When Religious Norms Become Law: The Cognitive Costs of Beef Bans in India

Farhan Majid [*†] Wafa Hakim Orman[†‡] Aparajita Dasgupta [§]

September 1, 2026

Abstract

Culturally rooted food regulation can shape human capital. In India, where cows are sacred in Hinduism, beef sale and possession bans restrict a nutrient-rich food consumed disproportionately by beef-consuming communities. Combining legal data, test scores, and Census demographics, we estimate a Callaway–Sant'Anna triple-difference event study exploiting timing, cohorts, and within-state exposure. Exposure around birth lowers mathematics and reading scores at ages 5–12 by 0.075 and 0.065 SD; pre-treatment placebos are null. Mathematics losses are larger when mothers never attended school. We do not observe an increase in dropout rates or school quality; weaker tutoring responses may reinforce the learning disadvantage from early-life beef restrictions.

1 Introduction

Religious beliefs and cultural norms shape economic life in profound ways, from daily consumption habits to national policies (Weber, 1904/05; Iannaccone, 1998; Kuran, 2004; Becker et al., 2024). When states legislate religious norms—whether to promote symbolic identity or enforce moral codes—they may generate unintended economic consequences. One common but understudied domain is diet: governments can regulate the production, sale, or consumption of particular foods in ways that reflect prevailing religious or cultural norms. Yet little is known about whether such policies have persistent consequences for human development. This paper asks whether religiously rooted dietary restrictions, when codified into state policy, can influence human-capital formation years later.

Beef regulation in India provides a particularly revealing setting. Restrictions on the sale and possession of beef apply broadly, but their material consequences depend on pre-existing dietary practices. Beef consumption is disproportionately concentrated among Muslims, Christians, and Scheduled Caste Hindus (Ferry, 2020). Districts with larger shares of these communities are therefore more likely to contain households for whom restrictions on beef meaningfully constrain the diet. This matters nutritionally because beef is a relatively inexpensive source of animal protein and heme iron, which is more readily absorbed than iron from plant sources. Thus, the same state-level restriction can generate very different nutritional shocks across places depending on the prevalence of likely beef consumers before the policy.

[*]Department of Accounting, Economics, and Finance, University of Alabama in Huntsville; Population Studies Center, University of Pennsylvania

[†]Majid and Orman contributed equally and share first authorship.

[‡]Department of Accounting, Economics, and Finance, University of Alabama in Huntsville

[§]Department of Economics, Ashoka University

We ask whether exposure to these restrictions in the birth year affects children's mathematics and reading performance several years later. Learning is an especially consequential margin in India: despite enrollment rates above 95 percent among children ages 6–14, nearly half of rural grade-5 students cannot read a grade-2 text and fewer than one-third can perform basic division (Muralidharan and Singh, 2021). The setting also provides an unusually specific window into the developmental consequences of early-life nutrition. Much of the existing evidence comes from shocks such as famine, fasting, income changes, or broad nutritional interventions that alter several dimensions of nutritional intake simultaneously, making it difficult to distinguish which features of the diet matter for later human capital. Beef restrictions instead change access to a particular nutrient-rich food, one that supplies heme iron as well as protein, vitamin B12, and other micronutrients. The design does not isolate the effect of any single nutrient, but it allows us to move inside the nutritional "black box" by studying a policy-induced shock to the composition of the early-life diet and its consequences for later learning.

To answer this question, we combine an original historical database of state restrictions on the sale and possession of beef with nearly two million child observations from the Annual Status of Education Report (ASER), district-level religion and caste composition from the 2011 Population Census and the 2011 Socio Economic and Caste Census (SECC), and ASER's separate village-level survey of government-school inputs. The legal data were assembled from state gazettes, legislation in Hindi and English, and federal sources (Dasgupta et al., 2023) and record the timing of changes in beef sale and possession restrictions across Indian states. Linking this legal history to children's birth cohorts allows us to measure exposure to these restrictions around birth and then observe mathematics and reading performance, schooling, and household educational investments several years later.

Our identification strategy exploits three dimensions of variation: state-level policy timing, children's birth cohorts, and within-state differences in the expected impact of the restrictions. We classify districts as HighBeef when their share of likely beef consumers exceeds the unweighted district mean within the state, and as LowBeef otherwise. Using never-treated states as the comparison group, we estimate cohort-specific policy effects separately in HighBeef and LowBeef districts and difference the two to obtain a triple-difference (DDD) estimate. Always-treated states are excluded. The baseline analysis focuses on children ages 5–12 and a common event window from two cohorts before through two cohorts after the first treated birth cohort, allowing us to compare the main post-ban differential with directly corresponding pre-treatment placebo estimates.

The results show economically meaningful losses in later learning. The main DDD estimate is -0.075 standard deviations for mathematics and -0.065 for reading, with both 95 percent wild-bootstrap confidence intervals excluding zero. The corresponding pre-treatment placebo estimates are -0.006 and 0.014 and are statistically insignificant. Event-time estimates show little differential movement for the two pre-treatment cohorts or the first treated cohort, followed by larger negative differentials at $T+1$ and $T+2$. The conclusions remain similar when we alter the exposure cutoff, restrict attention to Jharkhand, use all available event horizons, or retain ASER's full age range of 5–16. The mathematics loss is also substantially larger among children whose mothers never attended school, while the corresponding difference for reading is smaller and statistically insignificant.

An early-life nutritional pathway provides a natural interpretation. Beef is a rich source of highly bioavailable heme iron, whereas iron from predominantly cereal- and pulse-based diets is less readily absorbed in the presence of common dietary inhibitors (Hallberg and Hulthén, 2000; Hurrell and Egli, 2010). Iron availability is particularly important during fetal and early-childhood brain development (Georgieff, 2008; Cusick and Georgieff, 2016). Independent biomarker evidence from the same institutional setting supports the relevance of this channel: Dasgupta et al. (2023) show that beef restrictions lower women's hemoglobin and increase anemia among communities for whom beef is a customary food. The magnitudes we estimate are also

comparable to learning and cognitive effects documented for other prenatal nutritional shocks. We do not observe maternal biomarkers for the ASER children, however, and therefore treat this evidence as establishing quantitative and biological plausibility rather than mediation.

The auxiliary outcomes help distinguish the initial early-life shock from later responses to it. Because ASER tests children in their homes regardless of enrollment status, children do not mechanically disappear from the achievement sample when they leave school. We find no increase in dropout or never enrollment among more exposed children and no evidence of a post-ban deterioration in observed school inputs capable of accounting for the learning losses. At the same time, household educational investment responds differentially. Paid private tutoring is common in the sample, but following the bans its participation rate increases 1.67 percentage points less in HighBeef than in LowBeef districts, and private-tutoring expenditure displays a similar relative shortfall. These patterns are consistent with later household investments reinforcing rather than compensating for the initial learning disadvantage, although they do not identify tutoring as the causal channel.

Our first contribution is to show how cultural and religious norms can shape human capital through cross-group transmission via the state. Existing work shows that religious practices and culturally persistent food preferences can affect consumption, nutrition, and economic outcomes (Bell, 1968; Atkin, 2013, 2016; Campante and Yanagizawa-Drott, 2015; Becker et al., 2024). Most closely, Almond et al. (2015) show that prenatal Ramadan exposure lowers children's later academic achievement. We study a distinct mechanism: a culturally rooted norm is embodied in a broadly applicable state restriction, so its consequences can extend beyond those who share or practice the underlying norm.[1] Because beef consumption differed sharply across communities before the ban, the same legal restriction generates very different nutritional shocks across groups. We show that these cross-group effects can persist well beyond the initial dietary shock, becoming visible years later in children's learning.

Our second contribution is to provide, to our knowledge, the first causal evidence that a policy-induced restriction on access to beef around birth affects children's mathematics and reading skills years later. A large developmental-origins literature shows that early-life nutrition can have persistent effects on human capital (see, e.g., Almond et al., 2018; Majid and Behrman, 2021, for reviews). Yet much of the causal evidence on later cognitive and academic skills comes either from interventions that change bundles of nutrients or from broader shocks that alter several dimensions of nutrition and the early-life environment—for example, nutritional supplementation, prenatal fasting, and weather shocks (e.g., Maluccio et al., 2009; Almond et al., 2015; Majid, 2015; Shah and Steinberg, 2017; Webb, 2024). These studies establish the importance of early-life nutrition, but leave open the long-run consequences of policies that change access to a particular food within the diet. The closest causal evidence on beef comes from a small cluster-randomized school-feeding experiment involving 12 primary schools and 555 children in Kenya, which studies beef supplementation among children already in school and contemporaneous cognitive and academic outcomes (Whaley et al., 2003; Neumann et al., 2007). We instead study whether restricting access to beef around birth—a nutrient-rich animal-source food providing highly bioavailable heme iron, protein, vitamin B12, zinc, and other micronutrients—has persistent effects on learning years later.

Our third contribution is to the economics of private tutoring and household educational investment. Private tutoring—or "shadow education"—is a pervasive and socially salient feature of education across much of

[1] Religiously rooted norms have entered economic regulation in other settings. U.S. "blue laws," for example, prohibited Sunday retail activity and have been used to study how Sabbath-based regulation affects economic and religious behavior (Gruber and Hungerman, 2008). A food-specific example is Israel's *Festival of Matzot (Prohibition of Leaven) Law* of 1986, which restricts the public display of leavened products for sale or consumption during Passover in predominantly Jewish localities. These examples differ in institutional scope from the Indian setting, but illustrate the broader incorporation of culturally salient norms into regulation of economic activity.

