minimum pupil–teacher ratios, and formal recognition requirements. Private schools were granted three years to comply, with non-compliant schools subject to derecognition and closure.

School closures and enforcement. The impact of RTE enforcement on the private school sector was substantial but geographically uneven. By March 2014, approximately 4,500 private unaided schools had closed and over 15,000 had received closure notices for failing to meet infrastructure norms (Kingdon, 2020). Crucially, enforcement did not end with the initial compliance deadline. Between 2014 and 2016, nearly 5,000 low-cost private schools were closed in Punjab, Haryana, Delhi, and Rajasthan alone, and Maharashtra witnessed the closure of approximately 7,000 schools over the same period (Centre for Civil Society, 2018). A 2017 amendment extended the deadline for teacher qualification requirements, sustaining compliance pressure through the end of the decade. The enforcement process was thus *rolling and staggered* across states—some states enforced aggressively (Punjab, Delhi, Maharashtra), while others took limited action—generating substantial between-state and within-state variation in the regulatory shock to private school supply during the 2015–2020 period that overlaps with our estimation sample.

Unrecognised schools. A significant fraction of the private school sector operates without formal government recognition. Using school-level data from the UDISE portal, we compute that across 720 districts with at least one private school, unrecognised schools account for 2.1% of all schools on average but 11.8% of private schools (Table 1). The distribution is highly skewed: the median district has almost no unrecognised schools, but 103 districts (14%) have unrecognised schools exceeding 5% of the total, and the maximum reaches 40%.

Table 1: Prevalence of Unrecognised Private Schools Across Districts

	Unrecognised / all schools	Unrecognised / private schools
Mean	0.021	0.118
Median	0.001	0.005
Standard deviation	0.045	0.225
Minimum	0.000	0.000
Maximum	0.402	1.000
Districts with >5%	103	223
Total districts	720	720

Notes: Computed from school-level UDISE data (circa 2019–20). "Unrecognised" includes schools classified as "Unrecognized" and "Madarsa unrecognized" in the UDISE management type field. "Private schools" includes Private Unaided (Recognised) plus all unrecognised schools. The sample comprises 720 districts with at least one private school.

Regulatory heterogeneity. Beyond RTE enforcement, substantial variation exists in how states regulate private schools, particularly with respect to fee controls. Tamil Nadu operates one of the most stringent regimes: a government-appointed Fee Determination Committee sets binding fee ceilings for all private schools in a region, regardless of the school's cost structure. Delhi imposes strict penalties for fee violations (up to Rs 10 lakh for repeat offences) with provisions for de-recognition or administrative takeover. At the other end, states like Rajasthan rely on school-level Parent–Teacher Association committees to negotiate fees, and several states have no binding fee regulation at all. Maharashtra caps fee increases at 15% biennially, while Gujarat's High Court has upheld statutory caps on private school fees. This patchwork of regulatory regimes—ranging from binding price controls to near-complete laissez-faire—implies that the extent and nature of private school expansion depend not only on demand conditions but also on institutional constraints that vary substantially across states (Oxfam India, 2020).

2.6 Implications for Market Equilibrium and Household Behavior

Taken together, the evidence characterizes the private school sector as a system of localized, vertically differentiated markets in which entry, pricing, and quality provision respond to both demand conditions and institutional constraints. The expansion of this sector represents a shift from a largely free public system to one in which quality is increasingly priced.

For households, this shift changes the economic environment in which decisions about child investment and family size are made. The availability of fee-based schooling increases the marginal cost of child quality for households that opt into private provision.

In the presence of liquidity constraints and gender-biased preferences, this generates differential investment across children within the household.

As we show below, these market features give rise to equilibrium spillovers beyond the education sector, affecting fertility decisions and parity progression in ways that depend on both household composition and local market conditions.

2.7 Implications for Identification

The institutional context suggests several supply-side determinants of private school availability that are plausibly independent of household fertility decisions:

1. *State-level regulatory variation.* Differences across states in recognition norms, fee regulation, and enforcement intensity are determined by historical political and administrative factors, not by the fertility preferences of individual households within those states. When we include district fixed effects in our empirical specifications, we absorb the level of these state-specific regulatory environments; the remaining variation comes from within-district changes over time and from cross-household differences in exposure.
2. *The local supply of potential teachers.* Following Andrabi et al. (2013), who document this mechanism in Pakistan, the pre-existing stock of secondary-educated women in a district constrains where private school entrepreneurs can operate. Furthermore, Muralidharan and Kremer (2008) show that the sector's reliance on locally recruited, low-cost teachers creates geographic variation in school supply driven by the local labor market rather than contemporaneous household demand. This stock is determined by historical educational investments—themselves a product of earlier policy decisions, missionary presence, and geographic factors—rather than by current household fertility choices.
3. *RTE-induced supply shocks.* The mechanical closure of non-compliant private schools following the 2009 RTE Act provides variation in private school availability that is driven by the interaction between a national policy and pre-existing school characteristics (building type, presence of a playground), not by local fertility demand. Districts that experienced more closures effectively received a negative shock to private school supply.

While none of these arguments constitutes a formal instrument—and we accordingly present our main results as a structured reduced-form framework with district fixed effects—the supply-side drivers of private school expansion provide the institutional basis for our identifying assumption: that within a state, the variation in P_d across districts and over time reflects differences in the local schooling market that are not driven by household fertility preferences. We now test this assumption directly.

2.7.1 Empirical Balance Test

If the change in private-school prevalence across districts were driven by the same household-level forces that determine fertility—income growth, shifting gender norms, rising educational aspirations—then baseline socioeconomic and demographic characteristics should predict the subsequent expansion. We test this by regressing the change in district-level private-school share between 2015–16 and 2019–20 (ΔP_d) on a comprehensive set of baseline district characteristics measured from NFHS-4 (2015–16), with and without state fixed effects. The sample comprises 525 districts matched across both UDISE waves and NFHS-4.

Table 2 reports the results. Each row in Panel A presents the coefficient from a separate univariate regression of ΔP_d on a single baseline characteristic, controlling for state fixed effects. Panel B reports the joint specification including all baseline variables simultaneously.

The results are reassuring for our identification strategy. The variables most directly relevant to our fertility analysis—children ever born, acceptance of wife-beating, and the autonomy index—are individually insignificant predictors of private-school expansion within states (Panel A, $p = 0.77$, $p = 0.46$, and $p = 0.31$, respectively). A joint F-test for gender norms (beating acceptance and autonomy together) fails to reject the null of zero predictive power ($F = 0.53$, $p = 0.59$). Fertility-related variables as a group—children ever born, sterilisation prevalence, and modern contraceptive use—are jointly significant at the 5 percent level ($F = 2.82$, $p = 0.04$), driven primarily by the sterilisation variable ($\hat{\beta} = -0.015$, $p = 0.05$); the substantive magnitude is small, however, implying that a 10 percentage-point increase in district sterilisation prevalence is associated with only a 0.15 percentage-point smaller increase in private-school share.

The variables that *do* predict private-school growth within states are structural and demographic rather than attitudinal or fertility-related. Urban districts experienced faster expansion ($\hat{\beta} = 0.016$, $p = 0.04$), consistent with the supply-side logic that private schools require population density and infrastructure. Caste composition is a strong predictor: districts with larger Scheduled Caste populations saw more growth ($\hat{\beta} = 0.028$, $p < 0.01$), while Scheduled Tribe districts saw significantly less ($\hat{\beta} = -0.019$, $p < 0.001$)—a pattern consistent with the geographic remoteness of tribal areas limiting school entry. Female employment is negatively associated with expansion ($\hat{\beta} = -0.025$, $p < 0.01$), possibly reflecting that areas with more women in agriculture or manual labour have less demand for fee-paying schools.

The multivariate specifications in Panel B reinforce the distinction between between-state and within-state variation. Without state fixed effects (column 3), several baseline characteristics are significant predictors of ΔP_d: sterilisation prevalence (-0.021, $p < 0.001$), the autonomy index (-0.007, $p = 0.05$), urbanisation ($+0.030$, $p < 0.01$), the

Muslim population share (-0.021, $p < 0.001$), and child mortality ($+0.048$, $p = 0.02$). However, this specification conflates within-state and between-state variation, and many of these variables are correlated with broad regional development patterns.

When state fixed effects are added (column 4), most of these associations disappear: norms, education, religion, and child mortality variables all become insignificant at the 5% level. Children ever born remains marginally significant ($+0.0098$, $p < 0.10$), though small in magnitude. Only urbanisation ($+0.041$, $p < 0.01$), caste composition (SC: $+0.030$, $p = 0.02$; ST: -0.018, $p < 0.001$), and the lagged level of P_d (-0.178, $p < 0.01$) survive. The baseline private-school share enters with a strong negative coefficient, indicating convergence: districts that started with lower private-school prevalence caught up faster, a pattern more consistent with market entry dynamics than with demand-driven endogeneity. The R-squared rises from 0.18 without state fixed effects to 0.48 with them, confirming that the great majority of the explained variation in ΔP_d is attributable to state-level factors and baseline school supply rather than to within-state differences in socioeconomic or demographic characteristics.

These results support the view that private-school expansion across Indian districts was driven primarily by supply-side factors—urbanisation, caste-linked geography, and the pre-existing schooling landscape—rather than by the fertility preferences, gender norms, or reproductive behaviour of households within those districts. For our main specifications, which include district fixed effects and identify the key parameters from interactions between district-level P_d and individual-level son composition, the relevant exogeneity condition is even weaker: it requires only that the *interactions* of P_d with individual characteristics (DD status, wealth) be uncorrelated with unobserved determinants of parity progression, not that the level of P_d itself be exogenous.

Table 2: Balance Test: Do Baseline District Characteristics Predict Private-School Expansion?

	Panel A: Univariate		Panel B: Multivariate	
	No state FE	State FE	No state FE	State FE
Fertility and reproductive behaviour				
Children ever born	0.0083^{***}	-0.0010	0.0110^{*}	0.0098^{*}
	(0.0021)	(0.0033)	(0.0059)	(0.0057)
Sterilised	-0.0135^{***}	-0.0153^{**}	-0.0213^{***}	-0.0085
	(0.0043)	(0.0077)	(0.0060)	(0.0225)
Joint F-test		$F = 2.82$, $p = 0.04$		
Gender norms				
Justifies wife-beating	-0.0032	0.0057	-0.0004	0.0061
	(0.0089)	(0.0077)	(0.0079)	(0.0082)
Autonomy index	-0.0099^{***}	-0.0033	-0.0065^{**}	-0.0015
	(0.0033)	(0.0032)	(0.0033)	(0.0033)
Joint F-test		$F = 0.53$, $p = 0.59$		
Socioeconomic characteristics				
Education (years)	-0.0010^{**}	0.0008	-0.0006	0.0004
	(0.0005)	(0.0007)	(0.0008)	(0.0009)
Urban	0.0054	0.0157^{**}	0.0301^{***}	0.0405^{***}
	(0.0061)	(0.0076)	(0.0109)	(0.0129)
Low wealth (Q1–Q2)	0.0051	-0.0084	-0.0171^{*}	-0.0041
	(0.0039)	(0.0056)	(0.0099)	(0.0088)
Muslim	-0.0074^{*}	0.0031	-0.0210^{***}	-0.0067
	(0.0038)	(0.0051)	(0.0049)	(0.0063)
Scheduled Caste	0.0380^{***}	0.0278^{***}	0.0256^{*}	0.0303^{**}
	(0.0110)	(0.0105)	(0.0149)	(0.0132)
Scheduled Tribe	-0.0134^{***}	-0.0186^{***}	-0.0115^{***}	-0.0179^{***}
	(0.0025)	(0.0037)	(0.0039)	(0.0050)
Joint F-test		$F = 5.47$, $p < 0.001$		
Other				
Currently working	-0.0271^{***}	-0.0247^{***}	-0.0094	-0.0118
	(0.0093)	(0.0076)	(0.0082)	(0.0083)
Any child died	0.0662^{***}	-0.0306	0.0480^{**}	-0.0443
	(0.0138)	(0.0239)	(0.0198)	(0.0319)
Baseline P_d (2015–16)	-0.0078	-0.0569	0.0164	-0.178^{***}
	(0.0314)	(0.0382)	(0.0347)	(0.059)
State FE	No	Yes	No	Yes
Joint F: all baseline vars			$F = 9.71^{***}$	$F = 3.98^{***}$
R^2			0.176	0.477
Districts	525–534	525–534	525	525

Notes: Dependent variable: change in district private-school share (ΔP_d) between 2015–16 and 2019–20 (UDISE data). Baseline characteristics are district-level weighted means from NFHS-4 (2015–16). Panel A reports coefficients from separate univariate regressions of ΔP_d on each baseline variable; Panel B reports a single multivariate regression including all variables simultaneously. Modern contraception is excluded from Panel B due to collinearity with sterilisation. Joint F-tests in Panel A are computed from the full multivariate specification with state FE. Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

As a complementary assessment, we split districts into those above and below the median within-district change in P_d (0.96 percentage points) and compare baseline covariates between the two groups. To evaluate the economic magnitude of the differences, we compute normalised differences following Imbens and Wooldridge (2009): $\text{ND} = (\bar{X}_1 - \bar{X}_0) \,/\, \sqrt{(S_1^2 + S_0^2)/2}$, where $\bar{X}_g$ and S_g^2 are the group-specific mean and variance. Normalised differences are independent of sample size and measure how far apart the two distributions are in units of the pooled standard deviation; Imbens and Wooldridge recommend $|\text{ND}| < 0.25$ as the criterion for acceptable balance.

Table 3 reports the results. All 20 covariates satisfy the $|\text{ND}| < 0.25$ threshold recommended by Imbens and Wooldridge. Gender norms (beating acceptance: -0.02; autonomy: -0.06), the variable most directly relevant to our outcome (wants no more: 0.01), and most demographic covariates show normalised differences well below 0.10.

The one variable that warrants explicit discussion is *children ever born*, which at $\text{ND} = 0.20$ is the largest difference in the table and approaches the 0.25 threshold. Districts with above-median ΔP_d had a baseline mean of 2.90 children compared with 2.59 for below-median districts. This positive association raises a potential mean-reversion concern: districts with higher initial fertility may have experienced faster baseline-dependent fertility declines that coincided with private-school expansion, generating a spurious negative correlation between ΔP_d and subsequent fertility change. We take this concern seriously. For the *main effect* of P_d—which is identified from within-district variation across two waves—this is a legitimate source of residual confounding that our district fixed effects do not fully resolve if the convergence dynamics are district-specific and time-varying.

However, the paper's strongest evidence comes not from the main P_d effect but from its *interactions* with individual-level sex composition (DD $\times$ P_d). Mean reversion is a district-level phenomenon that would affect all women in a district symmetrically; it has no reason to operate differently for women with two daughters versus two sons within the same district and survey wave. The interaction coefficients are therefore identified from cross-individual variation within district-wave cells and are immune to the mean-reversion threat. This is the primary reason we emphasise the interaction-based evidence throughout the paper.

Table 3: Covariate Balance: Normalised Differences by Median Split on ΔP_d

	Mean		
	Above median ΔP_d	Below median ΔP_d	Normalised diff.
Demographics			
Age	35.68	35.97	−0.041
Education (years)	5.23	5.83	−0.114
No education	0.412	0.351	0.125
Higher education	0.093	0.097	−0.013
Age at first marriage	18.23	18.66	−0.104
Urban	0.336	0.298	0.081
Fertility and reproduction			
Children ever born	2.90	2.59	0.200
Sterilised	0.397	0.461	−0.129
Modern contraception	0.500	0.556	−0.112
Wants no more	0.839	0.834	0.012
Any child died	0.168	0.124	0.124
Gender norms and empowerment			
Justifies wife-beating	0.454	0.464	−0.021
Autonomy index (0–3)	2.246	2.317	−0.064
Currently working	0.268	0.286	−0.042
Household characteristics			
Low wealth (Q1–Q2)	0.410	0.348	0.127
Muslim	0.123	0.130	−0.022
Scheduled Caste	0.207	0.193	0.036
Scheduled Tribe	0.070	0.124	−0.184
OBC	0.481	0.441	0.081
Husband: higher education	0.137	0.138	−0.004
$\|\text{ND}\| > 0.25$	0 of 20 variables		

Notes: Districts are split at the median within-district change in private-school share ($\Delta P_d = 0.96$ pp). Means are weighted using NFHS-4 individual sample weights for ever-married women aged 25–49. The normalised difference is $\text{ND} = (\bar{X}_1 - \bar{X}_0)/\sqrt{(S_1^2 + S_0^2)/2}$, where $g \in \{1, 0\}$ indexes the above- and below-median groups. Following Imbens and Wooldridge (2009), $|\text{ND}| < 0.25$ indicates acceptable balance. All 20 covariates satisfy this criterion.

3 Data and Descriptive Statistics

Our empirical analysis combines two datasets: individual-level fertility and demographic data from India's National Family Health Survey (NFHS), and district-level school-type data from the Unified District Information System for Education (UDISE). This section describes each source, the matching procedure, and the resulting analysis sample.

3.1 National Family Health Survey (NFHS)

The NFHS is India's implementation of the Demographic and Health Surveys (DHS) programme, conducted by the International Institute for Population Sciences (IIPS) under the aegis of the Ministry of Health and Family Welfare. We use two rounds: NFHS-4 (2015–16) and NFHS-5 (2019–21). Each round surveys a nationally representative sample of women aged 15–49 and collects complete birth histories, socioeconomic characteristics, and information on reproductive health and behaviour[4] [5]

NFHS-4 interviewed 699,686 women across 640 districts in all 29 states and 7 union territories of India. NFHS-5 interviewed 724,115 women across 707 districts (reflecting district boundary changes in the intervening period). Both rounds identify the respondent's district of residence, enabling spatial linkage with external district-level data.

From the individual recode (IR) files, we extract three categories of variables. *First*, complete birth histories, which record the sex and date of each birth in birth order ($b4_k$ = sex of the k-th birth). This enables us to reconstruct the sex composition of the first n children at each parity n—a critical feature of our specification, as we discuss in Section 5. *Second*, individual and household characteristics: the mother's age, years of education, and employment status; the husband's education level (none, primary, secondary, or higher); household wealth quintile; urban or rural residence; religion; and caste/tribe category (Scheduled Caste, Scheduled Tribe, Other Backward Class, or General). *Third*, attitudinal variables: a set of five questions on whether the respondent considers wife-beating justified under various circumstances, a decision-making autonomy index, and an indicator for media exposure (frequency of reading newspapers, listening to the radio, and watching television).

We restrict the analysis sample to ever-married women aged 25–49. The lower bound ensures that women have had sufficient time to initiate and partly complete their child-bearing. The upper bound is the standard DHS eligibility age.

Figure 1 documents the fertility decline that motivates this paper. Panel (a) plots mean children ever born by single-year age for seven five-year birth cohorts, pooling across NFHS waves. The cohort profiles shift downward markedly: women born in 1960–64 had accumulated approximately 4.5 children by age 40, while those born in 1980–84—observable up to their mid-thirties in the most recent survey—have accumulated approximately 2.7 children by age 35, on a trajectory toward completed fertility well

[4]The DHS programme also conducted an earlier round, NFHS-3 (2005–06), which we use for descriptive purposes (Figure 1) but exclude from the main analysis. Two data limitations prevent its inclusion. First, NFHS-3 does not identify districts: the survey records only the respondent's state and urban/rural residence, with no sub-state geographic identifier, making it impossible to merge with district-level school data. Second, the UDISE school census begins in 2012–13; no comparable district-level data on school management type exists for 2005–06.

[5]Earlier rounds of the District Information System for Education (DISE), maintained by NIEPA, cover 2005–06 through 2017–18 and are available through the School Report Cards portal (`schoolreportcards.in`), but require institutional registration and were not accessible for this study. In practice, this limitation is less consequential than it may appear. The private school expansion we study accelerated sharply after 2006: ASER data show that rural private enrollment rose from 18.7% in 2006 to 30.8% by 2014, meaning that the two waves we use (2015–16 and 2019–21) capture the period in which the cost-of-quality channel was most salient. The earlier period, when private schools served fewer than one in five rural children, predates much of the mechanism our model describes.

below 3.0. The youngest observable cohorts (1990–94), still in their late twenties at the time of the most recent survey, are already on a trajectory that will yield completed fertility well below replacement level.