Asia, absorbing substantial household resources and reflecting intense competition over achievement and social mobility.[2] Existing work studies tutoring demand and shows that supplementary instruction can generate sizeable learning gains (Banerjee et al., 2007; Yuan and Zhang, 2015; Muralidharan et al., 2019; Guo and Qu, 2022; Burchardi et al., 2024; Meng et al., 2026). We add an early-life dimension in two respects: tutoring responds to a developmental shock experienced around birth and is already widespread well before the high-stakes secondary-school and college-admissions stage. Among children ages 5–12 in our sample, 27 percent receive paid tutoring, including 24 percent of those whose mothers never attended school. Following the beef restrictions, tutoring participation and expenditure increase less in more-exposed districts. Thus, an important investment in children's skills responds to shocks experienced years earlier and reinforces, rather than offsets, the emerging learning disadvantage.

2 Institutional Context and Conceptual Framework

2.1 Beef restrictions in India

Cattle regulation in India is primarily determined at the state level and has long varied substantially across states. Many states imposed restrictions on cow or cattle slaughter decades before the period studied here. A smaller set subsequently adopted additional restrictions criminalizing the sale or possession of beef. The distinction is important for our design: the treatment studied in this paper is not the original introduction of cow-slaughter prohibitions, but a later tightening of an already restrictive legal environment.

Two states that already restricted cattle slaughter introduced the sale-or-possession restrictions studied here: Jharkhand in 2005 and Gujarat in 2011. States in which these restrictions were already in force before the relevant cohort window are therefore always treated in our design, while states that do not adopt them during the analysis window provide the comparison group. The construction of the legal database and the longer history of these laws are described in Dasgupta et al. (2023).

The incidence of these restrictions is unlikely to be uniform across the population. Beef consumption in India is concentrated disproportionately among Muslims, Christians, and Scheduled Caste Hindus, for whom beef is a culturally accepted and relatively inexpensive source of animal-source nutrients (Ferry, 2020). Other communities consume substantially less beef for religious and cultural reasons. Accordingly, the same state-level restriction can have very different nutritional relevance across locations depending on local population composition. This motivates the within-state exposure dimension of our empirical design: districts with relatively larger populations belonging to groups for whom beef consumption is customary should be more affected by restrictions on its availability.

2.2 Conceptual framework

Our central hypothesis is that restricting access to beef around birth can affect later human capital through early-life nutrition. Exposure around birth encompasses maternal exposure during pregnancy and the child's early postnatal environment; the birth-cohort design does not separately identify prenatal and infancy exposure. Beef is a particularly rich source of heme iron, which is absorbed more readily than the non-heme iron found in plant foods. This distinction is potentially important in India, where diets rely heavily on

[2]Its prominence extends into popular culture: India's *Kota Factory* and *All India Rank* center on competitive coaching for elite-university entrance, while South Korea's *SKY Castle* depicts private education and intense parental pressure. Supplementary instruction is also global: Kumon, founded in Japan, opened its first overseas mathematics center in New York in 1974 and subsequently expanded internationally.

cereals and pulses and where phytates and other dietary components can inhibit non-heme iron absorption (Hallberg and Hulthén, 2000; Hurrell and Egli, 2010).

Iron availability is especially important during fetal and early-childhood development. Iron supports myelination, hippocampal development, neurotransmitter synthesis, and cellular energy metabolism, and deficiencies during sensitive developmental periods may have persistent consequences for cognition and skill formation (Georgieff, 2008; Cusick and Georgieff, 2016). This provides a direct link between a policy-induced nutritional shock around birth and learning outcomes measured several years later.

There is also setting-specific evidence that the policy can affect the relevant biological input. Dasgupta et al. (2023) show, using independent biomarker data, that beef restrictions reduce hemoglobin and increase anemia among women from communities in which beef consumption is customary. We therefore interpret maternal and early-life iron deficiency as a particularly plausible pathway linking the policy to later learning. Importantly, the present paper does not observe maternal biomarkers for the ASER children and therefore does not estimate this pathway directly.

The framework also allows for subsequent household responses. Parents may compensate for an early disadvantage by increasing investments in a child, or they may reinforce it if nutritional shocks reduce expected returns to later investments or coincide with tighter household resources. In our setting, private tutoring is a particularly salient margin of educational investment. We therefore examine paid private tutoring, tutoring expenditure, and the home reading environment as downstream responses. These outcomes should not be interpreted as alternative treatments or as mediators that identify the biological pathway; rather, they indicate whether later household investments tend to offset or reinforce an initial early-life disadvantage.

The framework thus yields two main empirical predictions. First, children exposed to stricter beef restrictions around birth should experience larger subsequent learning losses in locations where beef consumption is more prevalent. Second, later educational investments may either compensate for or reinforce those initial losses. The empirical analysis is designed around the first prediction, while the parental-investment outcomes help characterize how households respond after the initial shock.

3 Data and Empirical Strategy

3.1 Data

The empirical analysis brings together five sources of data that do not otherwise exist in a common data structure: an original historical panel of state-level beef legislation; nearly two million child observations from the Annual Status of Education Report (ASER); district-level religious composition from the 2011 Population Census; district-level caste composition from the 2011 Socio Economic and Caste Census (SECC)[3]; and ASER's separate village-level survey of government-school inputs. Linking the legal history to children's birth cohorts allows us to measure the policy environment around birth and observe learning and educational investments several years later.

Our first data source is an original historical database of Indian legislation governing cattle slaughter and, separately, the sale and possession of beef. We assembled the underlying legal record from state gazettes, legislation in Hindi and English, and federal sources for 26 states, recording policy changes at the month-year level from 1950 through 2018. The distinction between longstanding cattle-slaughter restrictions and later restrictions on the sale or possession of beef is important for the present analysis. During the birth-

[3]The SECC was conducted only once, in 2011; consequently, comparable district-level caste-composition data are unavailable for an earlier period.

cohort window relevant to our design, Jharkhand tightened its restrictions in 2005 and Gujarat in 2011 by introducing the sale-or-possession restrictions studied here; both states already had longstanding restrictions on cow slaughter.[4]

Our primary child data come from ASER for 2012, 2014, 2016, and 2018. We stop in 2018 to keep the analysis entirely pre-COVID, avoiding pandemic-era disruptions to schooling and learning (Singh et al., 2024). ASER is a large, repeated cross-sectional household survey designed to be representative of rural districts in India (Banerji et al., 2013). Importantly for our setting, children are interviewed and assessed in their homes rather than sampled through schools. The data therefore include children attending government or private schools as well as children who are not currently enrolled. ASER records children's age, sex, current schooling status, school type, household characteristics, and measures of household educational investment, including paid private tutoring.

ASER directly assesses children's basic mathematics and reading skills. Reading performance ranges from inability to recognize letters through reading a grade-two-level text, while mathematics ranges from basic number recognition through division. We transform the ordered achievement measures into age-standardized scores, so the learning outcomes are expressed in standard-deviation units relative to other children of the same age. Appendix Table A9 reports the underlying ASER assessment categories and their coding. Our baseline analysis focuses on children ages 5–12, providing a relatively homogeneous childhood sample and a common analysis population across the learning, heterogeneity, and mechanism exercises. As a robustness check, we retain ASER's full observed age range of 5–16.

ASER also administers a separate school survey in sampled villages. Surveyors visit the largest government primary school in the village, or the government primary school with the largest enrollment when more than one is present, and collect information on school infrastructure, enrollment, student and teacher attendance, staffing, and related inputs. These measures therefore characterize the observable local government-school environment; they are not observations on the particular school attended by each sampled child. We link these school data to household observations by village and survey year and use them only in the analysis of observable school inputs and alternative explanations.

Finally, we construct district-level exposure to beef consumption from population composition. Religion composition comes from the 2011 Population Census, while caste composition comes from the 2011 SECC. The district-level caste information required for the exposure measure has no comparable earlier SECC wave. Following the dietary patterns documented in the institutional literature, we define likely beef-consuming populations as Muslims, Christians, and Scheduled Caste Hindus and calculate their combined population share in each district. Because the measure is based on religious and caste composition rather than contemporaneous beef purchases, we treat it as a time-invariant proxy for the underlying prevalence of communities for whom beef is a culturally accepted food. This measure captures substantial within-state heterogeneity in the extent to which restrictions on beef are likely to bind.

3.2 Treatment and exposure

We define the 2005 birth cohort in Jharkhand and the 2011 birth cohort in Gujarat as the first treated cohorts. States in which the relevant restrictions were already in force before the usable cohort window are always treated and are excluded from the comparison group. The comparison group consists of states that remain untreated over the relevant window.

The third difference exploits within-state variation in the prevalence of likely beef consumers. For each

[4]The historical legal database was introduced in Dasgupta et al. (2023), which provides detailed documentation of the underlying legislation and uses it to study the contemporaneous nutritional effects of beef restrictions on women of reproductive age.

state, we classify a district as HighBeef when its beef-consumer share exceeds the *unweighted mean across districts in that state*; all other districts are LowBeef. The state-relative definition is the baseline because it creates a common within-state exposure contrast even in states, such as Gujarat, where no district exceeds the national mean. A national-cutoff definition is examined as a robustness check.

Figure 1 summarizes these two dimensions of exposure. District shading shows the prevalence of likely beef consumers; the labels identify the two states whose sale-or-possession restrictions begin during the relevant cohort window. Purple boundaries mark states in which these restrictions predate the window and which are consequently excluded as always treated.

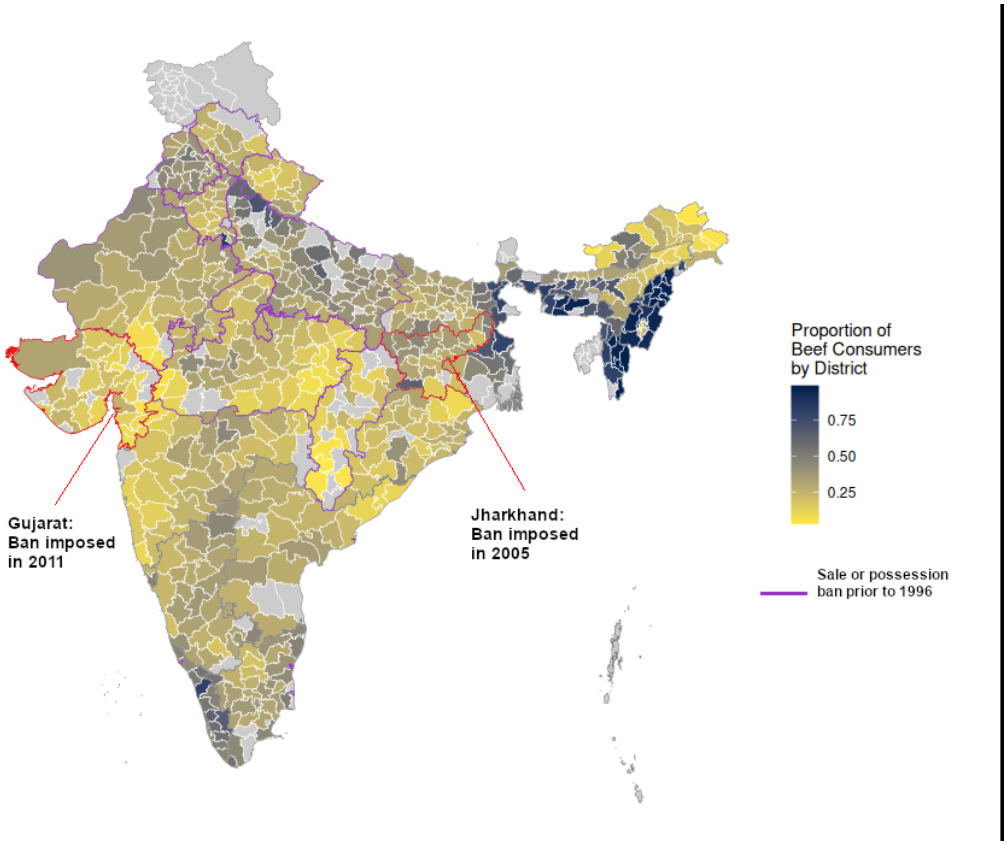


Figure 1: Proportion of likely beef consumers by district

Notes: District shading is the sum of the population shares of Muslims and Christians in the 2011 Population Census and Scheduled Caste Hindus in the 2011 SECC. The map identifies Gujarat and Jharkhand, which imposed the relevant sale-or-possession restrictions during the birth-cohort window used in the baseline analysis. Purple state boundaries denote states in which the restrictions were already in place before that window; these states are therefore always treated in the baseline design and are excluded from the comparison group. Source: Population Census (2011), Socio Economic and Caste Census (2011), and state legislation compiled by the authors.

The exposure measure is time-invariant. ASER identifies children's current district rather than district of birth, so individual migration histories are not observed directly. Differential policy-induced sorting across HighBeef and LowBeef districts is therefore a potential composition concern; the cohort-representation diagnostic reported in Appendix Table A11 and Figure A4 assesses whether such sorting is visible in the relative size of observed birth cohorts around the reforms. Additional details on the legal coding and exposure construction are provided in the Online Appendix.

3.3 Summary statistics

The main analysis sample contains 715,901 children ages 5–12 across 19 states. Appendix Table A10 reports summary statistics for this analysis population. Children are 8.6 years old on average, and 48.3 percent are girls. Using the state-relative cutoff, 47.8 percent of observations are in districts classified as high beef exposure. School enrollment is high in this household-based sample: 98.3 percent of children are currently enrolled, while 0.7 percent have dropped out and 1.0 percent have never enrolled. Outcome-specific sample sizes vary because reading and mathematics assessments and several household and school measures are not observed for every child.

3.4 Empirical strategy

We estimate the effects of beef-ban exposure using the multiple-period difference-in-differences framework of Callaway and Sant'Anna (2021), applied separately to HighBeef and LowBeef districts and then combined into a triple difference. This approach defines the causal contrasts directly by treatment cohort and birth cohort rather than relying on a pooled two-way fixed-effects interaction whose weighting can be difficult to interpret when treatment timing and effects vary.

The relevant time dimension is *birth cohort*, because treatment is exposure to the restriction around birth. Let G_s denote the first treated birth cohort in state s: $G_s = 2005$ for Jharkhand, $G_s = 2011$ for Gujarat, and $G_s = \infty$ for never-treated states. For child i living in district d, state s, and born in cohort b, define event time as

$$e = b - G_s.$$

The baseline common window is $e \in \{-2, -1, 0, 1, 2\}$. Always-treated states cannot contribute valid pre-treatment cohorts and are excluded.

For exposure group $h \in \{H, L\}$, corresponding to HighBeef and LowBeef districts, the Callaway–Sant'Anna building block is the group-time average treatment effect

$$ATT_h(g, b) = E\left[Y_{isdb}(1) - Y_{isdb}(0) \mid G_s = g,\ H_d = h\right], \tag{1}$$

where the untreated counterfactual is constructed from children in never-treated states. We estimate these parameters using ASER as repeated cross sections. For each event time e, the group-time effects for Jharkhand and Gujarat are aggregated over the treated groups for which that event time is observed, producing event-time effects $\theta_{H,e}$ and $\theta_{L,e}$.

Our target parameter is the difference between these two event-time effects:

$$\delta_e = \theta_{H,e} - \theta_{L,e}. \tag{2}$$

Thus, δ_e is the DDD at event time e: the treatment-induced change in the HighBeef–LowBeef learning gap in treated states relative to the corresponding change in never-treated states. The main DDD estimate averages the first three treated cohorts,

$$\delta^{\text{main}} = \frac{1}{3}\left(\delta_0 + \delta_1 + \delta_2\right), \tag{3}$$

while the pre-treatment placebo averages the two preceding event times,

$$\delta^{\text{pre}} = \frac{1}{2}\left(\delta_{-2} + \delta_{-1}\right). \tag{4}$$

Figure 2 reports the full vector $\{\delta_e\}_{e=-2}^{2}$.

The identifying restriction is a DDD parallel-trends condition. In the absence of the sale-or-possession restrictions, the evolution of the untreated HighBeef–LowBeef outcome gap in Jharkhand and Gujarat would have followed the same evolution as the HighBeef–LowBeef gap in never-treated states. The design also requires no anticipation by cohorts born before the restrictions take effect. The event-time estimates at $T-2$ and $T-1$ provide diagnostics for differential pre-treatment movements, although failure to reject those placebo effects does not establish the identifying assumption.

The pre- and post-treatment event coefficients have the standard Callaway–Sant'Anna interpretation. Post-treatment effects use the last untreated cohort $g-1$ as the base for a treated group, while the pre-treatment coefficients are placebo changes among untreated cohorts. In the implementation used here, the pre-period event coefficients use adjacent untreated cohorts as their comparison. Consequently, the pre-treatment average in equation (4) is a diagnostic placebo and is not subtracted from the main DDD in equation (3).

Inference accounts for the fact that the HighBeef and LowBeef DD estimates share the same state-level policy shocks and comparison states. We therefore do not treat the two components as independent. For each event time, we form the HighBeef-minus-LowBeef recentered influence-function contribution first and conduct inference on that joint DDD object. Standard errors and confidence intervals are clustered at the state level using a Mammen wild-cluster bootstrap with 999 repetitions. Tables report pointwise 95 percent confidence intervals; the event-study figure reports simultaneous 95 percent bands obtained from the joint bootstrap distribution of the event-time vector. Appendix Tables A6–A7 report Webb-weight wild-bootstrap robustness using the same DDD influence-function construction.

4 Results

4.1 Main and dynamic effects on learning

Table 1 reports the main estimates for children ages 5–12. The analysis includes Jharkhand and Gujarat as treated states, never-treated states as the comparison group, and a common event window from two cohorts before through two cohorts after each state's first treated birth cohort. Within each state, HighBeef districts are those with a beef-consumer share above the unweighted district mean.

The main DDD estimate is -0.075 standard deviations for mathematics and -0.065 for reading. Both 95 percent wild-bootstrap confidence intervals exclude zero. The corresponding pre-treatment placebo estimates are -0.006 and 0.014, respectively, and are small and statistically insignificant. Appendix Table A8 reports the underlying HighBeef and LowBeef DDs. Scores rise in LowBeef districts by 0.048 standard deviations in mathematics and 0.040 in reading, while the modest declines in HighBeef districts are not individually significant. The DDD therefore identifies a widening relative learning gap in more exposed districts, not an unconditional decline equal to the full DDD within HighBeef districts.