Panel (b) shows age-specific fertility rates (ASFR) computed from three-year birth windows in each NFHS wave, alongside the implied Total Fertility Rate (TFR).[6] The TFR declined from 2.67 in NFHS-3 (2005–06) to 2.18 in NFHS-4 (2015–16) and 1.98 in NFHS-5 (2019–21)—a 26% decline over fifteen years, with India crossing the replacement threshold of 2.1 between the last two surveys. The decline is visible across the entire age distribution but is most pronounced at younger ages: the ASFR for women aged 15–19 fell by 65%, from 44 to 15 per thousand, reflecting delayed initiation of childbearing. Peak fertility at ages 20–24 declined from 195 to 138 per thousand, while rates at ages 30 and above declined more modestly. This compression of the fertility schedule—earlier stopping rather than later starting—is a distinctive feature of India's demographic transition, consistent with the sterilisation-based stopping mechanism documented by Park et al. (2023).

[6]The total fertility rate (TFR) is a synthetic measure that summarizes the average number of children a woman would have over her lifetime if she were to experience the age-specific fertility rates observed in a given period. It is constructed by summing fertility rates across all reproductive ages (typically 15–49), where each age-specific fertility rate reflects the number of births per woman in that age group during the reference period. Because it is based on current age-specific rates rather than actual completed fertility of cohorts, the TFR provides a snapshot of fertility conditions at a point in time rather than a prediction of lifetime fertility for any particular cohort.

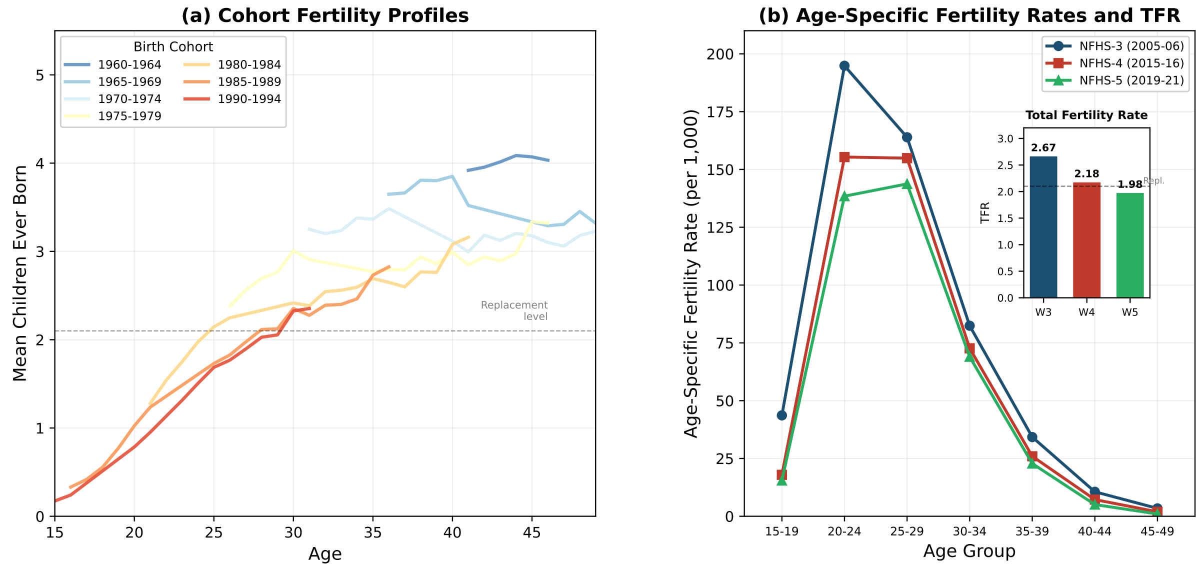

Figure 1: India's Fertility Decline: Cohort Profiles and Age-Specific Rates

Notes: Panel (a) plots mean children ever born by single-year age for ever-married women, separately by five-year birth cohort, pooling across NFHS-3 (2005–06), NFHS-4 (2015–16), and NFHS-5 (2019–21). Each cohort is observed at different ages across the three surveys; cells with fewer than 100 weighted observations are excluded. The dashed line indicates replacement-level completed fertility ($\approx$2.1). Panel (b) plots age-specific fertility rates (ASFR) per 1,000 women, computed from births in the 36 months preceding each survey. The inset shows the Total Fertility Rate (TFR $= 5\times\sum$ASFR) for each wave. All estimates use NFHS individual sample weights.

3.2 Unified District Information System for Education (UDISE)

The UDISE is a comprehensive school census maintained by the Ministry of Education, covering all recognised schools in India regardless of management type. We use data compiled from the UDISE Public Reports portal, which provides district-level school counts disaggregated by management category: Department of Education (government schools), private aided, private unaided (recognised), and other categories (local body, social welfare, etc.).

The compiled dataset covers 771 districts across 37 states and union territories for the academic years 2012–13 through 2019–20, yielding a balanced panel of 5,397 district-year observations (with 16 additional observations for 2012–13 covering a subset of districts). For each district-year, we compute the private-school share as:

$$P_d = \frac{\text{Private unaided schools} + \text{Private aided schools}}{\text{Total schools in district } d}.$$

A non-trivial segment of private provision operates without formal recognition: unrecognised private schools account for roughly 12% of private schools in the typical district (Table 1). Because such schools are recorded inconsistently in the district-level counts, our treatment variable P_d is built from *recognised* private unaided and aided schools and does not include them. To the extent that unrecognised but still fee-charging schools also raise households' exposure to the cost of quality, P_d understates that exposure, biasing our estimates toward zero.

We use P_d as a proxy for households' exposure to the cost of child quality rather than as a direct measure of school fees. The economic rationale is twofold. First, in a district with a higher private-school share, the probability that a household enrolls a given child in a fee-charging school is higher, either through direct choice or through social pressure and peer effects; P_d therefore captures the *expected* schooling cost facing a forward-looking household deciding on an additional child. Second, districts with more private schools are districts where the market for quality education is thicker and more visible: even families that do not (yet) enroll their children privately are more exposed to the cost of quality through neighbours, relatives, and community norms. The model in Section 4 formalises this by parameterising per-child schooling costs as linear in P_d ($c_b = \alpha_b P_d$, $c_g = \alpha_g P_d$), which should be understood as a reduced-form approximation: P_d indexes a household's expected exposure to the cost of quality, not the literal fee charged by a specific school.

Table 5 reports the distribution of this measure. In the 2015–16 cross-section (matched to NFHS-4), the mean private-school share across 661 districts with at least one private school is 9.1%, with substantial variation: the 10th percentile is 2.0% while the 90th percentile reaches 17.8%. By 2019–20 (matched to NFHS-5), the mean has risen to 10.5%, reflecting continued private-school expansion over the period.

3.3 Matching Procedure

We match NFHS individual records to UDISE district-level private-school shares using state and district names as the merge key. Both datasets identify districts by name (NFHS through value-labelled numeric codes, UDISE through string identifiers). We standardise all names to uppercase and harmonise minor spelling differences (e.g., replacing "and" with "&" to match NFHS conventions).

For each NFHS wave, we match to the contemporaneous UDISE cross-section: NFHS-4 (2015–16) to UDISE 2015–16, and NFHS-5 (2019–21) to UDISE 2019–20. The match rate is 82% of NFHS-4 districts (526 of 640) and 84% of NFHS-5 districts (593 of 707). Unmatched districts are primarily those created through boundary changes between the NFHS and UDISE reference periods, districts with name changes, and a small number of union territories with idiosyncratic naming conventions.

At the individual level, the matched sample comprises 747,739 ever-married women aged 25–49: 362,174 from NFHS-4 and 385,565 from NFHS-5. This represents approximately 83% of the full NFHS sample of eligible women, providing adequate power for district-level analyses.

3.4 Descriptive Statistics

Table 4 presents summary statistics for the analysis sample, separately by survey wave.

Demographics and fertility. The mean age of women in the sample is 36.0 years. Average years of education increased from 5.5 in NFHS-4 to 6.3 in NFHS-5, reflecting the continued educational expansion documented for India over this period. The mean age at first marriage is 18.6 years, barely above the legal minimum. The average woman has 2.7 children ever born, declining from 2.8 in NFHS-4 to 2.6 in NFHS-5. The proportion who had a birth in the three years preceding the survey fell from 17.5% to 16.0%. Approximately 84% of women want no more children or are already sterilised, and female sterilisation is the dominant contraceptive method: 42% in NFHS-4 rising to 46% in NFHS-5.

Household and social characteristics. About 32% of the sample resides in urban areas. The bottom two wealth quintiles account for 38% of the sample. The religious composition is 82% Hindu and 13% Muslim. By caste, 21% belong to Scheduled Castes, 9% to Scheduled Tribes, and 46% to Other Backward Classes. On husband's education, 20% have no formal education, 15% primary, 50% secondary, and 14% higher education.

Gender norms and empowerment. The proportion of women justifying wife-beating for at least one of five stated reasons declined from 45.7% in NFHS-4 to 39.8% in NFHS-5—a substantial shift in attitudes over a short period. Women's decision-making autonomy (an index counting participation in up to three household decisions) averages 2.4 out of 3, increasing from 2.3 to 2.5 between waves.

Sex composition at parity 2. Among women with at least two children, we classify the sex composition of the first two births using birth-history records. The distribution is: 16.4% had two daughters first (DD), 54.6% had one of each (Mixed), and 29.0% had two sons first (SS). Under India's sex ratio at birth of approximately 108 boys per 100 girls ($p_{\text{boy}} \approx 0.519$), the expected distribution among families reaching parity two would be roughly 23% DD, 50% Mixed, and 27% SS. The observed 29% SS is close to this benchmark, but the observed 16.4% DD represents a substantial deficit—approximately 30% below the expected share. Since families with a firstborn daughter are, if anything, *more* likely to progress to a second birth (owing to son preference), the DD share among parity-two women should be pushed *above* the unconditional benchmark, making the deficit even more striking. The most plausible explanation is sex-selective abortion following prenatal sex determination, which is well documented in India and disproportionately targets second-born daughters after a firstborn daughter (Jayachandran, 2017, 2023). This implies that the DD cell in our analysis is a selected sample: women who arrive at parity two with two daughters despite the availability of sex selection. This selection could attenuate or amplify the DD $\times$ P_d interaction depending on the correlation between the propensity for sex selection and the responsiveness to schooling costs, a caveat that should be borne in mind when interpreting the interaction coefficients.

Parity progression and son preference. The raw parity-progression ratios reveal strong son preference. Among women aged 25–49, 86% of those at parity 1 progressed to parity 2 (near-universal), 56% at parity 2 progressed to parity 3, and 50% at parity 3 progressed to parity 4. Critically, the parity-2-to-3 progression rate varies sharply by sex composition of the first two births: 63.5% among DD families, 56.2% among Mixed families, and 52.9% among SS families. The 10.6 percentage-point gap between DD and SS families is direct evidence of the son-preference stopping rule—households with only daughters are substantially more likely to continue childbearing.

Private-school prevalence. Panel I of Table 4 reports the distribution of the district-level private-school share P_d, which is sourced from the UDISE school census and merged onto each NFHS woman based on her district of residence. In the analysis sample, the woman-weighted mean P_d is 11.1%, with the 10th–90th percentile range spanning from 3.0% to 21.4%. The share increased between waves: from a mean of 10.3% in NFHS-4 districts to 11.9% in NFHS-5 districts, with a median of 9.3% and 10.7%, respectively. These woman-weighted values are close to the unweighted district-level means reported in Table 5 (9.1% and 10.4%), indicating that the NFHS matched sample is broadly representative of the national distribution of private-school prevalence and that the weighting does not substantially distort the distribution of the treatment variable.

Figure 2 maps the change in private-school share across Indian districts between 2013–14 and 2019–20. The mean increase was 2.6 percentage points over this seven-year period, but with pronounced geographic heterogeneity. The largest expansions (dark red shading) are concentrated in the northern Hindi belt—Uttar Pradesh, Madhya Pradesh, Rajasthan, and Bihar—where baseline private-school presence was lower and government school quality gaps wider. Parts of central India (Chhattisgarh, Jharkhand) and the western states (Gujarat, Maharashtra) also experienced substantial growth. By contrast, several northeastern states and portions of southern India saw more modest increases or, in a few districts, slight declines. This spatial variation is essential to the empirical strategy: it ensures that the private-school treatment is not simply a proxy for north–south cultural divides or uniform national trends, but reflects genuinely differential exposure across districts within the same state.

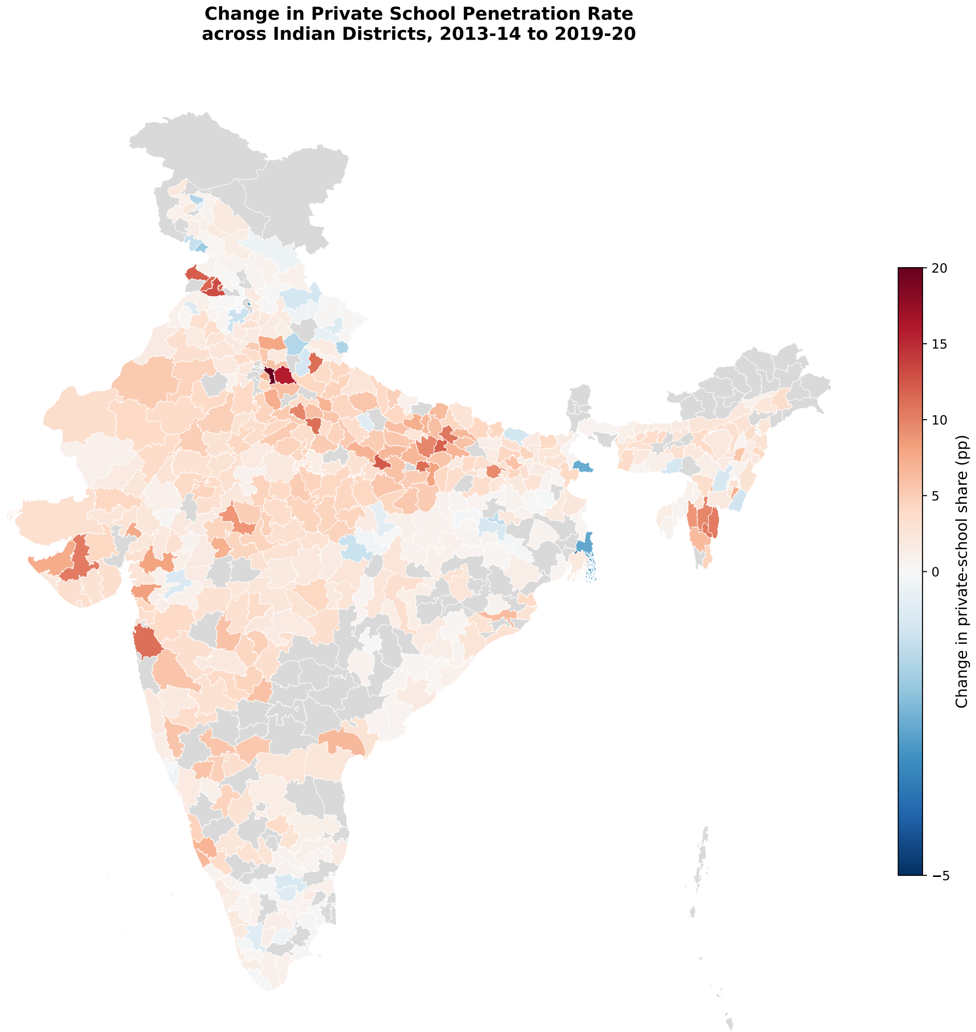

Figure 2: Change in Private School Penetration Rate across Indian Districts, 2013–14 to 2019–20

Notes: Each district is shaded according to the change in its private-school share (percentage points) between 2013–14 and 2019–20, computed from the UDISE school census. The colour scale spans the full data range (-5 to $+20$ pp), with a diverging palette centred at zero: red indicates an increase, blue a decrease, and white near-zero change. Grey districts could not be matched between the Census 2011 boundary file and the UDISE data. The mean change across matched districts is $+2.6$ percentage points (SD $= 2.7$; range: -4.7 to $+20.0$).

Table 4: Descriptive Statistics: Analysis Sample

	NFHS-4 (2015–16)		NFHS-5 (2019–21)	
	Mean	SD	Mean	SD
Panel A: Demographics				
Age (years)	35.8	7.1	36.2	7.2
Education (years)	5.5	5.3	6.3	5.4
Age at first marriage	18.4	4.2	18.7	4.1
Urban residence	0.321	0.467	0.313	0.464
Low wealth (Q1–Q2)	0.386	0.487	0.383	0.486
Panel B: Fertility				
Children ever born	2.77	1.62	2.60	1.47
Living children	2.56	1.43	2.44	1.32
Birth in last 3 years	0.175	0.380	0.160	0.367
Wants no more / sterilised	0.837	0.369	0.837	0.369
Female sterilisation	0.424	0.494	0.455	0.498
Modern contraception	0.524	0.499	0.602	0.489
Panel C: Mother's education				
No education	0.387	0.487	0.321	0.467
Primary	0.147	0.354	0.148	0.355
Secondary	0.372	0.483	0.408	0.492
Higher	0.095	0.293	0.122	0.328
Panel D: Husband's education				
No education	0.211	0.408	0.194	0.396
Primary	0.152	0.359	0.149	0.356
Secondary	0.498	0.500	0.504	0.500
Higher	0.137	0.344	0.150	0.357
Panel E: Religion and caste				
Hindu	0.813	0.390	0.821	0.383
Muslim	0.127	0.333	0.123	0.328
Scheduled Caste	0.203	0.402	0.219	0.414
Scheduled Tribe	0.092	0.289	0.094	0.292
Other Backward Class	0.462	0.499	0.462	0.499
Panel F: Gender norms and empowerment				
Justifies wife-beating (any reason)	0.457	0.498	0.398	0.489
Autonomy index (0–3)	2.28	1.07	2.47	0.94
Any media exposure	0.771	0.420	0.747	0.435
Currently working	0.276	0.447	0.306	0.461
Any child died	0.150	0.357	0.124	0.330
Panel G: Sex composition at parity 2				
Two daughters (DD)	0.166	0.372	0.162	0.369
One son, one daughter (Mixed)	0.542	0.498	0.550	0.498
Two sons (SS)	0.292	0.455	0.288	0.453
Panel H: Parity progression rates				
$P(2 \mid 1)$	0.865	0.342	0.857	0.350
$P(3 \mid 2)$	0.589	0.492	0.541	0.498
if DD	0.657	0.475	0.614	0.487
if Mixed	0.586	0.493	0.539	0.498
if SS	0.555	0.497	0.504	0.500
$P(4 \mid 3)$	0.521	0.500	0.478	0.500
Panel I: District private-school share (P_d, from UDISE)				
Mean	0.103	0.068	0.119	0.071
Median	0.093		0.107	
10th–90th percentile	[0.023, 0.195]		[0.036, 0.220]	
Women (N)	362,174		385,565	
Districts	526		593	

Notes: Sample restricted to ever-married women aged 25–49 in districts matched to UDISE school data. All means are weighted using NFHS individual sample weights. Education levels are categorical (0 = none, 1 = primary, 2 = secondary, 3 = higher). Autonomy index counts the number of household decisions (out of three: own health care, large purchases, family visits) in which the woman participates alone or jointly. Sex composition at parity 2 is based on the birth-order sex of the first two children ($b4_1$ and $b4_2$ from the birth history), not current living sons. Parity progression rates are computed as $P(n+1 \mid n) = \Pr(\text{CEB} > n \mid \text{CEB} \geq n)$ for women aged 25–49. Private-school share (P_d) is the ratio of private schools (aided and unaided) to total schools in the district, from the UDISE school census matched to the corresponding NFHS wave year.

Table 5: UDISE School Data: Summary Statistics

	2015–16				2019–20			
	Mean	SD	P10	P90	Mean	SD	P10	P90
Total schools	4,558	2,925	1,156	8,494	4,147	2,825	912	8,010
Government schools	1,237	1,123	31	2,603	1,080	986	41	2,265
Private schools	439	502	42	1,122	464	527	47	1,159
Private-school share	0.091	0.061	0.020	0.178	0.105	0.067	0.031	0.201
Districts	665 (661 with $P_d > 0$)				727 (all with $P_d > 0$)			

Notes: Data from the Unified District Information System for Education (UDISE), compiled from the UDISE Public Reports portal. Statistics are computed across districts, excluding state-aggregate rows and a small number of districts with incomplete coverage (fewer than 100 recorded schools). "Private schools" are private unaided (recognised) schools; the panel records no private aided schools. *Government and private schools do not sum to total schools*: the UDISE total also includes a large residual of other management types—other publicly managed schools (local-body and tribal/social-welfare departments), central-government schools, and madrasa or unrecognised schools—so that government plus private fall well short of the total. The private-school share P_d is private schools divided by total schools, reported for districts with at least one private school (661 of 665 in 2015–16; all 727 in 2019–20).

4 A Dynamic Fertility Model with Gendered Schooling Costs

4.1 Setup

Consider a household i in district d at parity n—that is, with n children already born, of whom s are sons and $n - s$ are daughters. The household sequentially decides whether to have an additional child. We model this parity-progression decision by specifying the *total* value and the *total* cost of children as functions of family size and composition, and then deriving the marginal benefit and marginal cost of one more child. The key insight is that marginal costs depend not only on the expected expenses of the next child, but also on how the *existing* cost burden—which is shaped by the gender composition of prior children—tightens the household's budget constraint.

4.2 Total Value of Children

The household derives utility from children according to:

$$V(n, s) = b\,n - \frac{\phi}{2}\,n^2 + \theta\,\mathbf{1}(s \geq 1), \tag{1}$$

where $b > 0$ is the baseline value per child, $\phi > 0$ generates diminishing marginal returns to fertility (concavity in n), and $\theta > 0$ is a one-time premium from having at least one son. The indicator $\mathbf{1}(s \geq 1)$ captures son preference in its simplest form: households value having *at least* one male heir, but additional sons beyond the first confer no further premium.

4.3 Total Cost of Children

The total cost of raising n children with sex composition $(s, n - s)$ in a district with private-school prevalence P_d is:

$$C(n, s, P_d) = c_0\, n + \big[s\, \alpha_b + (n - s)\, \alpha_g\big]\, P_d, \tag{2}$$

where $c_0 > 0$ is the baseline per-child rearing cost (food, clothing, healthcare) and the second term captures education expenditures. The parameters α_b and α_g represent the per-child education cost for boys and girls, respectively, scaled by the local prevalence of private schooling P_d. When $P_d = 0$ (no private schools), there are no additional education costs beyond c_0. When $P_d > 0$, sons generate higher schooling costs than daughters:

$$\alpha_b > \alpha_g \geq 0. \tag{3}$$

This ordering reflects the well-documented son-biased investment norm in India: boys are disproportionately enrolled in fee-charging private schools, while girls more often attend free government schools (see Kaul, 2018; Nandi et al., 2023).

4.4 The Household Budget Constraint

The household has total resources Y_i and allocates them between consumption x_i and child-rearing costs:

$$x_i = Y_i - C(n, s, P_d). \tag{4}$$

Total household utility combines the value of children with the utility from consumption:

$$U_i = V(n, s) + u(x_i), \tag{5}$$

where $u(\cdot)$ is strictly increasing and strictly concave ($u' > 0$, $u'' < 0$). The concavity of u is central to our mechanism: it implies that the *effective* cost of an additional child—measured in utility rather than rupees—depends on how stretched the household budget already is.

4.5 Expected Marginal Benefit of an Additional Child

At parity n with s sons, the household considers having one more child. The next child is a son with probability $p = 1/2$ and a daughter with probability $1/2$. The expected increase in the value of children from the $(n+1)$th birth is:

$$\begin{aligned}\mathbb{E}[\Delta V] &= \tfrac{1}{2}\big[V(n{+}1,\, s{+}1) - V(n,s)\big] + \tfrac{1}{2}\big[V(n{+}1,\, s) - V(n,s)\big] \\ &= b - \phi\big(n + \tfrac{1}{2}\big) + \tfrac{1}{2}\,\theta\,\mathbf{1}(s = 0). \end{aligned} \tag{6}$$

The first two terms are standard: the baseline value b net of the diminishing-returns penalty $\phi(n + 1/2)$, which increases with parity. The last term is the *son-preference bonus*: families with no sons receive an expected $\theta/2$ because there is a 50% chance the next child delivers the first son. Families that already have at least one son ($s \geq 1$) receive no additional premium.

4.6 Expected Direct Cost of an Additional Child

The expected increase in child-rearing costs from one more child is:

$$\mathbb{E}[\Delta C^{\text{direct}}] = c_0 + \tfrac{1}{2}(\alpha_b + \alpha_g)\, P_d. \tag{7}$$

This direct marginal cost depends on P_d (through expected schooling costs of the new child) but is *independent of the sex composition of existing children.* A family with two daughters and a family with two sons face identical expected costs for the *next* child. If costs were fully separable and the household faced no budget constraint, sex composition would affect fertility only through the benefit side (son preference), and no interaction between gender composition and P_d would arise.

4.7 The Shadow-Price Channel

The household's fertility decision, however, is not made in a vacuum: it is mediated by the budget constraint (4). Because $u(\cdot)$ is concave, the marginal utility of consumption—the *shadow price* of a rupee—is:

$$\lambda_i \equiv u'\big(Y_i - C(n, s, P_d)\big), \tag{8}$$

which is *increasing* in total costs $C(n, s, P_d)$ (since $u'' < 0$). A household that already devotes a large share of its income to children's schooling has lower residual consumption and therefore values each additional rupee more highly. For such a household, the direct cost $\mathbb{E}[\Delta C^{\text{direct}}]$ translates into a *larger* utility loss.

The effective marginal cost, in utility terms, is therefore:

$$\text{MC}_i = \lambda_i \cdot \mathbb{E}[\Delta C^{\text{direct}}] = u'\big(Y_i - C(n, s, P_d)\big) \cdot \big[c_0 + \tfrac{1}{2}(\alpha_b + \alpha_g)\, P_d\big]. \tag{9}$$

The household has another child if the expected marginal benefit exceeds the effective marginal cost:

$$\mathbb{E}[\Delta V] > \text{MC}_i + \varepsilon_i, \tag{10}$$

where ε_i is an idiosyncratic preference shock.

4.8 How Sex Composition Affects Effective Marginal Cost

The shadow price λ_i depends on the total cost $C(n, s, P_d)$, which in turn depends on the sex composition s of existing children. At a given parity n:

- A family with only daughters ($s = 0$) incurs total cost $C = c_0 n + n\,\alpha_g\,P_d$: all children attend cheaper government schools.
- A family with $s \geq 1$ sons incurs $C = c_0 n + [s\,\alpha_b + (n{-}s)\,\alpha_g]\,P_d > c_0 n + n\,\alpha_g\,P_d$: sons in private school raise the total bill.

Because $\alpha_b > \alpha_g$, families with sons have higher total costs, lower residual consumption, and therefore a *higher* shadow price λ_i. The same direct cost $\mathbb{E}[\Delta C^{\text{direct}}]$ translates into a higher effective marginal cost for families with sons than for families with only daughters.

This is the key mechanism: the gender composition of existing children affects fertility not through the expected cost of the *next* child (which is identical for all families at a given parity), but through the *budget-stretch* channel—the total cost burden of existing children raises the shadow price of the budget constraint, making any additional expenditure more painful in utility terms.

Our model treats children's existing school placements as a *predetermined commitment*: when the household contemplates an additional birth, the schooling of children already enrolled—sons in private, daughters in government—enters the budget as a fixed obligation rather than a margin that could be freely re-optimized. In principle a family could free up resources by moving an enrolled son to a cheaper school rather than by forgoing a birth; we assume only that doing so is *costly*. For any positive switching cost—whether transactions and social frictions or the sacrifice of an educational investment the family deliberately chose—a rise in P_d tightens a son-family's budget by more than it can cheaply undo, preserving the qualitative prediction that families with sons are more deterred than families with only daughters. Such frictions are substantial in the Indian context, where withdrawing a child from private schooling is uncommon and treated as a costly, status-laden reversal.

4.9 Linearisation and the Continuation Condition at Parity Two

To map the model onto an estimable reduced form, we linearise the shadow price and evaluate the continuation condition at parity $n = 2$, which is the empirically relevant margin for the transition from two to three children. Let $\bar{x}_i = Y_i - c_0 n$ denote residual income net of baseline child costs, and let $S(s, P_d) = [s\,\alpha_b + (n-s)\,\alpha_g]\,P_d$ denote the schooling-cost component of total child-rearing expenditure. A first-order Taylor expansion of u' gives:

$$\lambda_i = u'(\bar{x}_i - S) \approx u'(\bar{x}_i)\left[1 + \rho_i\,\frac{S(s, P_d)}{\bar{x}_i}\right], \tag{11}$$

where $\rho_i = -u''(\bar{x}_i)\,\bar{x}_i/u'(\bar{x}_i) > 0$ is the coefficient of relative risk aversion. Defining $\ell_i \equiv \rho_i/\bar{x}_i$, which is larger for poorer households (lower $\bar{x}_i$), and normalising $u'(\bar{x}_i) = 1$, the effective marginal cost (9) becomes:

$$\mathrm{MC}_i \approx \big[1 + \ell_i \cdot S(s, P_d)\big] \cdot \big[c_0 + \tfrac{1}{2}(\alpha_b + \alpha_g)\,P_d\big]. \tag{12}$$

4.10 Explicit Derivation of the Reduced Form

At parity $n = 2$, the continuation condition $\mathbb{E}[\Delta V] > \mathrm{MC}_i + \varepsilon_i$ can be written as the latent index $Y_i^* > 0$, where:

$$Y_i^* = \underbrace{\big[b - \tfrac{5}{2}\phi + \tfrac{\theta}{2}\,\mathrm{DD}_i\big]}_{\mathbb{E}[\Delta V]\text{ from (6)}} - \underbrace{\big[1 + \ell_i \cdot (2\alpha_g + s_i\,\Delta\alpha)\,P_d\big] \cdot \big[c_0 + \tfrac{1}{2}(\alpha_b + \alpha_g)\,P_d\big]}_{\mathrm{MC}_i\text{ from (12)}} + \varepsilon_i, \tag{13}$$

where $\mathrm{DD}_i \equiv \mathbf{1}(s_i = 0)$ indicates that both existing children are daughters, $\Delta\alpha \equiv \alpha_b - \alpha_g > 0$ is the gender gap in per-child schooling costs, and we have substituted $S(s, P_d) = (2\alpha_g + s_i\Delta\alpha)\,P_d$ at parity two.

Step 1: Expand the MC product. Multiplying through:

$$\begin{aligned}\mathrm{MC}_i &= c_0 + \tfrac{1}{2}(\alpha_b + \alpha_g)\,P_d + (2\alpha_g + s_i\Delta\alpha)\,c_0\,\ell_i\,P_d \\ &\quad + \underbrace{(2\alpha_g + s_i\Delta\alpha)\,\tfrac{1}{2}(\alpha_b + \alpha_g)\,\ell_i\,P_d^2}_{\text{second-order in } P_d\text{; dropped below}}.\end{aligned} \tag{14}$$

Dropping the P_d^2 term (a first-order approximation in P_d) and substituting into (13):

$$Y_i^* \approx \big(b - \tfrac{5}{2}\phi - c_0\big) + \frac{\theta}{2}\,\mathrm{DD}_i - \tfrac{1}{2}(\alpha_b + \alpha_g)\,P_d - (2\alpha_g + s_i\,\Delta\alpha)\,c_0\,\ell_i\,P_d + \varepsilon_i. \tag{15}$$

Step 2: Link s_i to DD. Since $\mathrm{DD}_i = 1$ implies $s_i = 0$ and $\mathrm{DD}_i = 0$ implies $s_i = \tilde{s} \geq 1$ (where $\tilde{s}$ is the number of sons among non-DD families), we can write $s_i\,\Delta\alpha = (1 -$

$\mathrm{DD}_i)\,\tilde{s}\,\Delta\alpha$. Substituting and rearranging:

$$
\begin{aligned}
Y_i^* \approx \big(b - \tfrac{5}{2}\phi - c_0\big) + \frac{\theta}{2}\,\mathrm{DD}_i - \tfrac{1}{2}(\alpha_b + \alpha_g)\,P_d \\
- 2\alpha_g\,c_0\,\ell_i\,P_d - \tilde{s}\,\Delta\alpha\,c_0\,\ell_i\,P_d + \mathrm{DD}_i \cdot \tilde{s}\,\Delta\alpha\,c_0\,\ell_i\,P_d + \varepsilon_i.
\end{aligned} \tag{16}
$$

Step 3: Decompose ℓ_i into mean and deviation. Write $\ell_i = \bar{\ell} + (\ell_i - \bar{\ell})$, where $\bar{\ell}$ is the population average and $(\ell_i - \bar{\ell})$ is the household-specific deviation—positive for poorer households. Each term involving $\ell_i \cdot P_d$ splits into a component at mean liquidity (absorbed into the P_d coefficient) and a component that varies with household resources. In the empirical specification, we proxy the continuous deviation $(\ell_i - \bar{\ell})$ with the binary indicator LowW_i (below within-district median wealth). Because this replaces a continuous variable (with units of inverse income) by a discrete indicator, the structural expressions for π_4 and π_5 below should be understood as determining the *signs* and *relative ordering* of the reduced-form coefficients rather than their exact magnitudes. Collecting terms:

$$
\begin{aligned}
Y_i^* = \pi_0 + \pi_1\,\mathrm{DD}_i + \pi_2\,P_d + \pi_3\,\mathrm{DD}_i \times P_d \\
+ \pi_4\,P_d \times \mathrm{LowW}_i + \pi_5\,\mathrm{DD}_i \times P_d \times \mathrm{LowW}_i + \varepsilon_i,
\end{aligned} \tag{17}
$$

with the observed outcome $Y_i = \mathbf{1}(Y_i^* > 0)$ indicating parity progression.

4.11 Mapping Coefficients to Structural Parameters

Matching (16) to (17) term by term, each reduced-form coefficient has an explicit structural expression:

$$\pi_0 = b - \tfrac{5}{2}\phi - c_0 \qquad \text{(net baseline value at parity 2)} \tag{18}$$

$$\pi_1 = \frac{\theta}{2} \qquad > 0 \quad \text{(son-preference premium)} \tag{19}$$

$$\pi_2 = -\tfrac{1}{2}(\alpha_b + \alpha_g) - (2\alpha_g + \tilde{s}\,\Delta\alpha)\,c_0\,\bar{\ell} \qquad < 0 \quad \text{(cost of } P_d \text{ at mean liquidity)} \tag{20}$$

$$\pi_3 = \tilde{s}\,\Delta\alpha\,c_0\,\bar{\ell} \qquad > 0 \quad \text{(DD families less deterred by } P_d\text{)} \tag{21}$$

$$\pi_4 = -(2\alpha_g + \tilde{s}\,\Delta\alpha)\,c_0 \qquad < 0 \quad \text{(liquidity amplifies } P_d \text{ cost)} \tag{22}$$

$$\pi_5 = \tilde{s}\,\Delta\alpha\,c_0 \qquad > 0 \quad \text{(DD cost advantage larger for poor)} \tag{23}$$

4.12 Structural Interpretation of Each Prediction

Model Prediction 1 ($\pi_1 > 0$: Son preference)**.** *From* (19)*:* $\pi_1 = \theta/2 > 0$*. The son-preference premium* θ *enters the marginal benefit* $\mathbb{E}[\Delta V]$ *with weight* $1/2$ *(the probability the next child is a son). Families with only daughters have a higher expected marginal benefit of continuing and are therefore more likely to progress to parity three.*

Model Prediction 2 ($\pi_2 < 0$: Private-school cost effect)**.** *From* (20)*:* $\pi_2 = -\frac{1}{2}(\alpha_b + \alpha_g) - (2\alpha_g + \tilde{s}\,\Delta\alpha)\,c_0\,\bar{\ell} < 0$*. This coefficient has two components. The first,* $-\frac{1}{2}(\alpha_b + \alpha_g)$*, is the* direct *expected schooling cost of the next child. The second,* $-(2\alpha_g + \tilde{s}\,\Delta\alpha)\,c_0\,\bar{\ell}$*, is the* shadow-price *amplification: higher* P_d *raises the total cost burden of existing children, which tightens the budget and makes the baseline rearing cost* c_0 *of the next child more painful. Both components are negative, so* P_d *unambiguously reduces fertility.*

Model Prediction 3 ($\pi_3 > 0$: DD $\times$ P_d interaction)**.** *From* (21)*:* $\pi_3 = \tilde{s}\,\Delta\alpha\,c_0\,\bar{\ell} > 0$*. This interaction arises entirely from the shadow-price channel. At parity two, a DD family incurs existing schooling costs of* $2\alpha_g\,P_d$ *(both daughters in government school), while a family with* $\tilde{s}$ *sons incurs* $(2\alpha_g + \tilde{s}\,\Delta\alpha)\,P_d$*. The difference,* $\tilde{s}\,\Delta\alpha\,P_d$*, means that non-DD families have a higher shadow price* λ_i*, which amplifies the baseline cost* c_0 *of the next child by an additional* $\tilde{s}\,\Delta\alpha\,c_0\,\bar{\ell}\,P_d$*. DD families are spared this amplification and are therefore* less deterred *by* P_d*.*

Model Prediction 4 ($\pi_4 < 0$: $P_d \times$ LowW interaction)**.** *From* (22)*:* $\pi_4 = -(2\alpha_g + \tilde{s}\,\Delta\alpha)\,c_0 < 0$*. Poorer households have a larger* ℓ_i *(higher marginal utility of consumption), so the shadow-price amplification of* P_d *costs is stronger. For a household below the median wealth, the existing schooling burden* $(2\alpha_g + \tilde{s}\,\Delta\alpha)\,P_d$ *amplifies the next child's baseline cost* c_0 *by more than it does for a wealthy household. Private-school prevalence therefore reduces fertility* more *for poor families.*

Model Prediction 5 ($\pi_5 \geq 0$: Triple interaction DD $\times$ $P_d \times$ LowW)**.** *From* (23)*:* $\pi_5 = \tilde{s}\,\Delta\alpha\,c_0 > 0$*. Among poor households, the DD cost advantage—lower existing schooling expenditure—translates into a larger shadow-price differential relative to families with sons, because the curvature of* $u(\cdot)$ *is more pronounced at low consumption levels. Formally, the DD* $\times$ P_d *effect* (π_3) *is the product* $\tilde{s}\,\Delta\alpha\,c_0\,\bar{\ell}$*, where* $\bar{\ell}$ *is evaluated at mean liquidity. For poor households,* $\ell_i > \bar{\ell}$*, so the DD* $\times$ P_d *advantage is larger by* $\tilde{s}\,\Delta\alpha\,c_0 \cdot (\ell_i - \bar{\ell})$*. In the continuous model,* π_3 *and* π_5 *share the same structural root* $(\tilde{s}\,\Delta\alpha\,c_0)$ *and satisfy* $\pi_3 = \pi_5 \cdot \bar{\ell}$*, with* π_3 *evaluated at mean liquidity and* π_5 *capturing the slope with respect to household resources. In the empirical specification, where the continuous* ℓ_i *is replaced by the binary* $LowW_i$*, this exact proportionality no longer holds, but the model's qualitative prediction is preserved:* π_3 *and* π_5 *should share the same sign, and both should be positive whenever* $\Delta\alpha > 0$*.*

Summary. The model generates a clear hierarchy of signed predictions:

- $\pi_1 > 0$: Son preference drives continuation at each parity.
- $\pi_2 < 0$: Private-school prevalence raises costs and reduces fertility.
- $\pi_3 > 0$: The reduction is *smaller* for DD families (daughters are cheaper to educate privately).
- $\pi_4 < 0$: Liquidity constraints amplify the deterrent effect of P_d.
- $\pi_5 > 0$: The DD cost advantage is amplified among the liquidity-constrained.

The coefficient π_1 reflects son preference ($\theta > 0$), which is a preference parameter independent of the schooling market. The remaining four coefficients (π_2 through π_5) originate from a single structural mechanism: the gender gap in private-school costs, $\Delta\alpha = \alpha_b - \alpha_g > 0$, transmitted through the household's budget constraint. The direct cost of the *next* child is gender-symmetric in expectation; it is the total cost burden of *existing* children—higher for families with sons in private school—that generates the heterogeneous response to P_d across family compositions and wealth levels. Note also that the model does not predict a standalone effect of household wealth on fertility through this channel: the liquidity parameter ℓ_i enters the continuation condition only through its interaction with P_d. Any direct effect of wealth on fertility in the empirical specification captures other mechanisms outside the model.