Table 1: Beef bans and learning outcomes: DDD estimates

	Mathematics z-score	Reading z-score
Main DDD estimate	−0.075 (0.012) [−0.102, −0.048]	−0.065 (0.015) [−0.095, −0.035]
Pre-treatment placebo estimate	−0.006 (0.010) [−0.025, 0.013]	0.014 (0.013) [−0.012, 0.040]
Observations	608,440	609,658
State clusters	19	19

Notes: The sample consists of children ages 5–12 in Jharkhand, Gujarat, and never-treated states; always-treated states are excluded. Jharkhand and Gujarat are first treated in the 2005 and 2011 birth cohorts. HighBeef equals one when a district's beef-consumer share exceeds the unweighted district mean within its state. Entries report HighBeef-minus-LowBeef DDD estimates. The main DDD estimate averages the first treated birth cohort and the following two birth cohorts; the pre-treatment placebo estimate averages the two immediately preceding birth cohorts. Each coefficient cell reports the estimate, the state-clustered Mammen wild-bootstrap standard error in parentheses, and the pointwise 95 percent wild-bootstrap confidence interval in brackets. Inference uses 999 repetitions and state clustering. DDD inference is computed after forming the joint HighBeef-minus-LowBeef recentered influence function. Appendix Table A8 reports the underlying HighBeef and LowBeef DD estimates.

Figure 2 shows no significant differential at $T-2$, $T-1$, or $T0$. The HighBeef–LowBeef differential turns negative at $T+1$, reaching −0.094 standard deviations in mathematics and −0.090 in reading, and widens at $T+2$ to −0.152 and −0.107, respectively.

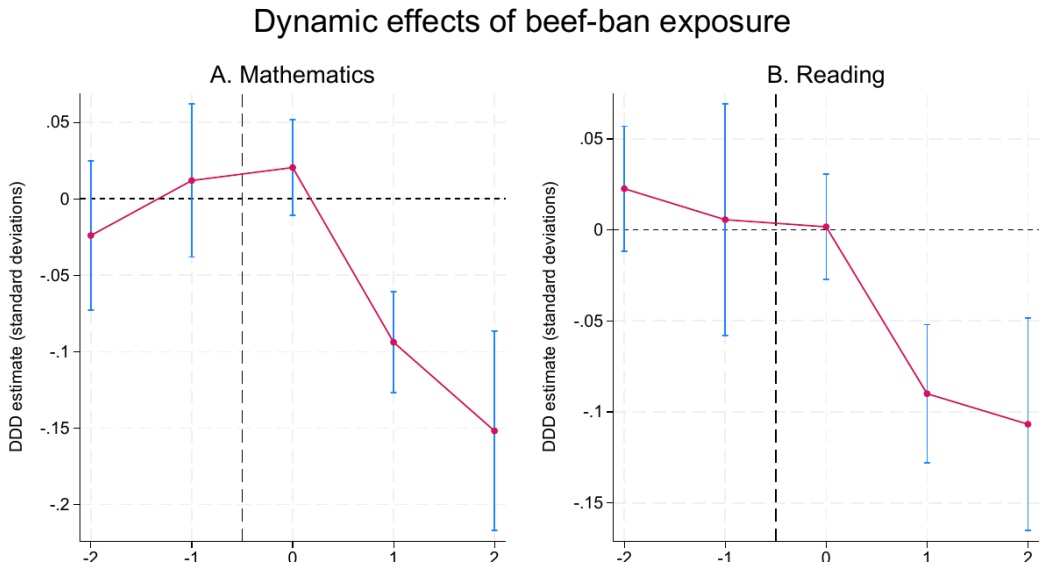


Figure 2: Dynamic effects of beef-ban exposure

Notes: Points are event-time HighBeef-minus-LowBeef DDD estimates from the baseline two-state specification. Bars are simultaneous 95 percent state-clustered Mammen wild-bootstrap confidence bands based on 999 repetitions. The vertical dashed line separates pre-treatment event times from the first exposed cohort; $T0$ denotes the first treated birth cohort.

4.2 Robustness and heterogeneity

Table 2 organizes the remaining test-score results around two questions. Panel A asks whether the main finding depends on the treated-state set, the exposure cutoff, the common event window, or the age restriction. The baseline row appears first as the benchmark, followed by one-at-a-time changes. Across specifications, the main DDD estimates remain negative, and all corresponding 95 percent confidence intervals exclude zero. The magnitudes range from -0.052 to -0.104 standard deviations across subjects and specifications. We use children ages 5–12 as the baseline sample to focus on a more homogeneous younger population and to maintain a common analysis sample across the learning, heterogeneity, and parental-investment results. Table 2 shows that the conclusions are similar when we instead retain ASER's full observed age range of 5–16.

The placebo evidence is mostly reassuring but not mechanically uniform. The baseline two-state estimates, the national-cutoff Jharkhand specification, the all-horizon two-state specification, and the full-age-range specification all have placebo confidence intervals that include zero. The Jharkhand-only state-relative specification produces a positive reading placebo estimate of 0.034 standard deviations. We therefore place greater weight on the prespecified two-state common-window design and its full event-study pattern than on any single restricted-sample specification.

Panel B examines heterogeneity by maternal human capital. Restricting the analysis to children with observed maternal-schooling information yields main DDD estimates of -0.077 standard deviations in mathematics and -0.056 in reading. The mathematics effect is substantially larger among children whose mothers never attended school (-0.130) than among those whose mothers attended some school (-0.043). The formal difference is -0.087 standard deviations and its 95 percent confidence interval excludes zero. For reading, the corresponding estimates are -0.082 and -0.054, but their difference is not statistically significant. The evidence therefore supports stronger mathematics losses among children with lower maternal human capital, but not a statistically distinct reading effect.

Table 2: Design robustness and heterogeneity by maternal human capital

Specification	Math DDD	Math placebo	Reading DDD	Reading placebo	N (math)	N (reading)	States
Panel A. Design, exposure, event-window, and age-sample robustness							
	−0.075	−0.006	−0.065	0.014			
	(0.012)	(0.010)	(0.015)	(0.013)			
Baseline: Jharkhand + Gujarat; state-relative HighBeef; ages 5–12; common window	[−0.102, −0.048]	[−0.025, 0.013]	[−0.095, −0.035]	[−0.012, 0.040]	608,440	609,658	19
	−0.076	−0.009	−0.072	0.034			
	(0.015)	(0.009)	(0.014)	(0.007)			
Jharkhand only; state-relative HighBeef; otherwise baseline	[−0.104, −0.047]	[−0.025, 0.008]	[−0.101, −0.042]	[0.019, 0.049]	566,885	567,910	18
	−0.071	−0.014	−0.092	−0.007			
	(0.017)	(0.010)	(0.019)	(0.014)			
Jharkhand only; national HighBeef cutoff; otherwise baseline	[−0.105, −0.036]	[−0.034, 0.005]	[−0.128, −0.056]	[−0.032, 0.018]	566,885	567,910	18
	−0.104	−0.004	−0.083	0.006			
	(0.024)	(0.004)	(0.023)	(0.004)			
All available event horizons; otherwise baseline	[−0.153, −0.055]	[−0.011, 0.004]	[−0.127, −0.039]	[−0.002, 0.013]	608,440	609,658	19
	−0.089	−0.015	−0.052	−0.003			
	(0.014)	(0.008)	(0.011)	(0.005)			
Ages 5–16; otherwise baseline	[−0.115, −0.063]	[−0.032, 0.003]	[−0.073, −0.030]	[−0.012, 0.007]	762,632	764,106	19
Panel B. Heterogeneity by maternal human capital							
	−0.077	−0.005	−0.056	0.003			
	(0.013)	(0.013)	(0.015)	(0.009)			
Maternal-schooling status observed	[−0.102, −0.052]	[−0.031, 0.020]	[−0.084, −0.028]	[−0.014, 0.020]	438,265	439,124	19
	−0.130	0.013	−0.082	0.031			
	(0.017)	(0.012)	(0.024)	(0.021)			
Mother never attended school	[−0.163, −0.097]	[−0.010, 0.037]	[−0.126, −0.039]	[−0.029, 0.091]	158,897	159,234	19
	−0.043	0.001	−0.054	0.003			
	(0.022)	(0.016)	(0.021)	(0.017)			
Mother attended some school	[−0.084, −0.001]	[−0.034, 0.035]	[−0.094, −0.013]	[−0.037, 0.044]	279,368	279,890	19
	−0.087	0.013	−0.028	0.028			
	(0.024)	(0.017)	(0.025)	(0.018)			
Equality test: mother never attended − mother attended some school	[−0.139, −0.035]	[−0.037, 0.063]	[−0.075, 0.019]	[−0.062, 0.118]	438,265	439,124	19

Notes: Entries are HighBeef-minus-LowBeef DDD estimates. The baseline design uses children ages 5–12, treats Jharkhand and Gujarat at their respective policy dates, uses never-treated states as controls, defines HighBeef relative to the unweighted district mean within each state, and uses the common event window from $T-2$ through $T+2$. The main DDD averages event times 0, 1, and 2; the placebo averages event times -2 and -1. Panel A changes only the feature named in each row: the national cutoff uses the national district mean; Jharkhand only excludes Gujarat; all available horizons uses all event times estimable in both exposure groups; and ages 5–16 changes only the age range. Panel B restricts to children with observed maternal-schooling status. The final row tests whether the DDD differs between the two maternal-schooling groups and reports mother never attended minus mother attended some school; the test is computed jointly to account for covariance between the subgroup estimates. Cells report the estimate, state-clustered Mammen wild-bootstrap standard error, and pointwise 95 percent wild-bootstrap confidence interval. Inference uses 999 repetitions.

4.3 Parental investment responses and alternative explanations

Table 3 studies parental investment responses and two competing explanations for the learning results. A substantial literature asks whether parents compensate for or reinforce fetal and early-life shocks through later investments in children's health and education (Almond and Mazumder, 2013; Yi et al., 2015; Adhvaryu and Nyshadham, 2016). Banerjee and Majid (forthcoming) further show that parental investment responses to fetal shocks tend to reinforce initial inequalities and vary with macroeconomic conditions. Against this background, we examine whether parents adjust educational inputs after beef-ban exposure, while also assessing selective school participation and changes in observable school inputs.

Panel A provides little support for selective school exit. The main DDD estimate for dropout is -0.0012 and statistically insignificant, while the estimate for never having enrolled is -0.0033. Because ASER assesses children at home regardless of enrollment, children do not mechanically disappear from the achievement sample when they leave school. Appendix Table A1 reports a small positive DDD for current enrollment—less than one-half of a percentage point—and an imprecise estimate close to zero for being on track for age. We do not interpret these small auxiliary estimates as evidence of a meaningful enrollment gain.

Panel B examines household educational investments. Paid private tutoring—commonly referred to as "tuition" in India—is an important form of supplementary human-capital investment across Asian education systems. Recent work shows that private tutoring can represent a substantial household expenditure, respond to educational competition and public-school inputs, and meaningfully affect learning (Banerjee et al., 2007; Yuan and Zhang, 2015; Muralidharan et al., 2019; Guo and Qu, 2022; Burchardi et al., 2024; Meng et al., 2026). In our sample, 27 percent of children receive paid private tutoring, including 24 percent of children whose mothers never attended school and 29 percent of those whose mothers attended some school. Following the bans, tutoring participation and expenditure increase less in HighBeef than in LowBeef districts: the main DDD estimate is -1.67 percentage points for participation and -0.039 standard deviations for expenditure, while the home reading environment changes little. Thus, the learning losses are accompanied by a relative shortfall in a widely used form of supplementary educational investment, although these reduced-form estimates do not establish private tutoring as the causal channel.

Panel C evaluates a prominent alternative explanation: deterioration in the schooling environment. The post-treatment DDD estimate for the observable school-input index is 0.005 and statistically indistinguishable from zero. The pre-treatment placebo is 0.019 and statistically different from zero, so we do not interpret the school-input series as a separate causal outcome. A formal comparison indicates that the post-treatment estimate is not statistically distinguishable from the placebo estimate: the post-treatment-minus-placebo contrast is -0.0145, with a 95 percent wild-bootstrap confidence interval of $[-0.0549, 0.0259]$. Appendix Figure A2 likewise provides no systematic evidence of deterioration after treatment. Our narrower conclusion is that the data provide no evidence of a discrete decline in observed school inputs that could account for the learning losses. In the school-linked sample, the mathematics and reading estimates remain essentially unchanged after adjustment for the observed school-input index, indicating that differences in measured school inputs do not account for the learning results.

Table 3: Parental investment responses and alternative explanations

Outcome	Main DDD estimate	Pre-treatment placebo	Mean	*N*
Panel A. School participation and sample selection				
	−0.0012	0.0005		
	(0.0012)	(0.0005)		
Dropped out	[−0.0033, 0.0009]	[−0.0024, 0.0034]	0.007	715,901
	−0.0033	0.0006		
	(0.0013)	(0.0008)		
Never enrolled	[−0.0059, −0.0007]	[−0.0014, 0.0027]	0.010	715,901
Panel B. Parental educational investments				
	0.008	−0.018		
	(0.016)	(0.010)		
Home reading-environment index (SD)	[−0.024, 0.040]	[−0.038, 0.002]	−0.028	701,289
	−0.0167	−0.0036		
	(0.0080)	(0.0046)		
Any paid private tutoring	[−0.0320, −0.0014]	[−0.0152, 0.0081]	0.272	653,261
	−0.039	−0.006		
	(0.018)	(0.014)		
Private tutoring expenditure (IHS, SD)	[−0.072, −0.006]	[−0.039, 0.028]	−0.024	489,498
Panel C. Observable school inputs				
	0.005	0.019		
	(0.017)	(0.009)		
Observable school-input index (SD)	[−0.028, 0.037]	[0.002, 0.036]	−0.015	686,259

Notes: Entries are HighBeef-minus-LowBeef DDD estimates. The main DDD estimate averages the first treated birth cohort and the following two birth cohorts; the pre-treatment placebo averages the two immediately preceding birth cohorts. Each coefficient cell reports the estimate, the state-clustered Mammen wild-bootstrap standard error in parentheses, and the pointwise 95 percent wild-bootstrap confidence interval in brackets. Inference uses 999 repetitions and 19 state clusters. For the observable school-input outcome in Panel C, the paired post-treatment-minus-placebo contrast is −0.0145 (state-clustered Mammen wild-bootstrap standard error 0.0206; 95 percent confidence interval [−0.0549, 0.0259]); the contrast is formed from the joint influence functions and accounts for covariance between the two averages. ASER tests children in their homes, including children who are not enrolled. The private-tutoring indicator equals zero only when the survey explicitly records no paid supplementary tutoring (reported as "tuition" in ASER); records with missing tutoring status are excluded. Private-tutoring expenditure is unavailable in 2012; in 2014–2018, recorded nonusers receive zero and users with missing reported amounts are excluded. The school-input index is measured at the village-by-survey-year level, standardized within survey year using never-treated village-years, and assigned to children according to village and survey year.

5 Additional Robustness

The baseline inference uses state-clustered Mammen wild-bootstrap intervals. Appendix Table A6 applies Webb's six-point auxiliary weights to the same recentered influence functions. The main DDD estimates remain statistically significant at the 1 percent level for both mathematics and reading, while both pre-treatment placebo estimates remain insignificant. Appendix Table A7 and Figure A3 show the same delayed event-time pattern: $T + 1$ and $T + 2$ are significant for both subjects, the two pre-treatment coefficients are jointly insignificant, and the post-treatment coefficients are jointly significant.[5]

The remaining appendix material documents the components and scope of the parental-investment and alternative-explanation analyses in Table 3. Appendix Table A1 restores the HighBeef and LowBeef com-

[5] A higher-replication audit using a second seed and alternative bootstrap options yields the same substantive conclusions.

ponent DDs and the additional participation outcomes. Appendix Table A2 varies the school-input index and observation unit and reports test-score estimates with and without adjustment for observed school inputs in the school-linked sample. Appendix Table A3 shows that the private-tutoring responses do not differ significantly by maternal schooling. Appendix Table A4 documents outcome availability, while Appendix Table A5 reports the complete event-time coefficients underlying the private-tutoring and school-input figures. Appendix Table A8 reports the HighBeef and LowBeef component DDs underlying the DDD-only main learning estimates.

Because district-level caste composition is available only from the 2011 SECC, we test for differential changes in observed cohort representation around treatment that could reflect migration or other demographic selection. The post-treatment DDD is -0.041 (95 percent confidence interval $[-0.163, 0.082]$), the placebo is 0.016 ($[-0.077, 0.108]$), and all simultaneous confidence bands include zero (Appendix Table A11 and Appendix Figure A4).

Because maternal schooling is strongly associated with children's learning, we also test whether the maternal-schooling composition of the observed sample changed differentially around treatment. The post-treatment DDD is -0.010, the placebo is -0.009, and the post-treatment-minus-placebo difference is -0.001 (95 percent confidence interval $[-0.061, 0.058]$); all simultaneous confidence bands include zero (Appendix Table A12 and Appendix Figure A5).

Finally, we examine household wealth to assess whether changes in economic resources could account for the learning results. Although wealth may itself respond to nutritional shocks, the post-treatment wealth differential is statistically indistinguishable from both zero and the pre-treatment differential, providing no evidence of a discrete change in household wealth at treatment (Appendix Table A13 and Appendix Figure A6).

6 Discussion

The results indicate that exposure to stricter beef restrictions around birth reduces later learning among children in districts with greater ex ante exposure to beef consumption. The main DDD estimates imply losses of 0.075 standard deviations in mathematics and 0.065 standard deviations in reading. The event-study estimates suggest that the learning differential becomes more pronounced for cohorts born after the restrictions had been in force for longer, while the corresponding pre-treatment estimates are small. The conclusions are also similar across alternative definitions of exposure, treatment samples, event windows, and when ASER's full age range is retained. The larger mathematics effect among children whose mothers never attended school further suggests that the consequences of the shock are more pronounced where maternal human capital is lower.

A natural interpretation is an early-life nutritional pathway. This interpretation is supported by our earlier evidence from the same policy setting. Dasgupta et al. (2023) show that beef restrictions lower women's hemoglobin and increase anemia among communities for whom beef is a customary food. The present results therefore extend the evidence from contemporaneous nutritional status to a human-capital outcome measured several years later.