From model to empirical specification. The theoretical latent index (17) is a stylised representation that abstracts from several features of the data. The empirical specifications in Section 5 differ in three ways: (i) they distinguish three sex-composition categories (DD, Mixed, SS) rather than two, since Mixed families (one son, one daughter) face intermediate schooling costs; (ii) they include district fixed effects δ_d that absorb the level of P_d, so that the P_d coefficient is identified from within-district variation across survey waves; and (iii) they include a rich set of individual-level controls $\mathbf{X}_i$. Despite these differences, the mapping from structural predictions to empirical coefficients is direct. Table 6 summarises the correspondence.

Table 6: Mapping model predictions to empirical coefficients

Model	Prediction	Empirical coefficient	Specification
$\pi_1 > 0$	Son preference	$\beta_1(\text{DD}) > 0$	A, B, C
$\pi_2 < 0$	P_d raises costs	$\beta_3(P_d) < 0$	A
$\pi_3 > 0$	DD $\times$ P_d	$\beta_4(\text{DD} \times P_d) > 0$	B
$\pi_4 < 0$	P_d $\times$ LowW	$\beta_7(P_d \times \text{LowW}) < 0$	C
$\pi_5 > 0$	DD $\times$ P_d $\times$ LowW	$\beta_9(\text{DD} \times P_d \times \text{LowW}) > 0$	C

5 Empirical Strategy and Specification

5.1 Data

We use individual-level microdata from the National Family Health Survey (NFHS), rounds 4 (2015–16) and 5 (2019–21), which are India's implementation of the Demographic and Health Surveys (DHS) programme. Each round surveys approximately 600,000–700,000 ever-married women aged 15–49 and contains complete birth histories, demographic characteristics, and household identifiers at the district level. We restrict our analysis sample to ever-married women aged 25–49 to focus on women with near-completed parity decisions.

District-level private-school prevalence is measured using the Unified District Information System for Education (UDISE), which provides annual school-level data on management type (government, private aided, private unaided) for all recognised schools in India. We compute

$$P_d = \frac{\text{Private unaided schools} + \text{Private aided schools}}{\text{Total schools in district } d}, \tag{24}$$

using the UDISE round contemporaneous with each NFHS wave (2015–16 for NFHS-4, 2019–20 for NFHS-5). The merged dataset covers approximately 500 matched districts per wave, yielding a pooled sample of roughly 900,000 women.

5.2 Measuring Son Status at Each Parity

A critical feature of our specification is that we measure the sex composition of children *at the time of the parity decision*, not at the time of the survey. We use the birth-order variables recorded in the NFHS birth history ($b4_k$ = sex of the k-th birth) to reconstruct the sex of the first n children for each woman at parity n.

At parity 2, we construct three mutually exclusive indicators for the sex composition of the first two births:

$$\begin{aligned} &\text{DD}_i = \mathbf{1}(\text{first 2 births both daughters}), \quad \text{Mixed}_i = \mathbf{1}(\text{one son, one daughter}), \\ &\text{SS}_i = \mathbf{1}(\text{first 2 births both sons}). \end{aligned} \tag{25}$$

Similarly, at parity 1 we define $\text{NoSon}_{i,1} = \mathbf{1}(\text{first birth is a daughter})$, and at parity 3 we define $\text{NoSon}_{i,3} = \mathbf{1}(\text{no son among first 3 births})$.

This construction is essential because using the *current* number of living sons (as commonly done) conflates the parity-n decision with subsequent fertility outcomes, generating mechanical reverse causality: women who continued past parity n have more children and therefore mechanically more sons.

5.3 Specification A: Contemporaneous Parity Progression Model

A standard approach to estimating parity progression models from cross-sectional data is to define the outcome as $\mathbf{1}(\text{CEB} > n)$ for women with $\text{CEB} \geq n$. This retrospective measure conflates the parity-n decision with the woman's entire subsequent fertility history and, critically, assigns the current district-level private-school share P_d to a decision that may have been made years or even decades earlier, when the schooling environment was very different. To ensure that the outcome and the treatment are contemporaneous, we adopt a flow-based measure.

For the parity-2-to-3 transition, we define the *at-risk* sample as women whose second child (identified by birth order from the DHS birth history) was born more than 36 months before the survey, and who had not yet had a third child at that point. These are women who were at parity exactly two at the start of a 36-month observation window ending at the survey date. The outcome is a binary indicator for whether the woman had a third birth within this 36-month window:

$$Y_i = \mathbf{1}\big(\text{3rd birth occurred in the 36 months before the survey}\big). \tag{26}$$

This construction has three advantages. First, the outcome is *contemporaneous* with the observed P_d: both refer to approximately the same two-year period. Second, right-censoring is eliminated by design—the observation window is closed at the survey date. Third, the sex composition of the first two children was determined well before the window opened and is therefore predetermined with respect to the outcome.

The baseline specification at parity 2 is:

$$Y_i = \alpha + \beta_1 \, \text{DD}_i + \beta_2 \, \text{Mixed}_i + \beta_3 \, P_d + \mathbf{X}_i' \mathbf{\Gamma} + \delta_d + \tau_t + \gamma_a + \varepsilon_{id}, \tag{27}$$

where SS_i (two sons) is the omitted reference category, P_d is the district-level private-school share, $\mathbf{X}_i$ is a vector of individual and household controls, δ_d are district fixed effects, τ_t are survey-wave fixed effects, and γ_a are single-year age fixed effects.

The inclusion of district fixed effects absorbs all time-invariant confounders at the district level. However, district-level private-school prevalence P_d may still be endogenous to time-varying local conditions (income growth, urbanisation, changing aspirations). Even though the interactions with individual-level characteristics (son composition, wealth) are less susceptible to this concern, we include in all specifications the following vector $\mathbf{X}_i$ of time-varying controls:

- Mother's years of education;
- Husband's education level (categorical: none, primary, secondary, higher);
- Urban/rural residence;
- Relative household wealth (an indicator for being below the within-district median

of the DHS wealth index score);

- Religion (Muslim indicator);
- Caste/tribe (Scheduled Caste, Scheduled Tribe, Other Backward Class indicators);
- Media exposure index (sum of newspaper, radio, and TV exposure frequencies).

We define the low-wealth indicator relative to the *within-district* wealth distribution rather than using national wealth quintiles. This avoids a composition problem: the share of women classified as "poor" under national quintiles varies systematically with P_d (63% in low-P_d districts versus 29% in high-P_d districts), confounding the wealth interaction. With within-district relative wealth, the share below the median is mechanically 50% in every district, ensuring that $P_d \times \text{LowWealth}$ interactions capture genuine differential responses by wealth rather than compositional differences in who is "poor" across districts.

All regressions are weighted by NFHS individual sampling weights. Standard errors are clustered at the district level throughout, since district-level private-school prevalence P_d is the treatment variable. District fixed effects absorb all time-invariant district characteristics—including persistent differences in income, culture, kinship systems, urbanisation, and infrastructure. With two survey waves, the P_d coefficient is identified from within-district changes in private-school share between 2015–16 and 2019–20. All interactions between P_d and individual-level variables (sex composition, wealth) are additionally identified from cross-individual variation within districts.

Identification as a continuous-treatment difference-in-differences. Specification (27) is a two-way fixed-effects difference-in-differences (DiD) design in which the treatment enters as a *continuous* dose—the district private-school share P_d. Because we observe only two survey waves and treatment intensity changes contemporaneously across all districts, the recent critiques of staggered-adoption DiD—which concern the "forbidden" comparisons between units treated at different times and the negative weighting these can induce (Goodman-Bacon, 2021; de Chaisemartin and D'Haultfœuille, 2020)—do not apply to our setting: there is no variation in treatment *timing*. The concerns specific to *continuous*-treatment DiD, however, do apply (Callaway et al., 2024). With a continuous dose, the two-way fixed-effects estimand recovers not a single causal effect but a weighted average of dose-specific causal responses (the average causal response on the treated). This average admits a clean causal interpretation only under conditions stronger than the textbook parallel-trends assumption: (i) *strong parallel trends*, requiring that, absent the expansion of private schooling, districts would have followed parallel fertility paths at *every* level of the dose, not merely on average; and (ii) *no selection into dose on gains*, requiring that districts experiencing larger increases in P_d are not systematically those that would have exhibited larger (or smaller) fertility responses. When the latter fails, the implicit estimand weights can turn negative and the weighted average loses its causal interpretation.

Limits of testability and our identifying strategy. We cannot fully validate these conditions with two survey waves. The parallel trends assumption is not directly testable in the absence of a pre-treatment period, and the balance diagnostics in Table 3 reveal a baseline difference in children ever born between high- and low-dose districts (normalised difference ≈ 0.20, the largest among all covariates), which raises a concrete mean-reversion concern: if high-dose districts begin from different fertility levels, regression toward the mean could be mistaken for a treatment effect. We confront this directly in Section 7, where we allow districts on different baseline-fertility trajectories to trend differentially and examine the dose–response across the distribution of ΔP_d. More fundamentally, we do not rest our causal claims on the level coefficient β_3 alone. Our preferred evidence comes from the son-composition interaction $\mathrm{DD}_i \times P_d$ introduced in Specification B (Section 5.4), which is identified from *within-district, cross-household* variation. District fixed effects, and any district-level mean reversion, difference out of this comparison, so the interaction is robust to precisely the dose-endogeneity and parallel-trends concerns that the level effect cannot fully escape. We therefore interpret β_3 as a dose–response that is informative under the conditions above, while treating the interaction evidence as the core of our identification.

A potential concern with our approach relates to the inexact timing alignment between the treatment and the fertility outcome. Throughout the paper, we refer to the 36-month parity progression measure as "contemporaneous," although this should be interpreted as referring to the survey period rather than to exact calendar-year alignment. In NFHS-5, births included in the outcome window may have occurred up to three years prior to the survey date, while treatment is measured using district-level private-school prevalence in 2019–20. This introduces some measurement error in exposure if fertility decisions respond to earlier values of private-school prevalence. However, two features of the data mitigate this concern. First, private-school prevalence evolves gradually within districts: annual within-district changes are small relative to the cross-wave variation that identifies our estimates. As a result, assigning a single survey-period value introduces limited misclassification. Second, our empirical strategy relies on within-district changes across survey waves, with district fixed effects absorbing level differences. Any remaining timing mismatch therefore acts as classical measurement error, which would tend to attenuate estimated effects toward zero.

5.4 Specification B: The Son-Composition–Private-School Interaction

The central test of the model is whether private-school prevalence modifies the son-preference stopping rule. At parity 2:

$$\begin{aligned}\Pr(\text{Birth}_3 \mid \text{CEB} \geq 2) = \alpha &+ \beta_1\,\text{DD}_i + \beta_2\,\text{Mixed}_i + \beta_3\,P_d \\ &+ \beta_4\,\text{DD}_i \times P_d + \beta_5\,\text{Mixed}_i \times P_d \\ &+ \mathbf{X}_i'\mathbf{\Gamma} + \delta_d + \tau_t + \gamma_a + \varepsilon_{id}.\end{aligned} \tag{28}$$

The coefficient β_4 is the key parameter (here β_3 is the P_d main effect—the response of the omitted two-son group). The model in Section 4 delivers an unambiguous prediction for its sign: $\beta_4 > 0$ (corresponding to $\pi_3 > 0$ in Table 6). The mechanism runs through the shadow-price channel: families with only daughters have lower total schooling expenditures on existing children, a less-stretched budget, and therefore a lower effective marginal cost of an additional child. They are *less deterred* by P_d than families with sons, who already bear the cost of private schooling. A positive β_4 thus confirms the "daughters are cheap" channel formalised in the model; a negative or zero β_4 would reject it.

This interaction is more informative than a simple $P_d \times$Parity specification because the model predicts that the relevant heterogeneity runs through *son composition*, not parity alone: once gender-differentiated schooling costs exist, the state variable that matters is whether the household has a son, not merely how many children it has.

5.5 Specification C: Triple Interaction with Liquidity Constraints

The richest reduced-form test incorporates the model's liquidity channel. At parity 2:

$$\begin{aligned}\Pr(\text{Birth}_3 \mid \text{CEB} \geq 2) = \alpha &+ \beta_1\,\text{DD}_i + \beta_2\,\text{Mixed}_i + \beta_3\,\text{LowWealth}_i + \beta_4\,P_d \\ &+ \beta_5\,\text{DD}_i \times P_d + \beta_6\,\text{Mixed}_i \times P_d + \beta_7\,P_d \times \text{LowWealth}_i + \beta_8\,\text{DD}_i \times \text{LowWealth}_i \\ &+ \beta_9\,\text{DD}_i \times P_d \times \text{LowWealth}_i + \mathbf{X}_i'\mathbf{\Gamma} + \delta_d + \tau_t + \gamma_a + \varepsilon_{id},\end{aligned} \tag{29}$$

where LowWealth_i is an indicator for being below the within-district median wealth score. The coefficient β_9 captures whether the DD cost advantage is amplified among liquidity-constrained households. The model predicts $\beta_9 > 0$ (corresponding to $\pi_5 > 0$ in Table 6): because DD families have lower existing schooling costs, the shadow-price differential between DD and non-DD families is larger when the budget is already tight. That is, the "daughters are cheap" force should be most visible among poor families. The coefficient β_7 ($P_d \times$ LowWealth) captures the direct liquidity amplification of P_d costs, predicted to be negative ($\pi_4 < 0$).

5.6 Testable Predictions

The model generates the following empirically testable predictions:

Prediction 1 (Son-Preference Stopping Rule)**.** *At a given parity, households whose first n births include no son are more likely to have another child:* $\beta_1(DD) > 0$ *in Specification A.*

Prediction 2 (Private-School Prevalence and Liquidity)**.** *Higher district private-school prevalence is associated with lower parity progression, particularly among poorer households:* $\beta_7(P_d \times LowWealth) < 0$ *in Specification C.*

Prediction 3 (Son Composition as the Relevant Heterogeneity)**.** *The effect of* P_d *on fertility progression depends on* son composition *rather than on parity alone. Specification B is more informative than* $P_d \times Parity$*.*

Prediction 4 (The DD $\times$ P_d Interaction)**.** *The model predicts* $\beta_4(DD \times P_d) > 0$ *in Specification B (corresponding to* $\pi_3 > 0$*): families with only daughters are less deterred by private-school prevalence because their lower existing schooling costs produce a lower shadow price of the budget constraint.*

The foundational assumption underlying Predictions 2–4—that boys are more likely to be enrolled in private schools than girls ($\alpha_b > \alpha_g$)—is not tested within the paper's estimating framework, but is documented using external evidence in Section 2 (see the discussion of gender differences in private-school enrolment in Section 2.4).

6 Results

This section presents the estimation results for the model's predictions using the contemporaneous parity progression measure described in Section 5.3. All specifications include district fixed effects, survey-wave fixed effects, and single-year age fixed effects, with standard errors clustered at the district level. The at-risk sample for the parity-2-to-3 transition comprises 37,251 women who were at parity exactly two at the start of the 36-month observation window, of whom 11.3 percent had a third birth within the window (12.4% in NFHS-4, 10.1% in NFHS-5).

6.1 First-Stage: Does P_d Predict Private Preschool Enrolment?

Before presenting the fertility results, we verify that the district-level private-school share P_d is a meaningful proxy for household exposure to the cost of child quality, and that this exposure is gendered. The model assumes that higher P_d raises the probability of private (fee-charging) enrolment and that this effect is stronger for sons ($\alpha_b > \alpha_g$). We test both claims directly in the NFHS-5 data using information on preschool attendance and type.

Data. We use the NFHS-5 Household Member Recode (PR) file, which records for each child aged 3–5 whether the child currently attends preschool (variable SH17) and, if so, the type of preschool (variable SH18: ICDS government, other government, private, or other). This information is collected only for the preschool margin; for school-age children aged 6 and above, NFHS records attendance and level but not whether the school is government or private.[7] The sample comprises 83,856 children aged 3–5 in 593 districts, of whom 22% attend some form of preschool and 15% attend a private preschool (with 30% of enrolled children attending private schools). We merge each child to the 2019–20 UDISE P_d in her district of residence.

Specification. We estimate two complementary specifications. The first uses cross-district variation with state fixed effects:

$$\text{PrivPre}_i = \gamma_1 P_d + \gamma_2 \text{Boy}_i + \gamma_3 \text{Boy}_i \times P_d + \mathbf{Z}_i'\mathbf{\Pi} + \eta_s + \nu_i, \tag{30}$$

where PrivPre_i is a dummy for attending private preschool, $\mathbf{Z}_i$ includes single-year age dummies and an urban indicator, η_s are state fixed effects, and standard errors are clustered by district. The coefficient γ_1 measures how cross-district variation in P_d translates into differences in private preschool enrolment; γ_3 tests whether boys respond differentially.

The second specification uses within-household (sibling) variation by replacing state fixed effects with household fixed effects, absorbing all household-level determinants of school choice (income, education, caste, parental preferences). With household fixed effects, the level effect of P_d and the direct effect of Boy are absorbed; only the interaction $\text{Boy}_i \times P_d$ is identified, from comparing the gender gap in private enrolment *within* the same household across districts with different P_d.

Results. Table 7 reports both sets of estimates.

[7] This limitation is precisely why our main fertility analysis uses the aggregate district-level UDISE P_d as a proxy for the cost of quality rather than household-level enrolment: child-level schooling type is not observed in the survey for school-age children. The preschool sample provides a direct test of the mechanism for the age group where the variable *is* observed.

Table 7: First-Stage: Private-School Prevalence and Preschool Enrolment (NFHS-5, ages 3–5)

	Panel A: State FE			Panel B: Household FE		
	(1) Any preschool	(2) Private preschool	(3) Private | enrolled	(4) Any preschool	(5) Private preschool	(6) Private | enrolled
P_d	0.197	0.559***	1.000***	—	—	—
	(0.130)	(0.094)	(0.134)			
Boy	−0.024***	0.018***	0.030***	−0.007	0.026	−0.023
	(0.009)	(0.007)	(0.012)	(0.022)	(0.019)	(0.036)
Boy × P_d	0.021	−0.096	−0.037	−0.104	−0.096	0.528*
	(0.057)	(0.060)	(0.118)	(0.179)	(0.146)	(0.294)
State FE	Yes	Yes	Yes	—	—	—
Household FE	—	—	—	Yes	Yes	Yes
Observations	83,856	83,856	39,785	9,502	9,502	3,358

Notes: Linear probability models. Sample: children aged 3–5 in NFHS-5 (2019–21). Outcome in columns (1) and (4) is a dummy for currently attending any preschool. In columns (2) and (5) it is a dummy for attending a private preschool (0 if not enrolled or in a government preschool). In columns (3) and (6) the sample is restricted to children currently enrolled in preschool and the outcome is a dummy for private (vs. government) preschool. The near-unity coefficient on P_d in column (3) is not mechanical: P_d measures the share of *schools* that are private (UDISE), while the outcome measures the share of enrolled *children* in private preschool (NFHS). The two are constructed from different data sources and units of observation. Estimates are stable across subsamples, with coefficients ranging from 0.76 to 1.23. P_d is the 2019–20 district private-school share from UDISE. Controls: single-year age dummies and an urban indicator in all columns. Panel A uses state fixed effects; Panel B uses household fixed effects, which absorb the level of P_d and the main effect of Boy. Standard errors clustered by district. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Pd strongly predicts private enrolment. Columns (2) and (3) of Panel A show that P_d is a powerful predictor of private preschool attendance. The coefficient of +0.559 ($p < 0.01$) in the unconditional specification implies that a 10 percentage-point increase in the district private-school share raises a child's probability of attending private preschool by 5.6 pp. Conditional on being enrolled in some form of preschool (column 3), the effect is essentially one-for-one: the coefficient of +1.000 ($p < 0.01$) means that the fraction of enrolled children attending private (vs. government) preschool tracks the fraction of schools that are private almost exactly. This is an economically very large effect and confirms that P_d is a strong proxy for actual exposure to the private-school market.

Boys are more likely than girls to attend private preschool. The unconditional Boy coefficient is +0.018 ($p < 0.01$) in column (2), and +0.030 ($p < 0.01$) in column (3) conditional on enrolment. These magnitudes match the descriptive gap in the raw data (15.6% of boys vs. 14.6% of girls attend a private preschool; among enrolees, 31.1% of boys vs. 28.1% of girls), and are consistent with the gendered enrolment evidence documented in Section 2.

Within households, boys are differentially more likely to be placed in private preschool in high-P_d districts. The sibling-comparison specification in Panel B is the cleanest test of the differential slope: it absorbs all household-level confounders and asks whether, within the same family, the brother–sister gap in private enrolment is larger in districts with more private schools. In column (6), conditional on both siblings being enrolled, the Boy $\times$ P_d coefficient is +0.528 ($p = 0.07$)—economically large and marginally significant. Within a household where both a boy and a girl are enrolled in preschool, a 10 pp increase in district P_d widens the brother–sister gap in private enrolment by about 5.3 pp. This finding directly matches Nandi et al. (2023)'s within-household estimate and provides mechanism-level evidence, in our own data and empirical setting, that P_d operates through a gendered cost channel. The cross-district Boy $\times$ P_d interaction in Panel A is insignificant, which is expected: cross-district variation conflates the gender gap with other district-level factors that household fixed effects absorb.

Implications. These first-stage results address two potential concerns about the fertility analysis that follows. First, they confirm that P_d is not merely a correlate of general district development—it translates one-for-one into actual private enrolment. Second, they confirm that within households the cost channel is gendered: boys face a higher probability of private-school costs than their sisters as P_d rises. The fertility estimates in the next subsections can therefore be interpreted as responses to a meaningful, gendered price of child quality rather than to a generic district-level modernisation proxy.

6.2 Son-Preference Stopping Rule (Prediction 1)

Table 8 reports the baseline contemporaneous parity progression regressions. At the parity-2-to-3 margin (column 2), son preference is stark: relative to families whose first two births were both sons (SS), families with two daughters (DD) are 16.8 percentage points more likely to have a third birth within the 36-month window ($p < 0.001$). Families with one son and one daughter (Mixed) are 2.0 percentage points more likely to continue ($p < 0.001$)—present but much smaller, consistent with the model's prediction that the son-preference motive is strongest when no son has yet been obtained.

Notably, the district-level private-school share P_d already shows a significant negative association with parity progression in this baseline specification: $\hat{\beta} = -0.306$ ($p = 0.02$). A one-standard-deviation increase in P_d (approximately 7 percentage points) is associated with a 2.1 percentage-point reduction in the 36-month third-birth probability—a quantitatively meaningful effect relative to the baseline rate of 11.3 percent.

At the parity-1-to-2 margin (column 1), the sex of the first child has a significant positive effect (+0.080, $p < 0.001$): women whose first birth was a daughter are more likely to have a second child within the window. The P_d coefficient at this margin is positive and insignificant (+0.164, $p = 0.51$), consistent with the fact that there are no existing schooling costs to tighten the budget—provided the first child is still below schooling age—so the shadow price is low, and the direct expected cost of the next child is smaller than at parity 2.

Table 8: Prediction 1: Son-Preference Stopping Rule (Contemporaneous)

	(1) 2nd birth in 36mo (at parity 1)	(2) 3rd birth in 36mo (at parity 2)
Birth-order sex composition		
Girl first	0.080***	
	(0.009)	
DD (two daughters)		0.168***
		(0.010)
Mixed (one son, one daughter)		0.020***
		(0.004)
Private-school prevalence		
P_d	0.164	−0.306**
	(0.251)	(0.128)
Controls		
Low wealth (within-district)	0.068***	0.020***
	(0.011)	(0.004)
Mother's education (years)	0.002*	−0.002***
	(0.001)	(0.001)
Husband: higher education	0.016	−0.027***
	(0.018)	(0.008)
Muslim	0.056***	0.087***
	(0.013)	(0.011)
District FE	Yes	Yes
Wave FE	Yes	Yes
Age FE	Yes	Yes
Observations	21,792	37,251

Notes: Linear probability models. Outcome: binary indicator for whether the woman had the next birth within a 36-month window ending at the survey date. The at-risk sample comprises women who had reached exactly parity n more than 36 months before the survey and had not yet progressed beyond parity n at the start of the window. Sex composition is measured using the birth-order sex of the first n children from DHS birth-history records. Reference category at parity 2: SS (two sons). Low wealth is defined as below the within-district median of the DHS wealth index score. Additional controls include husband's education, urban residence, caste/tribe, and media exposure. Standard errors clustered by district. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

6.3 Private-School Prevalence and Parity Progression (Predictions 2–4)

Table 9 reports the core specifications testing Predictions 2 through 4, all using the contemporaneous 36-month parity progression measure at the parity-2-to-3 margin.

Column 1 presents the baseline specification including P_d alongside the sex-composition dummies. The P_d coefficient is -0.306 ($p = 0.02$): a within-district increase in private-school prevalence significantly reduces the contemporaneous probability of a third birth. This is the direct cost-channel effect predicted by the model.

Column 2 adds interactions of P_d with the DD and Mixed indicators. The DD $\times$ P_d coefficient is +0.177 (SE = 0.138)—positive and economically meaningful, though not statistically significant at conventional levels ($p = 0.20$). The sign is consistent with the model's Prediction 3: families with only daughters, whose children are predominantly in free government schools, are less deterred by private-school prevalence than families who already have a son to educate. The Mixed $\times$ P_d coefficient is -0.071 (SE = 0.062)—negative and directionally consistent with the model (families with one son face the cost of privately schooling that son), though also imprecise ($p = 0.25$).

Column 3 adds the liquidity channel through the interaction of P_d with the within-district low-wealth indicator. The coefficient on $P_d \times$LowWealth is -0.053 (SE = 0.076)—negative, in the direction predicted by the model, but not statistically significant.

Column 4 presents the full triple-interaction specification (Specification C). Two results are notable. First, the DD $\times$ P_d interaction now reaches marginal significance at +0.288 (SE = 0.161, $p = 0.07$), providing suggestive evidence that families with only daughters are more likely to continue in high-P_d environments—the "daughters are cheap, keep trying" force from the model. Second, the DD $\times$ LowWealth coefficient is +0.066 ($p = 0.02$): within the same district, low-wealth families with two daughters are significantly more likely to try for a third child, consistent with tighter resource constraints intensifying the desire for a son. The DD $\times$ $P_d \times$ LowWealth triple interaction is -0.239 (SE = 0.242)—negative, in the predicted direction (cost pressure overwhelms son preference for constrained DD families), but not statistically significant. The magnitude is economically meaningful: at the mean P_d of 0.10, the implied differential effect for poor DD families is -2.4 percentage points. We interpret this as directional evidence consistent with the theory, while acknowledging that the data cannot sharply distinguish the triple-interaction effect from zero given the sample demands of the specification.

Table 9: Predictions 2–4: Private-School Prevalence, Son Composition, and Liquidity (Contemporaneous)

	(1) Baseline	(2) + DD×P_d	(3) + Liquidity	(4) Triple
Sex composition (ref: SS)				
DD	0.168***	0.149***	0.149***	0.119***
	(0.010)	(0.017)	(0.017)	(0.020)
Mixed	0.020***	0.028***	0.028***	0.016
	(0.004)	(0.009)	(0.009)	(0.010)
Private-school prevalence				
P_d	−0.306**	−0.287**	−0.264*	−0.288**
	(0.128)	(0.133)	(0.140)	(0.138)
Private-school × sex composition				
DD × P_d		0.176	0.177	0.288*
		(0.138)	(0.138)	(0.161)
Mixed × P_d		−0.071	−0.073	−0.058
		(0.062)	(0.063)	(0.066)
Liquidity channel				
Low wealth (within-district)	0.020***	0.020***	0.026***	0.001
	(0.004)	(0.004)	(0.009)	(0.012)
P_d × Low wealth			−0.053	−0.004
			(0.076)	(0.093)
DD × Low wealth				0.066**
				(0.028)
Mixed × Low wealth				0.027*
				(0.014)
DD × P_d × Low wealth				−0.239
				(0.242)
Mixed × P_d × Low wealth				−0.025
				(0.105)
Controls	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
Observations	37,251	37,251	37,251	37,251

Notes: Linear probability models. Outcome: binary indicator for having a third birth in the 36 months before the survey, among women at parity exactly two at the start of the window. P_d is the district-level private-school share from UDISE, contemporaneous with the outcome. The P_d coefficient is identified from within-district changes between NFHS-4 and NFHS-5, net of national time trends. Low wealth is defined as below the within-district median wealth score. Controls included in all specifications: mother's years of education, husband's education level (categorical), urban/rural residence, Muslim indicator, Scheduled Caste, Scheduled Tribe, Other Backward Class indicators, and media exposure index. The near-stability of the DD, Mixed, and DD × P_d coefficients between columns (2) and (3) reflects the near-orthogonality of P_d× LowWealth to the sex-composition interactions conditional on the fixed effects; the coefficients do change at the fourth decimal place. Standard errors clustered by district. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

6.4 Discussion of Findings

The main Pd effect. A consistent finding across specifications is the negative effect of district-level private-school prevalence on contemporaneous third-birth probability. The coefficient of approximately -0.29 implies that a one-standard-deviation increase in P_d (7 percentage points) reduces the 36-month third-birth rate by 2.0 percentage points—a roughly 18 percent reduction relative to the baseline rate of 11.3 percent. This effect survives the inclusion of district fixed effects, wave fixed effects, age fixed effects, and a rich set of individual controls, and is identified from within-district changes in private-school share between the two survey waves. We caution, however, that the *magnitude* of this district-level effect is not robust: as shown in Section 7.8, the P_d coefficient attenuates substantially and loses significance once we allow districts on different baseline-fertility trajectories to trend differentially. The level effect should therefore be read as suggestive; our most robust evidence is the within-district son-composition interaction discussed next, which is immune to this district-level confound.

Son composition and P_d. The DD $\times$ P_d interaction is positive across specifications, reaching marginal significance in the full triple specification ($+0.288$, $p = 0.07$). While imprecise in the pooled sample, it reaches conventional significance precisely where schooling costs are most salient—among older women, whose children are of school age (Table 13, column 4)—and once attenuation in P_d is addressed using supply-driven variation (Table 19), a pattern consistent with a genuine effect estimated with limited power rather than with the absence of an effect. The direction is consistent with the model's "daughters are cheap" force: families with only daughters, whose children are predominantly in free government schools, are less deterred by private-school prevalence. The implied total P_d effect for DD families (the sum of the P_d main effect and the DD $\times$ P_d interaction) is $-0.288 + 0.288 \approx 0$—private-school prevalence has essentially no effect on fertility for families with only daughters, while it strongly reduces fertility for SS and Mixed families. The Mixed $\times$ P_d coefficient is consistently negative (-0.058 to -0.073) and directionally consistent with the model's prediction that families with an existing son to educate are somewhat more deterred, though the coefficient does not reach conventional significance.

The liquidity channel. With the contemporaneous outcome and within-district wealth definition, the $P_d \times$LowWealth interaction is negative (-0.053) but imprecisely estimated. The point estimate is directionally consistent with the model (cost pressure operates for all wealth groups), though we cannot reject that the effect is identical across wealth levels. A notable finding in the full specification is that DD $\times$ LowWealth is positive and significant ($+0.066$, $p = 0.02$): low-wealth families with two daughters are more likely to continue than their wealthier counterparts within the same district, consistent with the model's prediction that resource-constrained families have stronger son-preference motives.

The triple interaction. The DD $\times$ $P_d\times$ LowWealth coefficient in the full specification is -0.239 (SE $= 0.242$). The model predicts a *positive* sign ($\pi_5 > 0$): the daughters-only cost advantage should be amplified among liquidity-constrained households. The point estimate is of the opposite sign but statistically insignificant and imprecisely estimated, reflecting the inherent difficulty of identifying a triple interaction in a sample of approximately 37,000 observations, particularly given that the DD $\times$ LowWealth cell comprises a small fraction of the population. We therefore do not read it as evidence for the liquidity-amplification prediction, though its imprecision means it is not strong evidence against it either.

Quantifying the contribution of private-school expansion. A back-of-envelope calculation provides a sense of the *potential* aggregate importance of the channel. Between NFHS-4 and NFHS-5, the mean district P_d increased by 1.3 percentage points. Taking the P_d coefficient from column 1 of Table 9 (-0.306), which captures the overall effect of private-school prevalence on parity progression, the predicted reduction in the 36-month third-birth rate is $0.306 \times 0.013 = 0.40$ percentage points. The observed decline in this rate between waves was 2.3 percentage points (from 12.4% to 10.1%), so the calculation implies a contribution of approximately **17 percent** of the observed decline in parity-2-to-3 progression over this period—an economically meaningful magnitude for a single mechanism operating at a single parity margin. We stress, however, that this figure is illustrative rather than a precise causal estimate: it rests on the P_d level coefficient, which Section 7.8 shows attenuates and loses significance once differential baseline-fertility trends are absorbed. The dose–response in Table 17 is consistent with a contribution of this order but is statistically imprecise, with most of the observed decline reflecting a common secular trend. We therefore read the figure as an upper bound on the channel's aggregate scale, and base our causal claims on the within-district interaction rather than on this decomposition.

Husband's education. Across all specifications, having a husband with higher education reduces the contemporaneous third-birth probability by approximately 2.7 percentage points ($p < 0.001$), even after controlling for the wife's own education. This is consistent with the household investment model in which both spouses' human capital matters for family-size decisions.

6.5 The Sex of the Third Birth: Quantity versus Selection

Our interpretation of the DD $\times$ P_d interaction as a *quantity*-margin effect rests on the model's assumption that the sex of the next child is a coin flip for every family type, so that the expected cost of the marginal child does not differ across DD, Mixed, and SS households. In India this assumption is not innocuous: families without a son are precisely those most likely to practise sex selection (Jayachandran, 2017; Anukriti, 2018), and if daughters-only families sex-select their third birth—and especially if they do so *differentially* where private schooling is more prevalent—the symmetry breaks and the interaction could partly reflect composition rather than quantity. We assess this directly.

Sex selection is indeed present. Among recent third births, the share that are male is 0.55 for daughters-only (DD) families and 0.56 for mixed families, against 0.50 for families that already have two sons; the DD advantage in third-birth maleness is statistically significant (Table 10, +0.08, $p < 0.05$, in both windows). This is the familiar son-targeting pattern.

The relevant question for our identification, however, is not whether sex selection exists but whether it *responds to* P_d. We therefore estimate, among women whose third birth occurred within the W months preceding the survey—so that the prevailing private-school environment is contemporaneous with the birth—a linear probability model for whether that birth is male, as a function of sex composition, P_d, and their interaction, with district, survey-wave, and age fixed effects. The coefficient on DD $\times$ P_d answers the question: did private-school expansion change son-targeting among daughters-only households?

It did not. The DD $\times$ P_d coefficient is *negative* in both the 36-month and 60-month windows (-0.47 and -0.33) and statistically insignificant; the more precisely estimated 60-month window bounds the effect below +0.17, ruling out any sizeable *increase* in selection. If anything, the point estimates suggest that expansion modestly *dampens* son-targeting—consistent with the logic of the model, since a higher P_d raises the effective price of a son and so weakens, rather than strengthens, the incentive to select one. We read this cautiously given the imprecision, but the direction is clear and the data provide no support for the opposite.

Two implications follow. First, the symmetry assumption holds where it matters: what the existing-cost interpretation requires is not that the next child be *literally* equally likely to be a boy or a girl, but that the sex composition of the *marginal* birth not shift with P_d—for only then could the DD $\times$ P_d quantity interaction be confounded by selection. Since the marginal birth's sex composition does not rise with P_d (and weakly falls), the quantity interaction is not a selection artifact, and the existing-cost channel remains its natural interpretation. Second, the finding speaks to a policy concern in the literature (Anukriti, 2018): we find no evidence that the expansion of fee-based schooling worsened son-targeting at this margin.

Table 10: Sex of the Third Birth: Does Private-School Expansion Affect Son-Targeting?

	(1) $W = 36$	(2) $W = 60$
DD (no son)	0.087**	0.078**
	(0.043)	(0.030)
Mixed	0.049	0.066**
	(0.041)	(0.031)
P_d	−1.204	−0.370
	(0.879)	(0.545)
DD × P_d	−0.472	−0.330
	(0.367)	(0.254)
Mixed × P_d	−0.111	−0.245
	(0.327)	(0.223)
Controls	Yes	Yes
District / Wave / Age FE	Yes	Yes
Observations	5,970	10,404

Notes: Linear probability models. Sample: women whose third birth occurred within the W months preceding the survey, so that district private-school prevalence P_d is contemporaneous with the birth. Outcome: third birth is male (= 1). Reference category for sex composition is SS (first two births both sons). Among recent third births, the share male is 0.55 (DD), 0.56 (Mixed), and 0.50 (SS). Controls: mother's years of education, husband's education level, urban/rural residence, Muslim, Scheduled Caste, Scheduled Tribe, and Other Backward Class indicators, and media exposure. District, survey-wave, and single-year-age fixed effects throughout. Standard errors clustered by district in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

6.6 Eldest-Child Preference and the Birth-Order Test

The Mixed (one-son-one-daughter) cell pools two types that differ only in birth order: boy-first (BG, the eldest is a son) and girl-first (GB, the eldest is a daughter). Both have the same son *count*, so the model—in which cost depends on the number of sons, not their order—predicts that BG and GB families behave identically. A leading alternative does not: eldest-son preference (Jayachandran and Pande, 2017) and selective entry into each cell could make boy-first and girl-first families differ. The split is therefore a discriminating test between the cost mechanism and a birth-order account. We first ask whether the cost channel could even depend on order, then test behaviour directly.

Does the cost channel depend on birth order? We examine whether eldest children—and eldest sons in particular—are differentially enrolled in private schooling, using the NFHS-5 preschool data (ages 3–5), the only level at which government and private attendance are distinguished. Table 11 documents a strong eldest-child premium: among boys, firstborn children are 5.5 percentage points more likely to attend private preschool unconditionally and 8.5 points more likely conditional on attendance (columns 1–2). Crucially, however, this premium is *not son-specific*: in the pooled sample of all children the firstborn × boy interaction is economically and statistically zero (column 3, −0.006), so firstborn daughters enjoy essentially the same private-schooling advantage as firstborn sons. Nor does the premium scale with P_d (column 4).

Because the eldest premium attaches to the firstborn *child* regardless of sex, a boy-first family (firstborn son, second-born daughter) and a girl-first family (firstborn daughter, second-born son) carry the *same* total private-schooling cost and should respond identically to private school expansion. We test this directly by splitting the Mixed cell in the parity-progression specification. Table 12 reports Specification B with Mixed disaggregated into BG and GB. The two are statistically indistinguishable on both margins the model restricts: the continuation main effects are 0.027 (BG) and 0.029 (GB), equal at $p = 0.88$; and the responses to P_d are −0.078 (BG) and −0.069 (GB), equal at $p = 0.90$. Boy-first and girl-first families behave identically, exactly as the cost mechanism predicts—given that there is no order-dependent cost for behavior to respond to.

Table 11: The Eldest-Child Premium in Private Schooling (NFHS-5 Preschool, Ages 3–5)

	(1) Private (boys)	(2) Private | enrol. (boys)	(3) Private (all kids)	(4) Private, boys (P_d)
Firstborn	0.055***	0.085***	0.060***	0.061***
	(0.005)	(0.009)	(0.005)	(0.008)
Boy			0.009*	
			(0.004)	
Boy × Firstborn			−0.006	
			(0.006)	
P_d				0.452***
				(0.102)
Firstborn × P_d				−0.031
				(0.064)
Sample	Boys	Enrol. boys	All kids	Boys
Fixed effects	District	District	District	State
Observations	42,176	19,622	81,527	42,176

Notes: Linear probability models on NFHS-5 children aged 3–5. Outcome is attendance at a *private* preschool: unconditional (= 1 if attending a private preschool, 0 otherwise) in columns (1), (3), (4); conditional on attending any preschool (= 1 if that preschool is private) in column (2). "Firstborn" indicates the child is the eldest of the mother's co-resident children (ages ≤ 18); birth order is proxied from co-residence. Among enrollees, the private share is 0.36 for firstborn boys versus 0.26 for later-born boys (0.33 versus 0.23 for girls). Controls: child age and urban residence. Column (4) uses state rather than district fixed effects so that P_d—fixed within district in a single wave—is estimable. Standard errors clustered by district in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 12: Birth-Order Split of the Mixed Cell: Boy-First versus Girl-First (Parity 2)

	(1) Baseline (Mixed)	(2) Mixed split
DD	0.149***	0.149***
	(0.017)	(0.017)
Mixed	0.028***	
	(0.009)	
BG (boy first)		0.027**
		(0.011)
GB (girl first)		0.029***
		(0.010)
P_d	−0.264*	−0.265*
	(0.140)	(0.141)
DD × P_d	0.177	0.177
	(0.138)	(0.138)
Mixed × P_d	−0.073	
	(0.063)	
BG × P_d		−0.078
		(0.072)
GB × P_d		−0.069
		(0.074)
Tests of the model's no-order predictions:		
H_0: BG = GB (level)		$p = 0.88$
H_0: BG×P_d = GB×P_d		$p = 0.90$
Controls / District / Wave / Age FE	Yes	Yes
Observations	37,251	37,251

Notes: Linear probability models, Specification B. Column (1) is the baseline with the Mixed (one-son-one-daughter) cell combined; column (2) splits it into boy-first (BG) and girl-first (GB). Reference category is SS (two sons). Low-wealth main effect and P_d× LowWealth, together with the full set of controls (mother's and husband's education, urban, Muslim, caste, media exposure) and district, survey-wave, and age fixed effects, are included throughout but suppressed for brevity. The two F-tests report the model's predictions that boy-first and girl-first families behave identically. Standard errors clustered by district in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

7 Robustness

We assess the sensitivity of our main findings through several exercises. The first six are reported together in Table 13 alongside the baseline estimates for direct comparison; each column uses the same Specification B from Table 9 column 3—with sex composition, P_d, DD × P_d, Mixed × P_d, P_d× LowWealth, and the full set of controls—applied to different subsamples or specifications. We then turn to two diagnostics motivated directly by the continuous-dose design discussed in Section 5.3—a test for differential baseline-fertility trends (mean reversion) and a dose–response decomposition—before presenting a supply-side consistency check in Section 7.10.