The magnitude of the learning effects is also consistent with evidence from other early-life nutritional shocks. In rural Nepal, randomized prenatal iron and folic-acid supplementation reduced third-trimester maternal anemia by approximately 54 percent and improved offspring intellectual performance at ages 7–9 by roughly one-quarter of a standard deviation (Christian et al., 2010). A simple linear scaling implies approximately 0.04 standard deviations in cognition for an anemia change of the relative magnitude documented in our earlier

beef-ban study, a sizeable fraction of the 0.065–0.075 standard-deviation learning losses estimated here. Evidence from prenatal Ramadan exposure provides another useful benchmark. Almond et al. (2015) find reductions of approximately 0.05–0.08 standard deviations in mathematics, reading, and writing achievement at age seven following exposure during pregnancy, remarkably close to the magnitudes we estimate. Using the published score coefficients and score dispersion in Indonesian data, the effects in Majid (2015) imply losses of roughly one-third of a standard deviation in mathematics and Raven cognitive performance among children exposed to Ramadan in utero. Together with prospective evidence linking maternal iron-deficiency anemia during pregnancy to subsequent cognitive development (Fararouei et al., 2010; Chang et al., 2013), these comparisons place our estimates comfortably within the range of developmental effects documented for prenatal nutritional conditions. We view them as evidence of quantitative plausibility rather than estimates of mediation: the underlying shocks and biological pathways differ, and our data do not link maternal biomarkers directly to the ASER children.

The household-investment results provide a complementary perspective on what happens after the initial early-life disadvantage. Paid private tutoring is common in our sample, with 27 percent of children receiving supplementary instruction. Following the bans, tutoring participation increases less in HighBeef than in LowBeef districts. The main DDD of -1.67 percentage points corresponds to about 6 percent of mean tutoring participation. Private-tutoring expenditure displays a similar relative shortfall, whereas the home reading environment changes little. This pattern is consistent with a reinforcing rather than compensating response: subsequent private educational investment moves in the same direction as the learning gap rather than offsetting it. This interpretation is consistent with a broader literature showing that private tutoring is an important and responsive component of household human-capital investment and educational competition (Yuan and Zhang, 2015; Guo and Qu, 2022; Burchardi et al., 2024; Meng et al., 2026).

A simple calibration also suggests that the extensive-margin tutoring response is unlikely to be the principal source of the learning loss. Muralidharan et al. (2019) study an unusually intensive personalized after-school program in India: offering free access increased participation in the program by approximately 58 percentage points and raised mathematics and Hindi achievement by 0.37 and 0.23 standard deviations, respectively. Linearly scaling those experimental effects to our 1.67-percentage-point tutoring differential implies effects of only about 0.011 standard deviations in mathematics and 0.007 standard deviations in language, or roughly 10–14 percent of our main learning estimates. This is deliberately a generous benchmark: the experimental program was substantially more intensive and structured than the heterogeneous paid tutoring reported in ASER, and our expenditure result suggests adjustment along the intensive margin as well. The calculation therefore should not be read as a mediation estimate. Rather, it indicates that the tutoring response can plausibly reinforce the initial disadvantage while being too small, by itself, to explain most of the learning loss.

Several findings make alternative explanations less compelling. Because ASER assesses children at home regardless of enrollment, selective school exit does not mechanically remove children from the achievement sample. Consistent with this feature of the survey, we find no increase in dropout or never enrollment among more exposed children in the baseline age range. Nor do we find evidence of a post-ban deterioration in observed school inputs capable of accounting for the learning losses. In the school-linked sample, the mathematics and reading estimates remain essentially unchanged when the observed school-input index is included. Because the school-input series displays a positive pre-treatment placebo and contemporaneous school inputs could themselves respond to treatment, we interpret this exercise as evidence against a deterioration in observed school inputs rather than as an estimate of a separate causal mechanism.

Taken together, the evidence points to an early-life human-capital shock whose consequences remain visible several years later and are accompanied by, rather than offset by, a relative shortfall in supplementary educational investment. The findings fit a broader literature showing that nutritional conditions during sensitive developmental periods can have persistent effects on later skills (Almond and Mazumder, 2011; Almond et al., 2015; Majid, 2015). What is distinctive in this setting is the source of the shock. The

restrictions were not designed as health or education policy; they tightened access to a culturally specific and nutritionally valuable food consumed disproportionately by particular religious and caste communities. The results therefore reveal a less visible distributional consequence of culturally motivated regulation: policies that differentially constrain access to nutritionally valuable foods can impose developmental costs that are not necessarily apparent when the regulation is enacted but emerge years later in children's human capital.

References

Achyuta Adhvaryu and Anant Nyshadham. Endowments at birth and parents' investments in children. *Economic Journal*, 126(593):781–820, June 2016. URL https://ideas.repec.org/a/wly/econjl/v126y2016i593p781-820.html.

Douglas Almond and Bhashkar Mazumder. Health capital and the prenatal environment: the effect of ramadan observance during pregnancy. *American Economic Journal: Applied Economics*, 3(4):56–85, 2011.

Douglas Almond and Bhashkar Mazumder. Fetal origins and parental responses. *Annual Review of Economics*, 5:37–56, 2013. doi: 10.1146/annurev-economics-082912-110145. URL https://doi.org/10.1146/annurev-economics-082912-110145.

Douglas Almond, Bhashkar Mazumder, and Reyn van Ewijk. In utero ramadan fasting and cognitive development: Evidence from england. *Labour Economics*, 34:135–145, 2015.

Douglas Almond, Janet Currie, and Valentina Duque. Childhood circumstances and adult outcomes: Act ii. *Journal of Economic Literature*, 56(4):1360–1446, 2018.

David Atkin. Trade, tastes, and nutrition in india. *American Economic Review*, 103(5):1629–1663, 2013.

David Atkin. Costly tastes: The economic consequences of cultural food preferences. *Journal of Political Economy*, 124(6):1591–1632, 2016.

Abhijit V. Banerjee, Shawn Cole, Esther Duflo, and Leigh Linden. Remedying education: Evidence from two randomized experiments in india. *The Quarterly Journal of Economics*, 122(3):1235–1264, 2007.

Rakesh Banerjee and Farhan Majid. Parental response to fetal shocks and macroeconomic volatility. *Review of Economics and Statistics*, forthcoming.

Rukmini Banerji, Suman Bhattacharjea, and Wilima Wadhwa. The annual status of education report (aser). *Research in Comparative and International Education*, 8(3):387–396, 2013.

Sascha O. Becker, Jared Rubin, and Ludger Woessmann. Religion in economic history. *Journal of Economic Literature*, 62(1):3–58, 2024.

Clive Bell. *Food, Culture, and Social Structure*. Cambridge University Press, 1968.

Konrad Burchardi, Jonathan de Quidt, Selim Gulesci, and Munshi Sulaiman. Borrowing constraints and demand for remedial education: Evidence from tanzania. *The Economic Journal*, 134(662):2621–2637, 2024. doi: 10.1093/ej/ueae024. URL https://doi.org/10.1093/ej/ueae024.

Brantly Callaway and Pedro H. C. Sant'Anna. Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230, 2021.

Filipe Campante and David Yanagizawa-Drott. Does religion affect economic growth and happiness? evidence from ramadan. *The Quarterly Journal of Economics*, 130(2):615–658, 2015.

Suying Chang, Lingxia Zeng, Inge D Brouwer, Frans J Kok, and Hong Yan. Effect of iron deficiency anemia in pregnancy on child mental development in rural china. *Pediatrics*, 131(3):e755–e763, 2013.

Parul Christian, Laura E Murray-Kolb, Subarna K Khatry, Joanne Katz, Barbara A Schaefer, Pamela M Cole, Steven C LeClerq, and James M Tielsch. Prenatal micronutrient supplementation and intellectual and motor function in early school-aged children in nepal. *Jama*, 304(24):2716–2723, 2010.

Sarah E. Cusick and Michael K. Georgieff. The role of nutrition in brain development: The golden opportunity of the first 1000 days. *The Journal of Pediatrics*, 175:16–21, 2016.

Aparajita Dasgupta, Farhan Majid, and Wafa Hakim Orman. The nutritional cost of beef bans in india. *Journal of Development Economics*, 163:103104, 2023. ISSN 0304-3878. doi: 10.1016/j.jdeveco.2023.103104. URL https://www.sciencedirect.com/science/article/pii/S0304387823000597.

Mohammad Fararouei, Claire Robertson, John Whittaker, Ulla Sovio, Aimo Ruokonen, Anneli Pouta, Anna-Liisa Hartikainen, Marjo-Riitta Jarvelin, and Elina Hyppönen. Maternal hb during pregnancy and offspring's educational achievement: a prospective cohort study over 30 years. *British journal of nutrition*, 104(9):1363–1368, 2010.

Mathilde Ferry. Who eats beef in india? caste, religion, and dietary preferences. *Development and Change*, 51(4):1012–1038, 2020.

Michael K. Georgieff. The role of iron in neurodevelopment: Fetal and early childhood periods. *The American Journal of Clinical Nutrition*, 89(2):693S–696S, 2008.

Jonathan Gruber and Daniel M. Hungerman. The church versus the mall: What happens when religion meets secular competition? *The Quarterly Journal of Economics*, 123(2):831–862, 2008.

Juncong Guo and Xi Qu. Competition in household human capital investments: Strength, motivations and consequences. *Journal of Development Economics*, 158:102937, 2022. doi: 10.1016/j.jdeveco.2022.102937. URL https://doi.org/10.1016/j.jdeveco.2022.102937.