Before turning to the individual exercises, Figure 3 summarises the central object of the paper—the DD $\times$ P_d interaction—across every specification we estimate. The coefficient is positive in all twenty cases, with point estimates clustered between roughly 0.12 and 0.29 in the pooled and robustness specifications and larger still where the private market is most active (urban households, +0.52) or where P_d is restricted to its supply-driven component (2SLS, +0.42 to +0.59). Most pooled estimates are individually imprecise—their confidence intervals include zero—but they are tightly bunched and uniformly signed, and the interaction attains conventional significance precisely in the cases where the model predicts the cost channel should bite hardest: among older women whose children have reached school age, among urban households facing the most developed private market, and once attenuation in P_d is addressed. We read the figure as the visual counterpart of our identification argument: the sign and rough magnitude of the son-composition interaction are stable across subsamples, alternative windows, migrant exclusions, development and modernisation controls, and the frozen-denominator and supply-driven treatments, while its precision varies with the informativeness of the subsample.

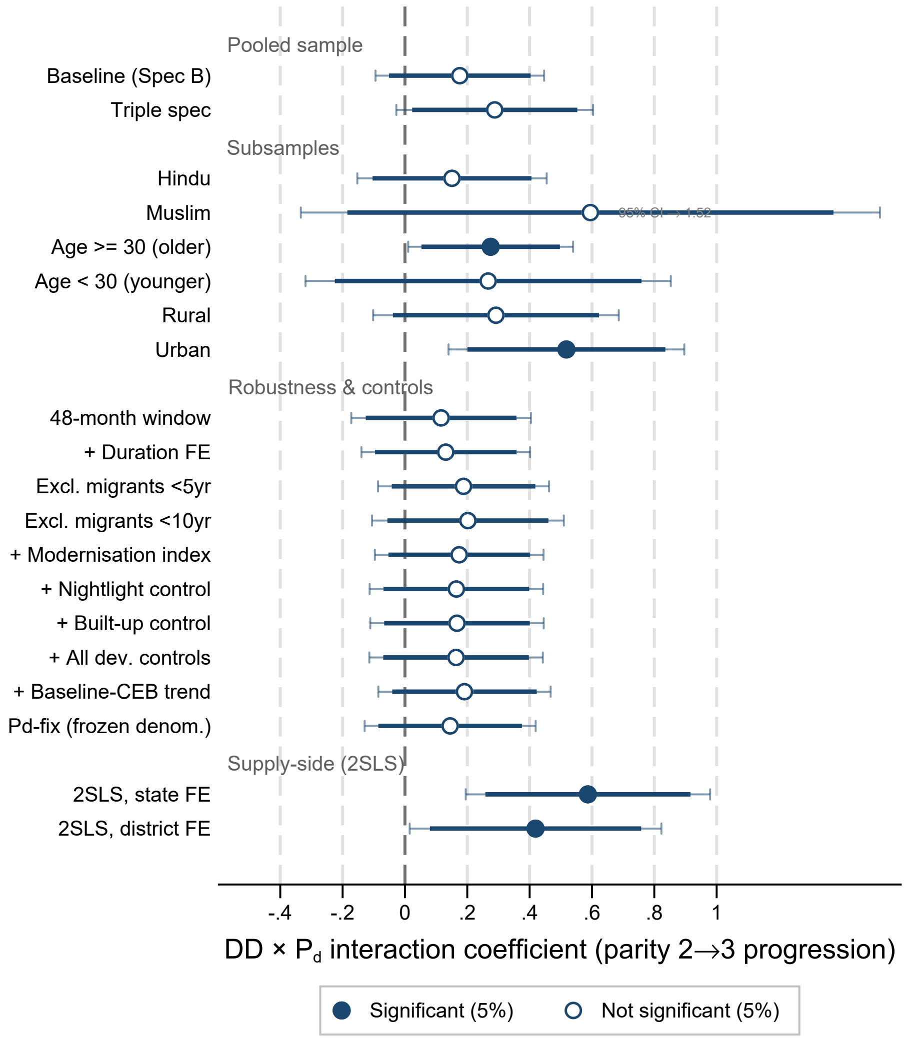

Figure 3: Stability of the DD × P_d Interaction Across Specifications

Notes: Each row plots the estimated DD × P_d coefficient (parity-2-to-3 progression) from a single specification, with its 90% (thick) and 95% (thin) district-clustered confidence interval. Solid markers denote significance at the 5% level; hollow markers are not significant at 5%. Estimates are taken from Tables 9 (pooled baseline and triple specification), 13 (religion, age, 48-month window, duration controls, migrant exclusions, modernisation index), 15 (urban/rural), 16 (baseline-fertility trend), 14 (satellite and asset development controls), 18 (P_d^{fix}, frozen denominator), and 19 (supply-side 2SLS). The point estimate is positive in every specification. The Muslim subsample's 95% confidence interval extends to approximately 1.52 (truncated at the axis for readability), reflecting its small sample (N = 2,837).

7.1 Hindu versus Muslim Heterogeneity

If the mechanism operates through the interaction of son preference with gendered private-school costs, the effects should be concentrated among Hindu families—who exhibit stronger son preference and higher rates of private-school enrolment—and attenuated among Muslim families, who have different schooling norms (including greater use of government schools and madrassas) and somewhat weaker son preference in the fertility-stopping dimension.

Columns (2) and (3) of Table 13 report the results. The DD coefficient is 0.164 ($p < 0.001$) among Hindu women but an insignificant 0.025 among Muslim women ($N = 2{,}837$). The P_d main effect is -0.411 ($p = 0.02$) for Hindus—strongly negative—and an insignificant $+0.919$ for Muslims. The cost channel operates with particular force among Hindus, who are the primary users of the private-school market. The DD $\times$ P_d interaction is $+0.151$ for Hindus (positive but imprecise, $p = 0.33$), directionally consistent with the "daughters are cheap" force. Among Muslims, the DD $\times$ P_d coefficient is $+0.595$—large in magnitude but imprecisely estimated (SE $= 0.474$), reflecting the small sample and correspondingly wide confidence interval. Formal tests cannot reject equality of either coefficient across the two groups (DD $\times$ P_d: $p = 0.38$; P_d: $p = 0.32$), a direct consequence of the very imprecise Muslim estimates; we therefore read the religion contrast as suggestive.

The religion split is informative despite the limited Muslim sample: the cost-channel effect (the strongly negative P_d coefficient) operates among Hindu women, who face the Indian private-school market most directly. Hindu and Muslim women in the same district face the same P_d; the differential response must reflect household-level differences in schooling norms and son preference rather than district-level confounders.

7.2 Age Heterogeneity

If private-school costs become salient only once children approach school age, the effect should be concentrated among older women (age ≥ 30) whose first two children are more likely to be of school age during the observation window. Columns (4) and (5) report the results. Among older women ($N = 27{,}898$), the P_d coefficient is -0.238 ($p = 0.05$) and the DD $\times$ P_d interaction is $+0.275$ ($p = 0.04$)—the only subsample in which this interaction reaches conventional significance. Among younger women ($N = 9{,}347$), the P_d coefficient is similar in magnitude (-0.234) but imprecise ($p = 0.55$), reflecting the smaller sample. The DD coefficient is much larger for younger women (0.287 vs. 0.074), consistent with stronger son-preference motives earlier in the reproductive life course. The two DD $\times$ P_d point estimates are in fact very close (0.275 versus 0.267) and a formal test cannot reject their equality ($p = 0.84$): the age contrast is one of *precision*—the older-women estimate is sharper—rather than of magnitude. This is consistent with, though it does not sharply confirm, the model's prediction that the cost channel operates through the existing schooling burden of children who have reached school age.

7.3 Alternative Observation Window: 48 Months

Column (6) extends the observation window from 36 to 48 months, allowing more time for parity progression at the cost of some reduction in contemporaneity. The at-risk sample comprises 35,268 women. The P_d coefficient strengthens to -0.341 ($p = 0.02$), consistent with the cost channel cumulating over a longer horizon. The DD coefficient rises to 0.189 ($p < 0.001$), and the DD $\times$ P_d interaction is +0.116 (positive, as in the baseline, though not significant). The pattern across windows is monotonic and stable: longer observation periods produce larger main effects while preserving the same sign structure for all interaction terms.

7.4 Controlling for Duration at Parity Two

A potential concern is that women in the at-risk sample differ in how long they have been at parity two when the observation window opens. Column (7) augments Specification B with a full set of fixed effects for years since the second birth (16 categories: 0 to 15+ years). The DD coefficient is virtually unchanged (0.140, $p < 0.001$). The P_d coefficient attenuates from -0.264 to -0.164 ($p = 0.22$), and the DD $\times$ P_d interaction attenuates from 0.177 to 0.131. Importantly, the attenuation of the P_d coefficient is *consistent with the mechanism* rather than evidence against it: if P_d genuinely reduces the probability of a third birth, women in high-P_d districts will mechanically exhibit longer durations at parity two—precisely because they are not progressing. Duration since the second birth is therefore partly a post-treatment variable, and conditioning on it absorbs part of the causal channel. The specification without duration controls provides the unconditional reduced-form effect, while the specification with duration controls provides a conservative lower bound. The qualitative pattern of all coefficients—negative P_d, positive DD $\times$ P_d—is preserved.

7.5 Excluding Recent Migrants

A potential identification threat is that households with low-fertility preferences selectively migrate to districts with expanding private schooling, generating a spurious negative association between P_d and fertility. The NFHS records how many years each woman has lived in her current place of residence (variable V104). In our analysis sample, 10.7% of women moved to their current district within the last four years; the remaining 89.3% have lived there for five or more years. Columns (8) and (9) of Table 13 re-estimate Specification B after progressively excluding recent migrants. Column (8) drops women who moved within the last four years ($N = 33{,}946$); column (9) applies a stricter cutoff, dropping all women who moved within the last ten years ($N = 26{,}197$—a 30% sample reduction, reflecting that most moves are patrilocal marriage migration rather than economic sorting). If fertility-selective migration were driving the P_d–fertility association, we would expect the coefficients to attenuate as we remove progressively more mobile women. Instead, the results are stable or *strengthen*: the DD $\times$ P_d interaction increases monotonically from +0.177 (baseline) to +0.188 (excluding <5 years) to +0.202 (excluding <10 years), and the $P_d \times$ LowWealth coefficient moves from −0.053 to −0.086 to −0.107. The cost-channel mechanism is, if anything, more pronounced among women who have been rooted in their district for a decade or more. This rules out fertility-selective sorting as a primary driver of the results.

7.6 Controlling for Local Modernisation

District fixed effects absorb time-invariant district characteristics but not within-district trends, so the P_d coefficient could partly reflect broader modernisation occurring alongside private-school expansion. To probe this, we construct a district-by-wave modernisation index—the standardised average of the district shares of households with electricity, a television, a refrigerator, a motorcycle/scooter, and a car, each measured identically in NFHS-4 and NFHS-5—and add it to Specification B. These are precisely the household correlates that satellite measures such as night-time luminosity and built-up area are designed to capture, so the index serves as an in-data analogue of those proxies. Because district fixed effects absorb its level, the index enters as a within-district modernisation *trend*. Column (10) of Table 13 reports the result. The index is itself a strong negative predictor of parity progression (-0.070, $p < 0.01$): districts modernising faster experienced larger declines in third-birth probability, consistent with a genuine modernisation channel. Controlling for it attenuates the P_d main effect from −0.264 to −0.192 (and to −0.13 when the index's components enter separately, driven by electrification and vehicle ownership), confirming that part of the district-level association does reflect contemporaneous modernisation. The DD $\times$ P_d interaction, however, is entirely unaffected ($0.177 \rightarrow 0.174$), as expected for a coefficient identified from within-district, cross-household variation. Reassuringly, the within-district change in the modernisation index is only weakly correlated with the change in P_d ($r = 0.08$), indicating that private-school expansion is not merely a proxy for local development. We confirm this directly with two independent satellite measures—night-time luminosity and built-up surface area—presented immediately after the main robustness table (Table 14).

Several conclusions emerge from Table 13. First, the son-preference stopping rule (the DD coefficient) is present across all subsamples, ranging from 0.074 (older women) to 0.287 (younger women), always highly significant except in the small Muslim sample. Second, the P_d main effect is consistently negative, reaching conventional significance in the baseline (-0.264, $p < 0.10$), for Hindus (-0.411, $p = 0.02$), for older women (-0.238, $p = 0.05$), and in the 48-month window (-0.341, $p = 0.02$). When duration-at-parity fixed effects are added, the P_d coefficient attenuates to -0.164; this attenuation is expected, since duration at parity two is partly a post-treatment variable—women in high-P_d districts who do not progress mechanically accumulate longer durations—and conditioning on it absorbs part of the causal channel. The sign is preserved across all specifications. Third, the DD $\times$ P_d interaction is consistently positive, reaching significance among older women ($+0.275$, $p = 0.04$): families with only daughters are less deterred by private-school prevalence, consistent with the "daughters are cheap" force. Fourth, the $P_d \times$LowWealth coefficient is negative or near-zero across subsamples, confirming that the cost channel operates broadly across the wealth distribution. Finally, progressively excluding recent migrants (columns 8–9) leaves all coefficients stable or strengthened—the DD $\times$ P_d interaction increases monotonically as more mobile women are removed—ruling out fertility-selective migration as a driver of the results. Adding a district-by-wave modernisation index (column 10) reinforces the same pattern: it modestly attenuates the P_d level effect while leaving the DD $\times$ P_d interaction unchanged.

Table 13: Robustness: Specification B Across Subsamples and Specifications (Parity 2)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	Baseline	Hindu	Muslim	Age $\geq$ 30	Age < 30	48-month	+ Duration FE	Excl. <5 yr	Excl. <10 yr	+ Modern. index
Sex composition (ref: SS)										
DD	0.149***	0.164***	0.025	0.074***	0.287***	0.189***	0.140***	0.162***	0.125***	0.149***
	(0.017)	(0.018)	(0.060)	(0.017)	(0.033)	(0.018)	(0.016)	(0.017)	(0.019)	(0.017)
Mixed	0.028***	0.035***	−0.032	0.008	0.073***	0.024**	0.030***	0.034***	0.023***	0.029***
	(0.009)	(0.009)	(0.039)	(0.008)	(0.022)	(0.010)	(0.008)	(0.010)	(0.009)	(0.009)
Private-school prevalence										
P_d	−0.264*	−0.411**	0.919	−0.238**	−0.234	−0.341**	−0.164	−0.234	−0.234	−0.192
	(0.140)	(0.172)	(1.334)	(0.120)	(0.394)	(0.152)	(0.134)	(0.147)	(0.146)	(0.134)
Private-school × sex composition										
DD × P_d	0.177	0.151	0.595	0.275**	0.267	0.116	0.131	0.188	0.202	0.174
	(0.138)	(0.155)	(0.474)	(0.135)	(0.299)	(0.147)	(0.138)	(0.140)	(0.157)	(0.138)
Mixed × P_d	−0.073	−0.088	0.073	−0.026	−0.090	−0.071	−0.103*	−0.097	−0.069	−0.078
	(0.063)	(0.062)	(0.292)	(0.053)	(0.179)	(0.075)	(0.060)	(0.067)	(0.062)	(0.063)
Liquidity channel										
Low wealth	0.026***	0.027***	0.047	0.014*	0.052**	0.032***	0.022**	0.029***	0.028***	0.026***
	(0.009)	(0.010)	(0.036)	(0.008)	(0.024)	(0.010)	(0.009)	(0.009)	(0.010)	(0.009)
P_d × Low wealth	−0.053	−0.046	−0.362	0.007	−0.157	−0.040	−0.046	−0.086	−0.107	−0.056
	(0.076)	(0.079)	(0.262)	(0.062)	(0.192)	(0.084)	(0.072)	(0.073)	(0.083)	(0.076)
Local modernisation										
Modernisation index										−0.070***
										(0.014)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Wave FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Duration FE	No	No	No	No	No	No	Yes	No	No	No
Observations	37,251	29,932	2,837	27,898	9,347	35,268	37,251	33,946	26,197	37,251

Notes: Linear probability models. All columns use Specification B from Table 9 column 3. Outcome: binary indicator for a third birth in the 36-month window (columns 1–5, 7) or 48-month window (column 6). Controls included in all specifications: mother's years of education, husband's education level (categorical), urban/rural residence, Scheduled Caste, Scheduled Tribe, Other Backward Class indicators, and media exposure index. The Muslim indicator is included in all columns except (2) and (3), where the sample is restricted to a single religion. Low wealth is defined as below the within-district median wealth score. Column (7) additionally includes fixed effects for years since the second birth (16 categories: 0 to 15+ years). Columns (8) and (9) exclude women who moved to their current place of residence within the last four and ten years, respectively, as recorded by NFHS variable V104. Column (10) adds a district-by-wave modernisation index: the standardised average of the district shares of households with electricity, a television, a refrigerator, a motorcycle/scooter, and a car, each measured identically in both waves. Its level is absorbed by the district fixed effects, so it enters as a within-district modernisation trend. Standard errors clustered by district in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Satellite development proxies. As satellite complements to the in-data asset index, we use two district measures from independent remote-sensing sources: VIIRS night-time luminosity (mean radiance) from the Harmonized Global Night-Time Lights series (Li et al., 2020), and built-up surface area (as a share of district area) from the Global Human Settlement Layer (Pesaresi and Politis, 2023). Both match our analysis districts by name for 94.5% of the estimation sample (the unmatched districts are concentrated among Telangana's post-2016 reorganised districts), and both are assigned to waves by year; with district fixed effects each enters as a within-district development *trend*. The within-district changes in the three development measures—luminosity, built-up area, and the household asset index—are only weakly correlated with one another (pairwise

$|r| \leq 0.28$), so they capture somewhat distinct dimensions of local development. Table 14 reports the results on the common matched sample. Each measure is independently a negative predictor of parity progression—night-time luminosity -0.049 ($p < 0.05$) and built-up area -0.073 ($p < 0.10$)—confirming that faster-developing districts experienced larger fertility declines, exactly as the asset index implies. Adding either attenuates the P_d main effect only modestly (to -0.211 and -0.222, respectively) and leaves the DD $\times$ P_d interaction unchanged ($0.166 \rightarrow 0.165$ and 0.167). When all three proxies enter jointly (column 5), the household asset index dominates and the satellite measures lose significance, indicating that they capture overlapping development variation. The satellite data thus reinforce the in-data result: local development independently lowers fertility, the P_d level is only mildly sensitive to it, and the son-composition interaction is robust throughout.