Leif Hallberg and Lena Hulthén. Prediction of dietary iron absorption: An algorithm for calculating absorption and bioavailability of dietary iron. *The American Journal of Clinical Nutrition*, 71(5):1147–1160, 2000.

Richard Hurrell and Ines Egli. Iron bioavailability and dietary reference values. *The American Journal of Clinical Nutrition*, 91(5):1461S–1467S, 2010.

Laurence R. Iannaccone. Introduction to the economics of religion. *Journal of Economic Literature*, 36(3): 1465–1496, 1998.

Timur Kuran. Why the middle east is economically underdeveloped: Historical role of commercial legal institutions. *Journal of Economic Perspectives*, 18(3):71–90, 2004.

Farhan Majid. The persistent effects of in utero nutrition shocks: The case of ramadan in indonesia. *Journal of Development Economics*, 117:125–143, 2015.

Farhan Majid and Jere R. Behrman. Early life shocks and human capital formation. In *Handbook of Labor, Human Resources and Population Economics*, pages 1–35. Springer, 2021.

John A. Maluccio, John Hoddinott, Jere R. Behrman, Reynaldo Martorell, Agnes R. Quisumbing, and Rafael Flores. The impact of improving nutrition during early childhood on education among guatemalan adults. *The Economic Journal*, 119(537):734–763, 2009.

Juanjuan Meng, Hui Wang, Yu Alan Yang, and Mingshan Zhang. Education competition and fertility intention: Evidence from china's private tutoring ban. *The Economic Journal*, 136(676):1504–1526, 2026. doi: 10.1093/ej/ueaf096. URL https://doi.org/10.1093/ej/ueaf096.

Karthik Muralidharan and Abhijeet Singh. India's primary schooling system: Structural challenges and possible solutions. *Indian Economic Review*, 56(1):141–164, 2021.

Karthik Muralidharan, Abhijeet Singh, and Alejandro J. Ganimian. Disrupting education? experimental evidence on technology-aided instruction in india. *American Economic Review*, 109(4):1426–1460, 2019.

Charlotte G. Neumann, Shannon E. Whaley, Susan P. Murphy, Nimrod O. Bwibo, Jonathan H. Siekmann, Robert E. Weiss, M. Bwibo, and Lindsay H. Allen. Meat supplementation improves growth, cognitive development, and activity of kenyan school children. *The Journal of Nutrition*, 137(4):1119–1123, 2007.

Manisha Shah and Bryce Millett Steinberg. Droughts, schooling, and child labor in rural india. *American Economic Review*, 107(2):421–453, 2017.

Abhijeet Singh, Mauricio Romero, and Karthik Muralidharan. Covid-19 learning loss and recovery. *Journal of Human Resources*, 2024. ISSN 0022-166X. doi: 10.3368/jhr.0723-13025R2. URL https://jhr.uwpress.org/content/early/2024/04/01/jhr.0723-13025R2.

Duncan Webb. Critical periods in cognitive and socioemotional development: Evidence from weather shocks in indonesia. *The Economic Journal*, 134(660):1637–1665, 2024. doi: 10.1093/ej/uead105. URL https://doi.org/10.1093/ej/uead105.

Max Weber. *The Protestant Ethic and the Spirit of Capitalism*. Allen & Unwin, 1904/05.

Shannon E. Whaley, Charlotte G. Neumann, Nimrod O. Bwibo, Robert E. Weiss, Marian Sigman, and Charlotte G. Neumann. The micronutrient impacts on cultural performance (micp) project: Meat supplementation improves cognition in kenyan school children. *The Journal of Nutrition*, 133(11):3965S–3971S, 2003.

Junjian Yi, James J Heckman, Junsen Zhang, and Gabriella Conti. Early health shocks, intra-household resource allocation and child outcomes. *The Economic Journal*, 125(588):F347–F371, 2015.

Cheng Yuan and Lei Zhang. Public education spending and private substitution in urban china. *Journal of Development Economics*, 115:124–139, 2015. doi: 10.1016/j.jdeveco.2015.02.006. URL https://doi.org/10.1016/j.jdeveco.2015.02.006.

A Online Appendix

A.1 Treatment and exposure construction

The legal database follows the state-level legislation assembled in Dasgupta et al. (2023) from state gazettes and other official sources and distinguishes long-standing cattle-slaughter prohibitions from later restrictions on the sale or possession of beef. The treatment used in this paper is the latter. Within the cohort window that can be linked to the 2012–2018 ASER rounds, Jharkhand first enters treatment in the 2005 birth cohort and Gujarat in the 2011 birth cohort. States whose sale-or-possession restrictions predate the usable cohort window are classified as always treated and excluded; the comparison states are never treated over the window.

District exposure is constructed using religion composition from the 2011 Population Census and caste composition from the 2011 SECC. The likely beef-consumer share is the combined district population share of Muslims, Christians, and Scheduled Caste Hindus. The baseline HighBeef indicator equals one when this share is above the unweighted mean across districts within the same state. This state-relative threshold is fixed across cohorts and is chosen so that the exposure contrast is defined within every state, including Gujarat. The national-mean cutoff used in the earlier version of the analysis is retained only as a robustness specification. Because ASER records current residence rather than birthplace, treatment assignment uses the child's current state and district together with birth year.

Table A1: Full components of school participation and household educational-investment estimates

Outcome	Period	HighBeef DD	LowBeef DD	DDD: High − Low	Mean	N
Panel A. School participation						
		−0.0025	−0.0013	−0.0012		
		(0.0018)	(0.0010)	(0.0012)		
Dropped out	Main estimate	[−0.0059, 0.0010]	[−0.0032, 0.0007]	[−0.0033, 0.0009]	0.007	715,901
		0.0002	−0.0003	0.0005		
		(0.0005)	(0.0005)	(0.0005)		
	Pre-treatment placebo	[−0.0009, 0.0013]	[−0.0026, 0.0019]	[−0.0024, 0.0034]		
		0.0175	0.0166	0.0009		
		(0.0179)	(0.0244)	(0.0094)		
On track for age	Main estimate	[−0.0186, 0.0536]	[−0.0357, 0.0689]	[−0.0200, 0.0217]	0.899	618,766
		0.0018	−0.0130	0.0148		
		(0.0093)	(0.0088)	(0.0046)		
	Pre-treatment placebo	[−0.0229, 0.0265]	[−0.0319, 0.0058]	[0.0050, 0.0247]		
		0.0070	0.0024	0.0045		
		(0.0013)	(0.0020)	(0.0022)		
Currently enrolled	Main estimate	[0.0047, 0.0093]	[−0.0013, 0.0062]	[0.0004, 0.0086]	0.983	715,901
		−0.0013	−0.0002	−0.0011		
		(0.0005)	(0.0005)	(0.0007)		
	Pre-treatment placebo	[−0.0024, −0.0003]	[−0.0011, 0.0007]	[−0.0025, 0.0003]		
		−0.0045	−0.0012	−0.0033		
		(0.0021)	(0.0026)	(0.0013)		
Never enrolled	Main estimate	[−0.0081, −0.0009]	[−0.0061, 0.0037]	[−0.0059, −0.0007]	0.010	715,901
		0.0011	0.0005	0.0006		
		(0.0002)	(0.0008)	(0.0008)		
	Pre-treatment placebo	[0.0008, 0.0014]	[−0.0015, 0.0025]	[−0.0014, 0.0027]		
Panel B. Household educational investments						
		−0.0052	−0.0132	0.0080		
		(0.0178)	(0.0090)	(0.0157)		
Home reading-environment index (SD)	Main estimate	[−0.0419, 0.0316]	[−0.0317, 0.0053]	[−0.0243, 0.0403]	−0.028	701,289
		−0.0308	−0.0125	−0.0183		
		(0.0121)	(0.0105)	(0.0103)		
	Pre-treatment placebo	[−0.0547, −0.0069]	[−0.0436, 0.0187]	[−0.0384, 0.0017]		
		0.0084	0.0251	−0.0167		
		(0.0072)	(0.0126)	(0.0080)		
Any paid private tutoring	Main estimate	[−0.0048, 0.0216]	[0.0018, 0.0484]	[−0.0320, −0.0014]	0.272	653,261
		−0.0078	−0.0042	−0.0036		
		(0.0058)	(0.0036)	(0.0046)		
	Pre-treatment placebo	[−0.0190, 0.0034]	[−0.0114, 0.0029]	[−0.0152, 0.0081]		
		0.0206	0.0599	−0.0393		
		(0.0166)	(0.0335)	(0.0182)		
Private tutoring expenditure (IHS, SD)	Main estimate	[−0.0138, 0.0550]	[0.0047, 0.1151]	[−0.0725, −0.0062]	−0.024	489,498
		−0.0097	−0.0038	−0.0058		
		(0.0151)	(0.0084)	(0.0144)		
	Pre-treatment placebo	[−0.0416, 0.0222]	[−0.0193, 0.0116]	[−0.0392, 0.0275]		

Notes: This table restores the HighBeef and LowBeef component DDs underlying the compact DDD-only results in Table 3 and reports the two additional participation outcomes. The main estimate averages event times 0, 1, and 2; the pre-treatment placebo averages event times −2 and −1. Each coefficient cell reports the estimate, state-clustered Mammen wild-bootstrap standard error in parentheses, and pointwise 95 percent confidence interval in brackets. The current-enrollment DDD is positive but small, while the on-track estimate is close to zero and imprecise. For private tutoring, the component estimates show increases in both exposure groups but smaller increases in HighBeef districts; the negative DDDs therefore represent relative shortfalls rather than unconditional declines. Remaining design and inference details follow Table 3.