Table 14: Satellite and Asset-Based Development Controls (Specification B, Parity 2)

	(1) Baseline	(2) + Nightlight	(3) + Built-up	(4) + Assets	(5) + All
P_d	-0.235^{*}	-0.211	-0.222	-0.174	-0.165
	(0.140)	(0.137)	(0.140)	(0.135)	(0.135)
DD $\times$ P_d	0.166	0.165	0.167	0.163	0.164
	(0.143)	(0.142)	(0.142)	(0.142)	(0.142)
Night-time luminosity (ln, std.)		-0.049^{**}			-0.022
		(0.019)			(0.019)
Built-up area (ln, std.)			-0.073^{*}		-0.039
			(0.043)		(0.045)
Asset/infrastructure index				-0.067^{***}	-0.057^{***}
				(0.014)	(0.015)
DD	0.152^{***}	0.152^{***}	0.152^{***}	0.152^{***}	0.152^{***}
	(0.018)	(0.018)	(0.018)	(0.018)	(0.018)
Controls / District / Wave / Age FE	Yes	Yes	Yes	Yes	Yes
Observations	35,193	35,193	35,193	35,193	35,193

Notes: Linear probability models, Specification B (parity-2-to-3 progression), on the districts matched to both satellite panels (94.5% of the estimation sample). "Night-time luminosity" is the standardised log of district mean radiance from the Harmonized Global Night-Time Lights series (Li et al., 2020) (VIIRS), with the 2016 value assigned to wave 4 and 2020 to wave 5. "Built-up area" is the standardised log of the district built-up surface share from the Global Human Settlement Layer (GHS-BUILT-S R2023A) (Pesaresi and Politis, 2023), with the 2015 epoch assigned to wave 4 and the 2020 epoch to wave 5. With district fixed effects, both satellite measures enter as within-district development trends. "Asset/infrastructure index" is the standardised average of district shares of households with electricity, a television, a refrigerator, a motorcycle/scooter, and a car (the measure of Table 13, column 10). All columns include the full Specification B controls and interactions and district, survey-wave, and single-year-age fixed effects. Standard errors clustered by district in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

7.7 Urban versus Rural

A natural heterogeneity prediction follows from the structure of the private market. The gendered cost channel can operate only where private schooling is actually prevalent and where the boy–girl gap in private enrolment is salient—and both conditions are far more strongly met in urban areas. In the NFHS-5 preschool data (ages 3–5, the only level at which government and private attendance are distinguished), the private share among enrollees is roughly 50 percent in urban areas versus 21 percent in rural areas, and the district private-school share P_d averages 0.13 versus 0.10. The boy–girl gap in private enrolment is also wider in urban areas (a male advantage of 4.8 versus 2.3 percentage points among enrollees), though this difference is only marginally significant. Taken together, the far greater development of the private market in urban areas and the somewhat wider gendered enrolment gap lead us to expect the son-composition interaction, DD $\times$ P_d, to be stronger in urban areas.

Table 15 reports Specification B estimated separately for rural and urban households, and the point estimates move in the predicted direction. The DD $\times$ P_d interaction is $+0.518$ ($p < 0.01$) among urban households—the largest and most precisely estimated interaction anywhere in the paper—versus an insignificant $+0.292$ in rural areas. The urban point estimate is thus markedly larger, but a formal test cannot reject equality of the two: the urban–rural difference is $+0.23$ ($p = 0.42$). We therefore read the split as *suggestive* rather than decisive: the interaction is significantly positive among urban households, where the private market is most developed, and indistinguishable from zero among rural ones, even though the two estimates are not themselves statistically distinguishable. Among urban families with only daughters, the implied total effect of P_d is close to zero (the P_d main effect of -0.310 is roughly offset by the interaction), while urban families with sons are deterred—the gender-differentiated pattern the model predicts, sharpest where the private market is most active.

The P_d *main* effect, by contrast, is statistically indistinguishable across the two (urban–rural difference $p = 0.84$)—and, if anything, larger in rural areas in point terms (-0.371, $p < 0.10$, versus an insignificant -0.310 in urban). This is consistent with the affordability logic: the broad deterrent effect of schooling costs need not be concentrated in urban areas, because the *burden* of fees relative to income is heaviest for poorer households, who are disproportionately rural. The two patterns together—a broad deterrent effect that is, if anything, rural, and a gender-differentiated effect concentrated in urban areas—are hard to reconcile with a story based on absolute fee levels, but fit naturally with a cost channel that operates through household budgets and gendered schooling investment.

Table 15: Heterogeneity by Urban/Rural Residence (Specification B, Parity 2)

	(1) Full	(2) Rural	(3) Urban
DD	0.149***	0.174***	0.033
	(0.017)	(0.021)	(0.027)
Mixed	0.028***	0.041***	−0.015
	(0.009)	(0.012)	(0.011)
P_d	−0.264*	−0.371*	−0.310
	(0.140)	(0.206)	(0.198)
DD × P_d	0.177	0.292	0.518***
	(0.138)	(0.201)	(0.193)
Mixed × P_d	−0.073	−0.115	0.122
	(0.063)	(0.096)	(0.079)
Low wealth	0.026***	0.024**	−0.007
	(0.009)	(0.011)	(0.018)
P_d × Low wealth	−0.053	−0.037	0.152
	(0.076)	(0.092)	(0.131)
Test of equality, urban vs rural (p-value):			
DD × P_d		0.42	
P_d		0.84	
Controls / District / Wave / Age FE	Yes	Yes	Yes
Observations	37,251	25,651	11,595

Notes: Linear probability models, Specification B (parity-2-to-3 progression), estimated on the full sample (column 1) and separately for rural (2) and urban (3) households. SS (two sons) is the reference sex-composition category. All columns include mother's years of education, husband's education level, Muslim, Scheduled Caste, Scheduled Tribe, Other Backward Class, and media-exposure controls, and district, survey-wave, and single-year-age fixed effects. Standard errors clustered by district in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

7.8 Mean Reversion and the Continuous-Dose Design

As discussed in Section 5.3, our baseline is a continuous-treatment difference-in-differences design, and two threats specific to that design require attention: differential trends across dose levels (which, given the baseline imbalance in fertility, take the concrete form of a mean-reversion concern), and the possibility that the two-way fixed-effects estimand places pathological weights on dose-specific responses. We address each in turn.

Differential baseline-fertility trends. The balance diagnostics in Table 3 showed that districts with larger private-school expansion begin from higher fertility (children ever born of 2.90 versus 2.59; normalised difference ≈ 0.20, the largest among all covariates). If high-fertility districts were declining faster regardless of private-school growth, the district-level P_d coefficient could reflect regression toward the mean rather than a treatment effect. To assess this, we augment Specification B with an interaction between baseline district fertility—the NFHS-4 district mean of children ever born among ever-married women aged 25–49—and the wave-5 indicator, allowing districts on different baseline-fertility trajectories to trend differentially. Table 16 reports the results. The baseline-fertility trend is itself strongly significant (-0.039, $p < 0.001$), confirming that higher-fertility districts did decline faster on their own. Once this differential trend is absorbed, the P_d level coefficient falls from -0.261 to -0.062 and loses significance, while the DD $\times$ P_d interaction is essentially unchanged ($0.192 \rightarrow 0.191$). A same-sample decomposition (columns 1–3) confirms that this attenuation is driven by the control and not by the modest change in sample: on the identical 35,276-observation sample, P_d is -0.261 without the control and -0.062 with it. We read this as direct evidence that the aggregate *level* effect of P_d is fragile to district-level mean reversion. The within-district interaction, by contrast, differences out any district-level trend—including mean reversion—and is therefore immune to this concern. We emphasize that baseline fertility is itself correlated with the dose, so this is a demanding and arguably over-controlling specification, hence the appropriate conclusion is that the level effect is *not robust*, not that it is zero. This is because private schools enter where the school-age population is larger and where public provision is most strained (Muralidharan and Kremer, 2008). In other words, both conditions are mechanically more binding where fertility is higher, since more children per family imply a larger pupil market and greater pressure on government-school capacity (Section 2). Because higher baseline fertility is itself a determinant of where private schooling expands, allowing high-fertility districts their own differential trend mechanically absorbs part of the schooling effect, not merely a confounding trend.

Dose–response. Continuous-treatment difference-in-differences additionally warns that the two-way fixed-effects estimand is a weighted average of dose-specific causal responses whose implicit weights can be negative, in which case the average would mask offsetting heterogeneity (Callaway et al., 2024). To check that our estimate is not contaminated in this way, we group districts into terciles of the within-district change in private-school share (ΔP_d) and estimate the wave-4-to-wave-5 change in parity-2-to-3 progression separately by tercile. Table 17 shows that the dose–response is monotone and correctly signed: progression falls by 3.1, 3.3, and 3.9 percentage points moving from the lowest to the highest dose tercile, and the implied per-unit slope (≈ -0.23) is consistent with the main P_d coefficient (-0.26). This ordered pattern rules out the negative-weighting pathology. Two caveats temper the magnitude. First, the lowest-dose tercile—districts with essentially no change in private-school share—still exhibits a 3.1-point decline, indicating that most of the progression decline over this period reflects a common secular trend rather than private-school expansion. Second, the incremental decline associated with a higher dose (the T3$-$T1 difference of 0.8 points) is not statistically significant ($p = 0.39$). Together with the mean-reversion result, this reinforces our decision to rest causal inference on the within-district interaction rather than on the aggregate dose–response.

Table 16: Mean Reversion: Robustness to Differential Baseline-Fertility Trends

	(1) Baseline (full)	(2) Baseline (restricted)	(3) + Baseline CEB × Wave 5
P_d	-0.264^{*}	-0.261^{*}	-0.062
	(0.140)	(0.140)	(0.129)
DD × P_d	0.177	0.192	0.191
	(0.138)	(0.141)	(0.141)
Baseline CEB × Wave 5			-0.039^{***}
			(0.007)
Controls	Yes	Yes	Yes
District / Wave / Age FE	Yes	Yes	Yes
Observations	37,251	35,276	35,276

Notes: Linear probability models with the Specification B regressors (sex composition, P_d, DD × P_d, Mixed × P_d, P_d× LowWealth) and the full set of controls; district, survey-wave, and single-year-age fixed effects throughout. Outcome: third birth within the 36-month window. "Baseline CEB" is the NFHS-4 district mean of children ever born among ever-married women aged 25–49, interacted with the wave-5 indicator (its level is absorbed by district fixed effects). Column (1) uses the full estimation sample; columns (2)–(3) use the sample with a non-missing baseline-CEB value, so the difference between them isolates the effect of the control. Standard errors clustered by district in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

Table 17: Dose–Response: Wave-4-to-Wave-5 Change in Parity-2-to-3 Progression by ΔP_d Tercile

	T1 (low)	T2 (mid)	T3 (high)
Mean ΔP_d (pp)	-0.3	$+1.1$	$+3.3$
Change in progression (pp)	-3.07^{***}	-3.30^{***}	-3.90^{***}
	(0.58)	(0.60)	(0.83)
T3 − T1 difference: -0.83 (SE 0.96), $p = 0.39$			

Notes: Districts in the estimation sample are split into terciles by the within-district change in private-school share between waves ($\Delta P_d = P_{d,2019\text{–}20} - P_{d,2015\text{–}16}$). Each "change in progression" entry is the estimated wave-4-to-wave-5 change in the probability of a parity-2-to-3 progression for that tercile, from a single regression of the outcome on a wave-5 indicator interacted with tercile, with district and single-year-age fixed effects and the full set of controls (N = 34,879). Standard errors clustered by district in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

7.9 What Drives the Change in P_d?

Because P_d is a count share, it can rise either when private schools open (the numerator grows) or when public schools close (the denominator shrinks), and several states consolidated their government-school networks over this period. We therefore examine what actually drives the within-district change in P_d and confirm that our estimates reflect private expansion rather than public contraction.[8]

Using the district-year UDISE counts, private-school openings are the dominant driver of the rise in P_d. A ceteris-paribus decomposition attributes roughly 68 percent of the mean within-district increase to growth in the private count, and in a regression of ΔP_d on the underlying count changes, private growth accounts for essentially all of the cross-district variation that identifies our estimates (standardised $\beta = 0.98$, versus -0.41 for the change in the non-private stock). Public consolidation is real but secondary: the (narrowly defined) government count fell in two-thirds of districts, concentrated in the expected states (Assam, Madhya Pradesh, Odisha, Uttar Pradesh, West Bengal), yet this denominator channel carries little of the identifying variation. Decisively, consolidation *alone* does not raise P_d: in districts where the public stock shrank but private schools did not open, P_d if anything fell. Private expansion is the necessary ingredient.

As a direct check, Table 18 re-estimates Specification B replacing P_d with $P_d^{fix} = \text{private}_t/\text{total}_{2016}$, which holds the denominator fixed at its baseline value so that the measure can rise *only* when private schools open—it is mechanically immune to public-school closures. The estimated effect is negative, of essentially the same magnitude as the baseline (-0.32), and in fact more precisely estimated ($p < 0.01$ versus $p < 0.10$): freezing the denominator strips out consolidation-driven noise and sharpens the estimate. The son-composition interaction is unchanged. The fertility effect of P_d thus reflects the expansion of private schooling, not the contraction of the public sector.

[8] The UDISE panel's government category captures Department-of-Education schools; other publicly managed schools (local-body, tribal and social-welfare departments) and madrasa or unrecognised schools fall into a large residual category. We therefore decompose the change into private versus non-private (everything else) rather than relying on a finer government classification.

Table 18: Private-Expansion-Only Treatment: Freezing the Denominator

	(1) Baseline P_d	(2) $P_d^{fix} = \text{Priv}_t/\text{Tot}_{2016}$
DD	0.148***	0.153***
	(0.017)	(0.017)
Mixed	0.030***	0.030***
	(0.009)	(0.009)
Treatment	−0.326*	−0.316***
	(0.183)	(0.119)
DD × Treatment	0.194	0.145
	(0.144)	(0.140)
Controls / District / Wave / Age FE	Yes	Yes
Observations	35,879	35,879

Notes: Linear probability models, Specification B (parity-2-to-3 progression). "Treatment" is the baseline count share P_d in column (1) and $P_d^{fix} = \text{private}_t/\text{total}_{2016}$ in column (2), which freezes the denominator at its 2015–16 value so the measure rises only when private schools open. Both columns include the full Specification B interactions and controls (Mixed × Treatment, low wealth, Treatment × LowWealth, mother's and husband's education, urban, Muslim, caste, media exposure) and district, survey-wave, and age fixed effects; the sample is restricted to districts with a non-missing baseline value ($N = 35{,}879$). Standard errors clustered by district in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

7.10 Supply-Side Variation: A Consistency Check

A central concern with our identification strategy is that district-level private-school prevalence P_d may be endogenous: districts with faster private-school expansion may also have experienced changes in income, urbanisation, or fertility preferences that independently affected parity progression. The district fixed effects in our baseline specification absorb all time-invariant district confounders, and the individual-level controls—mother's education, husband's education, household wealth, urban residence, religion, caste, and media exposure—absorb the most important time-varying channels through which district-level changes could affect individual fertility decisions. The interaction coefficients (DD × P_d) are further protected, as they are identified from within-district cross-individual variation in child gender composition, which is orthogonal to district-level shocks. Nevertheless, the main P_d effect remains potentially exposed to time-varying district-level changes not captured by these individual controls—for instance, district-specific shifts in family-planning programme intensity or local labour market conditions that are correlated with private-school expansion.

As a *supplementary* check on the P_d main effect—the component of our estimates most exposed to this concern—we isolate the variation in P_d that is driven by a predetermined, supply-side determinant of private-school entry: the local stock of potential private-school teachers, following the logic of Andrabi et al. (2013) for Pakistan. Private schools in India rely on locally recruited women with at least secondary education, who accept wages far below government-school levels (Kingdon, 2020; Chudgar and Sakamoto, 2021). Districts with a larger pool of educated women can therefore sustain more private schools—a supply-side constraint that is predetermined with respect to contemporaneous fertility decisions. We re-estimate Specification B by two-stage least squares, using this supply measure to isolate the supply-driven component of P_d. We stress at the outset that this is *not* a clean instrumental-variables design: the exclusion restriction cannot be guaranteed. We therefore present it as a *consistency check*—asking whether our results survive when P_d is restricted to its supply-driven variation, which addresses attenuation and reverse causality—rather than as a source of clean causal identification, and none of our central conclusions rests on it.

Supply-side measure. For each district, we compute the share of women with at least ten years of education (secondary or above) from the NFHS-4 survey (2015–16). To generate time variation, we use different age cohorts for each wave: women aged 30–50 for the NFHS-4 supply measure and women aged 24–44 for the NFHS-5 supply measure. The latter captures the same birth cohorts that will be aged 30–50 at the time of NFHS-5, but measured six years earlier—before any private-school expansion between waves could have attracted educated women to high-P_d districts through endogenous migration. Crucially, *both* values of the supply measure are computed entirely from NFHS-4, ensuring that it is fully predetermined with respect to the 2019–21 schooling environment and fertility outcomes.

Results. Table 19 reports the supply-driven (two-stage least squares) estimates alongside the OLS baseline for comparison. All four columns use the same sample of 35,276 women in districts present in both NFHS waves.[9] The first-stage F-statistic (Kleibergen-Paap *rk* Wald F) is 45.1 with state fixed effects and 9.2 with district fixed effects.

The supply-driven estimates confirm the sign pattern of the OLS results across both fixed-effects structures. The P_d coefficient is negative in all specifications, reaching -0.667 ($p < 0.01$) with state fixed effects and -3.091 ($p = 0.10$) with district fixed effects, though this last magnitude is unreliable given the weak first stage in that specification. The DD $\times$ P_d interaction—the model's central prediction—is positive and statistically significant in both 2SLS specifications: $+0.587$ ($p < 0.01$) with state fixed effects and $+0.419$ ($p = 0.04$) with district fixed effects. The Anderson-Rubin weak-identification-

[9] The sample is slightly smaller than the main analysis (N = 37,251) because the supply measure requires NFHS-4 coverage: districts appearing only in NFHS-5 are dropped. In the main district-FE specification, these single-wave districts are included in the observation count but contribute nothing to the identification of the P_d coefficient, since there is no within-district variation for them.

robust test rejects the null of no P_d effect at $p = 0.049$ in the district-FE specification, providing inference that is valid even if the first stage is borderline weak.

We interpret these results as supportive of the main findings. The supply-driven estimates are larger in magnitude than OLS, consistent with attenuation from measurement error in P_d or from reverse causality (high-fertility districts attracting more schools). The sign and significance of the DD $\times$ P_d interaction—positive in all four columns and significant in the 2SLS specifications—reinforces the model's prediction that families with only daughters, whose existing children impose lower schooling costs, are less deterred by private-school prevalence.

Finally, although the supply-driven estimates are reassuring, we retain the OLS specification with district fixed effects as our preferred approach for three reasons. First, the district fixed effects absorb all time-invariant district confounders — including persistent differences in culture, caste composition, urbanisation, and historical development — providing a more conservative identification strategy than the state fixed effects required by the supply-side check. Second, the paper's most distinctive predictions concern the interaction coefficients (DD $\times$ P_d, P_d $\times$ LowWealth), which are identified from cross-individual variation in child gender composition and household wealth within the same district and wave. These interaction coefficients do not require P_d itself to be exogenous: they require only that the differential response to P_d across family types is not confounded by unobserved district-level factors, a substantially weaker condition. Third, while the supply-side first stage is strong with state fixed effects (F = 45), it is borderline with district fixed effects (F = 9.2), raising concerns about weak-identification bias if this design were adopted as the main specification. We therefore present the supply-side analysis as a consistency check that speaks to attenuation and reverse causality in the P_d main effect, while emphasising that the interaction-based evidence — which is the core of our empirical contribution — does not depend on it.

Table 19: Supply-Side Variation: Two-Stage Least Squares Using Predetermined Teacher Supply

	(1) OLS District FE	(2) OLS State FE	(3) 2SLS State FE	(4) 2SLS District FE
Sex composition (ref: SS)				
DD	0.147***	0.149***	0.104***	0.123***
	(0.017)	(0.017)	(0.023)	(0.024)
Mixed	0.029***	0.028***	−0.020	0.002
	(0.009)	(0.009)	(0.019)	(0.020)
Private-school prevalence				
P_d	−0.261*	0.008	−0.667***	−3.091*
	(0.141)	(0.071)	(0.240)	(1.864)
Private-school × sex composition				
DD × P_d	0.192	0.180	0.587***	0.419**
	(0.141)	(0.142)	(0.200)	(0.206)
Mixed × P_d	−0.078	−0.072	0.363**	0.166
	(0.064)	(0.064)	(0.164)	(0.170)
Liquidity channel				
Low wealth	0.026***	0.021**	−0.001	0.012
	(0.010)	(0.009)	(0.013)	(0.014)
P_d × Low wealth	−0.058	−0.042	0.185	0.075
	(0.078)	(0.076)	(0.114)	(0.120)
Controls	Yes	Yes	Yes	Yes
District FE	Yes	No	No	Yes
State FE	No	Yes	Yes	No
Wave FE	Yes	Yes	Yes	Yes
Age FE	Yes	Yes	Yes	Yes
First-stage F	—	—	45.14	9.18
AR test p-value	—	—	<0.001	0.049
Observations	35,276	35,276	35,276	35,276

Notes: Linear probability models. Outcome: 3rd birth in 36-month window. Columns (3)–(4) report two-stage least squares estimates that isolate the supply-driven component of P_d, using a predetermined supply-side measure: the district-level share of ever-married women with at least secondary education (10+ years of schooling), computed from NFHS-4 (2015–16). For wave 4, women aged 30–50 are used; for wave 5, women aged 24–44 are used (the same birth cohorts measured six years earlier, before any between-wave private-school expansion). Both values of the supply measure are computed entirely from the NFHS-4 survey. P_d is projected onto the supply measure; the interactions use actual P_d. The first-stage F-statistic is the Kleibergen-Paap rk Wald F from `ivreg2`. The Anderson-Rubin (AR) test is a weak-identification-robust test of the null that P_d has no effect. As discussed in the text, the exclusion restriction cannot be guaranteed, so these estimates are presented as a supply-side consistency check rather than as clean causal identification. Controls: mother's education, husband's education, urban, Muslim, SC, ST, OBC, media exposure. Standard errors clustered by district. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.10$.

8 Conclusion

India reached below-replacement fertility under conditions that do not fit the standard demographic-transition narrative. This paper proposes and tests a complementary mechanism: the expansion of private schooling made higher-quality, higher-cost education widely available, increasing households' exposure to the cost of child quality and thereby reshaping fertility decisions.

We study this mechanism in a setting with vertical differentiation and persistent son preference. In our framework, the expansion of fee-based schooling raises the effective marginal cost of children for households investing in quality, with heterogeneous effects because boys are more likely than girls to be enrolled in private schools. Consistent with the model, we find that higher private-school prevalence is associated with lower probabilities of an additional birth at higher parities, and that this effect varies systematically with child gender composition. The negative association is concentrated at later fertility margins, stronger among Hindu households, and suggestively larger for lower-income families, although the latter estimates are imprecise.

These findings point to a broader insight: the expansion of a fragmented service market can generate equilibrium spillovers outside the market itself. When a higher-quality, higher-cost option becomes widely available, households adjust not only their investment per child but also their fertility choices. In this sense, market expansion alters the effective budget constraint governing household decisions, even for those who do not fully participate in the higher-quality segment.

Our results also speak to the conditions under which the quality–quantity tradeoff becomes binding. In the Indian context, the key shift was not primarily rising income or female labour supply, but the increased availability of costly quality through private schooling. This helps explain how fertility can decline in settings where standard income-based mechanisms are weak.

We interpret these results as evidence consistent with a cost-based channel operating through the expansion of private schooling, while acknowledging that our empirical strategy is based on a reduced-form design and cannot fully rule out all sources of time-varying district-level confounding. The strongest support for the mechanism comes from the interaction between private-school prevalence and child gender composition, which aligns closely with the model's predictions.

More broadly, the paper highlights how changes in market structure in essential service sectors can shape household behavior in ways that extend beyond the focal market. In India, the expansion of private schooling did not affect only education choices; it also altered the economic tradeoffs underlying fertility decisions by making costly quality salient at scale.

References

Andrabi, Tahir, Jishnu Das, and Asim Ijaz Khwaja, "A Dime a Day: The Possibilities and Limits of Private Schooling in Pakistan," *Comparative Education Review*, 2008, *52* (3), 329–355. `https://doi.org/10.1086/588796`.

_ , _ , and _ , "Students today, teachers tomorrow: Identifying constraints on the provision of education," *Journal of Public Economics*, 2013, *100*, 1–14. `https://doi.org/10.1016/j.jpubeco.2012.12.003`.

_ , _ , _ , Selcuk Ozyurt, and Niharika Singh, "Upping the Ante: The Equilibrium Effects of Unconditional Grants to Private Schools," *American Economic Review*, 2020, *110* (10), 3315–3349. `https://doi.org/10.1257/aer.20180924`.

_ , Natalie Bau, Jishnu Das, and Asim Ijaz Khwaja, "Heterogeneity in School Value Added and the Private Premium," *American Economic Review*, 2025, *115* (1), 147–182. `https://doi.org/10.1257/aer.20221422`.

_ , _ , _ , Naureen Karachiwalla, and Asim Ijaz Khwaja, "Crowding in Private Quality: The Equilibrium Effects of Public Spending in Education," *Quarterly Journal of Economics*, 2024, *139* (4), 2525–2577. `https://doi.org/10.1093/qje/qjae014`.

Angrist, Joshua D, Victor Lavy, and Analia Schlosser, "Multiple experiments for the causal link between the quantity and quality of children," *Journal of Labor Economics*, 2010, *28* (4), 773–824. `https://doi.org/10.1086/653830`.

Anukriti, S, "Financial incentives and the fertility-sex ratio trade-off," *American Economic Journal: Applied Economics*, 2018, *10* (2), 27–57. `https://doi.org/10.1257/app.20150234`.

Azam, Mehtabul and Geeta Gandhi Kingdon, "Are Girls the Fairer Sex in India? Revisiting Intra-Household Allocation of Education Expenditure," *World Development*, 2013, *42*, 143–164. `https://doi.org/10.1016/j.worlddev.2012.09.003`.

_ , _ , and Kin Bing Wu, "Impact of Private Secondary Schooling on Cognitive Skills: Evidence from India," *Education Economics*, 2016, *24* (5), 465–480. `https://doi.org/10.1080/09645292.2015.1110116`.

Bagde, Surendrakumar, Dennis Epple, and Lowell Taylor, "The emergence of private high schools in India: The impact of public-private competition on public school students," *Journal of Public Economics*, 2022, *215*, 104749. `https://doi.org/10.1016/j.jpubeco.2022.104749`.

Becker, Gary S, "An economic analysis of fertility," *Demographic and Economic Change in Developed Countries*, 1960, pp. 209–240. https://www.nber.org/chapters/c2387.

— **and H Gregg Lewis**, "On the interaction between the quantity and quality of children," *Journal of Political Economy*, 1973, *81* (2), S279–S288. https://doi.org/10.1086/260166.

— **, Kevin M Murphy, and Robert Tamura**, "Human capital, fertility, and economic growth," *Journal of Political Economy*, 1990, *98* (5, Part 2), S12–S37. https://doi.org/10.1086/261723.

Bhalotra, Sonia and Damian Clarke, "The twin instrument: Fertility and human capital investment," *Journal of the European Economic Association*, 2020, *18* (6), 3090–3139. https://doi.org/10.1093/jeea/jvz058.

Black, Sandra E, Paul J Devereux, and Kjell G Salvanes, "The more the merrier? The effect of family size and birth order on children's education," *Quarterly Journal of Economics*, 2005, *120* (2), 669–700. https://doi.org/10.1093/qje/120.2.669.

Burlando, Alfredo and Edward Bbaale, "Fertility responses to schooling costs: Evidence from Uganda's universal primary education policy," *Economic Development and Cultural Change*, 2022, *70* (3), 1017–1039. https://doi.org/10.1086/713938.

Callaway, Brantly, Andrew Goodman-Bacon, and Pedro H. C. Sant'Anna, "Difference-in-Differences with a Continuous Treatment," NBER Working Paper 32117, National Bureau of Economic Research 2024. https://www.nber.org/papers/w32117.

Central Square Foundation, "State of the Sector Report on Private Schools in India," Technical Report, Central Square Foundation and Omidyar Network India, New Delhi 2020. https://www.centralsquarefoundation.org/reports/state-of-the-sector-report-private-schools-in-india.

Centre for Civil Society, "India School Closure Report," Technical Report, Centre for Civil Society, New Delhi 2018. Based on Right to Information requests and media reports, covering April 2015 to March 2018 https://ccs.in/sites/default/files/2022-08/school-closure-report.pdf.

Chatterjee, Shoumitro and Tom Vogl, "Escaping Malthus: Economic growth and fertility change in the developing world," *American Economic Review*, 2018, *108* (6), 1440–1467. https://doi.org/10.1257/aer.20170748.

Chudgar, Amita and Elizabeth Quin, "Relationship between Private Schooling and Achievement: Results from Rural and Urban India," *Economics of Education Review*, 2012, *31* (4), 376–390. `https://doi.org/10.1016/j.econedurev.2011.12.003`.

_ **and Jutaro Sakamoto**, "Similar work, different pay? Private school teacher working conditions in India," *International Journal of Educational Development*, 2021, *86*, 102478. `https://doi.org/10.1016/j.ijedudev.2021.102478`.

Datta, Sandip and Geeta Gandhi Kingdon, "Gender Bias in Intra-Household Allocation of Education in India: Has It Fallen over Time?," IZA Discussion Paper 12671, IZA – Institute of Labor Economics, Bonn, Germany 2019. `https://www.iza.org/publications/dp/12671/gender-bias-in-intra-household-allocation-of-education-in-india-has-it-fallen-over`

de Chaisemartin, Clément and Xavier D'Haultfœuille, "Two-way fixed effects estimators with heterogeneous treatment effects," *American Economic Review*, 2020, *110* (9), 2964–2996. `https://doi.org/10.1257/aer.20181169`.

Department of School Education and Literacy, "Unified District Information System for Education Plus (UDISE+) 2021–22: Flash Statistics," Technical Report, Ministry of Education, Government of India, New Delhi 2022. `https://www.education.gov.in/udise-report-2021-22`.

Dizon-Ross, Rebecca, "Parents' Beliefs about Their Children's Academic Ability: Implications for Educational Investments," *American Economic Review*, 2019, *109* (8), 2728–2765. `https://doi.org/10.1257/aer.20171172`.

Drèze, Jean and Mamta Murthi, "Fertility, education, and development: Evidence from India," *Population and Development Review*, 2001, *27* (1), 33–63. `https://doi.org/10.1111/j.1728-4457.2001.00033.x`.

Duflo, Esther, Pascaline Dupas, and Michael Kremer, "Education, HIV, and early fertility: Experimental evidence from Kenya," *American Economic Review*, 2015, *105* (9), 2757–2797. `https://doi.org/10.1257/aer.20121607`.

Endow, Tanuka, "Low Cost Private Schools: How Low Cost Really Are These?," *Indian Journal of Human Development*, 2019, *13* (1), 102–108. `https://doi.org/10.1177/0973703019838704`.

FSG, "Understanding the Affordable Private School Market in India," Technical Report 2013. `https://www.fsg.org/wp-content/uploads/2021/10/Understanding-the-Affordable-Private-School-Market-in-India_1013.pdf`.

Goodman-Bacon, Andrew, "Difference-in-differences with variation in treatment timing," *Journal of Econometrics*, 2021, *225* (2), 254–277. https://doi.org/10.1016/j.jeconom.2021.03.014.

Imbens, Guido W and Jeffrey M Wooldridge, "Recent developments in the econometrics of program evaluation," *Journal of Economic Literature*, 2009, *47* (1), 5–86. https://doi.org/10.1257/jel.47.1.5.

Jayachandran, Seema, "Fertility decline and missing women," *American Economic Journal: Applied Economics*, 2017, *9* (1), 118–139. https://doi.org/10.1257/app.20150576.

_ , "Ten facts about son preference in India," 2023. NBER Working Paper No. 31883 https://www.nber.org/papers/w31883.

_ **and Ilyana Kuziemko**, "Why do mothers breastfeed girls less than boys? Evidence and implications for child health in India," *Quarterly Journal of Economics*, 2011, *126* (3), 1485–1538. https://doi.org/10.1093/qje/qjr029.

_ **and Rohini Pande**, "Why are Indian children so short? The role of birth order and son preference," *American Economic Review*, 2017, *107* (9), 2600–2629. https://doi.org/10.1257/aer.20151282.

Kaul, Tara, "Intra-household allocation of educational expenses: Gender discrimination and investing in the future," *World Development*, 2018, *104*, 336–343. https://doi.org/10.1016/j.worlddev.2017.12.017.

Kingdon, Geeta Gandhi, "The Gender Gap in Educational Attainment in India: How Much Can Be Explained?," *Journal of Development Studies*, 2002, *39* (2), 25–53. https://doi.org/10.1080/00220380412331322741.

_ , "The private schooling phenomenon in India: A review," *Journal of Development Studies*, 2020, *56* (10), 1795–1817. https://doi.org/10.1080/00220388.2020.1715943.

_ **and Mohd Muzammil**, "Government Per Pupil Expenditure in Uttar Pradesh: Implications for the Reimbursement of Private Schools under the RTE Act," CSAE Working Paper 2015-18, Centre for the Study of African Economies, University of Oxford 2015. https://ideas.repec.org/p/csa/wpaper/2015-18-2.html.

_ **and Sarmistha Pal**, "Can Private School Growth Foster Universal Literacy? Panel Evidence from Indian Districts," IZA Discussion Paper 5274, IZA – Institute of Labor Economics, Bonn, Germany 2010. https://www.iza.org/publications/dp/5274/can-private-school-growth-foster-universal-literacy-panel-evidence-from-indian-dist

Kremer, Michael, Nazmul Chaudhury, F Halsey Rogers, Karthik Muralidharan, and Jeffrey Hammer, "Teacher absence in India: A snapshot," *Journal of the European Economic Association*, 2005, *3* (2-3), 658–667. https://doi.org/10.1162/jeea.2005.3.2-3.658.

Kugler, Adriana D and Santosh Kumar, "Preference for boys, family size, and educational attainment in India," *Demography*, 2017, *54* (3), 835–859. https://doi.org/10.1007/s13524-017-0575-1.

Li, Xuecao, Yuyu Zhou, Min Zhao, and Xia Zhao, "A harmonized global nighttime light dataset 1992–2018," *Scientific Data*, 2020, *7* (1), 168. DMSP-OLS and VIIRS harmonized series; https://doi.org/10.1038/s41597-020-0510-y.

Mahler, Lukas, Michèle Tertilt, and Minchul Yum, "Policy concerns in an era of low fertility: The role of social comparisons and intensive parenting," *Brookings Papers on Economic Activity*, 2025. Conference Draft, Fall 2025 https://www.brookings.edu/articles/policy-concerns-in-an-era-of-low-fertility/.

Maitra, Pushkar, Sarmistha Pal, and Anurag Sharma, "Absence of Altruism? Female Disadvantage in Private School Enrollment in India," *World Development*, 2016, *85*, 105–125. https://doi.org/10.1016/j.worlddev.2016.04.005.

Meng, Juanjuan, Hui Wang, Yu Alan Yang, and Mingshan Zhang, "Education Competition and Fertility Intention: Evidence from China's Private Tutoring Ban," *Economic Journal*, 2025, *136* (676), 1504–1526. https://doi.org/10.1093/ej/ueaf096.

Munshi, Kaivan and Mark Rosenzweig, "Traditional Institutions Meet the Modern World: Caste, Gender, and Schooling Choice in a Globalizing Economy," *American Economic Review*, 2006, *96* (4), 1225–1252. https://www.aeaweb.org/articles?id=10.1257/aer.96.4.1225.

Muralidharan, Karthik and Michael Kremer, "Public and private schools in rural India," in Rajashri Chakrabarti and Paul E Peterson, eds., *School Choice International: Exploring Public–Private Partnerships*, MIT Press, 2008, pp. 91–110. https://doi.org/10.7551/mitpress/9780262033763.003.0005.

_ **and Venkatesh Sundararaman**, "The aggregate effect of school choice: Evidence from a two-stage experiment in India," *Quarterly Journal of Economics*, 2015, *130* (3), 1011–1066. https://doi.org/10.1093/qje/qjv013.

Nandi, Arindam, Soham Sahoo, Nicole Haberland, and Thoai D. Ngo, "A glass ceiling at the playhouse? Gender gaps in public and private preschool enrollment in

India," *International Journal of Educational Development*, 2023, *103*, 102905. https://doi.org/10.1016/j.ijedudev.2023.102905.

National Statistical Office, "Household Social Consumption on Education in India: NSS 75th Round (July 2017–June 2018)," Technical Report, Ministry of Statistics and Programme Implementation, Government of India, New Delhi 2019. https://mospi.gov.in/sites/default/files/publication_reports/Report_585_75th_round_Education_final_1507_0.pdf.

_ , "Comprehensive Modular Survey: Education, 2025 (April–June 2025)," Technical Report, Ministry of Statistics and Programme Implementation, Government of India, New Delhi 2025. NSS 80th Round https://www.pib.gov.in/PressReleasePage.aspx?PRID=2160863.

Oxfam India, "Regulation of Private Schools in India," Technical Report, Oxfam India, New Delhi 2020. Comprehensive review of state-level fee regulation regimes https://www.oxfamindia.org/knowledgehub/workingpaper/regulation-private-schools-india-analysis-legal-frameworks.

Park, Narae, Sangita Vyas, Kathleen Broussard, and Dean Spears, "Near-universal marriage, early childbearing, and low fertility: India's alternative fertility transition," *Demographic Research*, 2023, *48* (34), 945–956. https://doi.org/10.4054/demres.2023.48.34.

Pesaresi, Martino and Panagiotis Politis, "GHS-BUILT-S R2023A: GHS Built-Up Surface Grid," European Commission, Joint Research Centre (JRC), Global Human Settlement Layer 2023. https://doi.org/10.2905/9F06F36F-4B11-47EC-ABB0-4F8B7B1D72EA.

PROBE Team, *Public Report on Basic Education in India*, New Delhi: Oxford University Press, 1999. https://searchworks.stanford.edu/view/4115275.

Reddy, Bheemeshwar A., Prakash Kumar Shukla, Mayank Agrawal, and Shamsher Singh, "Caste, Class, City and Choice: Unravelling Caste-Based School Segregation in Urban India," *Urban Studies*, 2026. https://doi.org/10.1177/00420980251407970.

Rosenzweig, Mark R and Kenneth I Wolpin, "Testing the quantity-quality fertility model: The use of twins as a natural experiment," *Econometrica*, 1980, *48* (1), 227–240. https://doi.org/10.2307/1912026.

Sahai, Harshil, "The Welfare Implications of School Voucher Design: Evidence from India," Technical Report, University of Chicago 2023. Job Market Paper https://socialsciences.uchicago.edu/sites/default/files/2023-11/JMP_Sahai_1.pdf.

Shah, Manisha and Bryce Millett Steinberg, "The Right to Education Act: Trends in enrollment, test scores, and school quality," *AEA Papers and Proceedings*, 2019, *109*, 232–238. https://doi.org/10.1257/pandp.20191060.

Singh, Abhijeet, "Private School Effects in Urban and Rural India: Panel Estimates at Primary and Secondary School Ages," *Journal of Development Economics*, 2015, *113*, 16–32. https://doi.org/10.1016/j.jdeveco.2014.10.004.

Srivastava, Prachi, "Low-Fee Private Schooling: Issues and Evidence," in Prachi Srivastava, ed., *Low-Fee Private Schooling: Aggravating Equity or Mitigating Disadvantage?*, Oxford: Symposium Books, 2013, pp. 7–35. https://www.symposium-books.co.uk/bookdetails/68/.

_ **and Claire Noronha**, "The Myth of Free and Barrier-Free Access: India's Right to Education Act—Private Schooling Costs and Household Experiences," *Oxford Review of Education*, 2016, *42* (5), 561–578. https://doi.org/10.1080/03054985.2016.1220087.

The Economist, "India's surprise baby bust is a warning to the world," *The Economist*, Leaders section, 4 June 2026. https://www.economist.com/leaders/2026/06/04/indias-surprise-baby-bust-is-a-warning-to-the-world.

Tooley, James, *The Beautiful Tree: A Personal Journey into How the World's Poorest People Are Educating Themselves*, Washington, DC: Cato Institute, 2009. https://www.cato.org/books/beautiful-tree.

Vinod, Athira, "Affirmative action and private education expenditure by disadvantaged groups: Evidence from India," Technical Report, University of Nottingham 2024. CREDIT Discussion Paper No. 2024-02 https://www.nottingham.ac.uk/credit/documents/papers/2024/2402.pdf.

World Bank, "World Development Indicators," https://data.worldbank.org/indicator/SL.TLF.CACT.FE.ZS?locations=IN 2024. Indicator: Labor force participation rate, female (% of female population ages 15+), modeled ILO estimate; accessed April 2026.