Table A2: Observable school inputs and test-score adjustment

Specification	Main DDD estimate	Pre-treatment placebo	N
Panel A. School-input outcome specifications			
Mean school-input index, survey-year standardized; village-year-cohort cells	0.0057 (0.0091) [−0.0124, 0.0238]	0.0052 (0.0116) [−0.0191, 0.0295]	244,236
PCA school-quality index; village-year-cohort cells	−0.0075 (0.0091) [−0.0257, 0.0106]	−0.0000 (0.0082) [−0.0161, 0.0161]	155,405
Panel B. Test-score estimates with and without adjustment for observed school inputs			
Mathematics: baseline sample	−0.0751 (0.0124) [−0.1017, −0.0485]	−0.0060 (0.0098) [−0.0247, 0.0127]	608,440
Mathematics: school-linked sample, unadjusted	−0.0771 (0.0142) [−0.1047, −0.0495]	−0.0033 (0.0085) [−0.0201, 0.0136]	584,087
Mathematics: same sample, adjusting for school-input index	−0.0804 (0.0117) [−0.1027, −0.0581]	−0.0082 (0.0069) [−0.0213, 0.0049]	584,087
Reading: baseline sample	−0.0650 (0.0149) [−0.0949, −0.0351]	0.0142 (0.0127) [−0.0117, 0.0400]	609,658
Reading: school-linked sample, unadjusted	−0.0724 (0.0178) [−0.1047, −0.0400]	0.0151 (0.0141) [−0.0115, 0.0417]	585,263
Reading: same sample, adjusting for school-input index	−0.0757 (0.0175) [−0.1083, −0.0432]	0.0080 (0.0171) [−0.0277, 0.0438]	585,263

Notes: Panel A reports the two retained school-input robustness specifications at the village-year-birth-cohort cell level. The school-linked sample in Panel B consists of children observed in village-years for which the school survey provides a nonmissing school-input index. Panel B compares the baseline test-score estimate with an unadjusted estimate in the school-linked sample and an estimate in the identical sample adjusting for the standardized mean school-input index. These specifications assess whether the test-score estimates are sensitive to adjustment for the observed school-input index. Cells report estimates, standard errors, and pointwise 95 percent Mammen wild-cluster-bootstrap confidence intervals. All estimates use 19 state clusters.

Table A3: Household educational investments by maternal schooling

Outcome	Period	Mother never attended school	Mother attended some school	Difference: never − some
		−0.0176	0.0224	−0.0399
		(0.0189)	(0.0289)	(0.0226)
Home reading-environment index (SD)	Main estimate	[−0.0586, 0.0235]	[−0.0366, 0.0813]	[−0.0826, 0.0027]
		−0.0086	−0.0296	0.0210
		(0.0111)	(0.0198)	(0.0210)
	Pre-treatment placebo	[−0.0300, 0.0128]	[−0.0686, 0.0094]	[−0.0286, 0.0706]
Observations		190,528	319,256	509,784
		−0.0035	−0.0196	0.0161
		(0.0110)	(0.0101)	(0.0137)
Any paid private tutoring	Main estimate	[−0.0242, 0.0172]	[−0.0388, −0.0003]	[−0.0097, 0.0418]
		−0.0017	−0.0021	0.0004
		(0.0044)	(0.0078)	(0.0060)
	Pre-treatment placebo	[−0.0109, 0.0075]	[−0.0179, 0.0136]	[−0.0115, 0.0124]
Observations		177,530	305,840	483,370
		−0.0125	−0.0381	0.0256
		(0.0273)	(0.0329)	(0.0418)
Private tutoring expenditure (IHS, SD)	Main estimate	[−0.0640, 0.0391]	[−0.1081, 0.0318]	[−0.0658, 0.1171]
		−0.0077	−0.0085	0.0008
		(0.0101)	(0.0178)	(0.0149)
	Pre-treatment placebo	[−0.0289, 0.0134]	[−0.0496, 0.0326]	[−0.0301, 0.0316]
Observations		175,803	303,144	478,947

Notes: Maternal-schooling groups are defined by whether the child's mother ever attended school; observations with missing maternal-schooling information are excluded. The difference column reports mother never attended minus mother attended some school and is computed jointly to account for covariance between subgroup estimates. Each coefficient cell reports the estimate, state-clustered RIF standard error, and pointwise 95 percent Mammen wild-cluster-bootstrap confidence interval. None of the main private-tutoring difference intervals excludes zero, so the household educational-investment response is not clearly concentrated in either maternal-schooling group.

Table A4: Outcome availability in the baseline child sample

Survey year	Child records	Reading index	Private tutoring status	Private tutoring expenditure	School-input index
2012	184,195	180,066 (97.8%)	159,223 (86.4%)	–	166,240 (90.3%)
2014	174,799	171,003 (97.8%)	159,719 (91.4%)	158,540 (90.7%)	170,703 (97.7%)
2016	180,469	176,536 (97.8%)	167,809 (93.0%)	166,410 (92.2%)	176,591 (97.9%)
2018	176,438	173,684 (98.4%)	166,510 (94.4%)	164,548 (93.3%)	172,725 (97.9%)

Notes: Percentages are relative to the baseline child sample in the corresponding survey year. Private-tutoring status is missing for some children; these records are not coded as nonusers. Private-tutoring expenditure was not collected in 2012. In 2014–2018, recorded nonusers receive zero expenditure, while tutoring users with missing reported amounts are excluded. The school-input index cannot be constructed in village-years lacking sufficient component information.

Table A5: Supplementary event-study coefficients

Outcome	$T-2$	$T-1$	$T0$	$T+1$	$T+2$	N
	−0.0028	−0.0044	0.0017	−0.0073	−0.0443	
	(0.0076)	(0.0091)	(0.0109)	(0.0061)	(0.0154)	
Any paid private tutoring	[−0.0213, 0.0158]	[−0.0268, 0.0180]	[−0.0252, 0.0285]	[−0.0223, 0.0076]	[−0.0821, −0.0065]	653,261
	−0.0162	0.0045	−0.0105	−0.0120	−0.0954	
	(0.0298)	(0.0163)	(0.0238)	(0.0114)	(0.0366)	
Private tutoring expenditure (IHS, SD)	[−0.0859, 0.0535]	[−0.0335, 0.0426]	[−0.0662, 0.0452]	[−0.0387, 0.0147]	[−0.1809, −0.0099]	489,498
	−0.0374	0.0755	−0.0202	0.0484	−0.0146	
	(0.0223)	(0.0365)	(0.0159)	(0.0167)	(0.0213)	
Observable school-input index (SD)	[−0.1086, 0.0338]	[−0.0409, 0.1919]	[−0.0708, 0.0304]	[−0.0051, 0.1018]	[−0.0825, 0.0534]	686,259

Notes: Entries are HighBeef-minus-LowBeef event-time DDD estimates. Parentheses contain standard errors and brackets contain simultaneous 95 percent state-clustered Mammen wild-bootstrap confidence bands across the five displayed event times. Inference uses 999 repetitions.

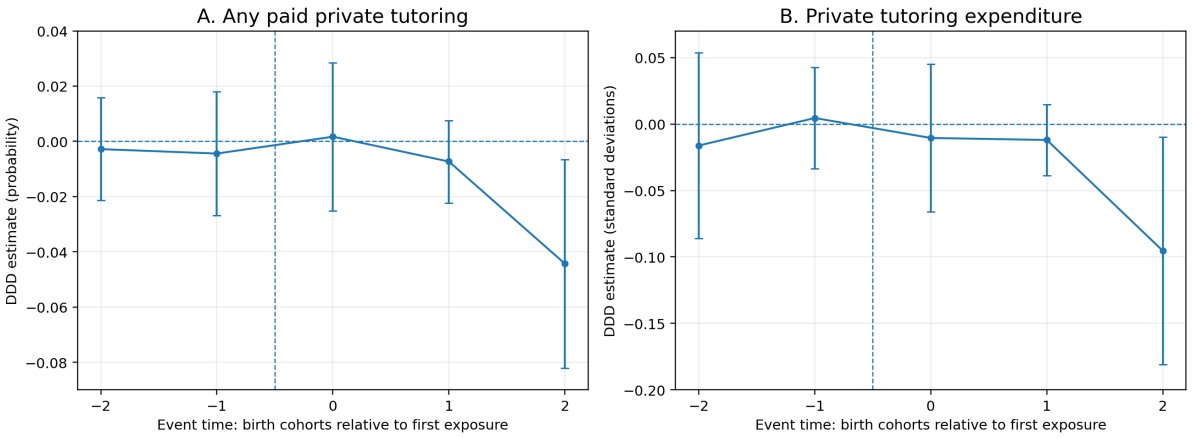


Figure A1: Supplementary private-tutoring responses

Notes: Points are HighBeef-minus-LowBeef event-time DDD estimates. Bars are simultaneous 95 percent state-clustered Mammen wild-bootstrap confidence bands based on 999 repetitions. Panel A reports the probability of any paid private tutoring; Panel B reports private-tutoring expenditure measured using the IHS transformation and standardized outcome.

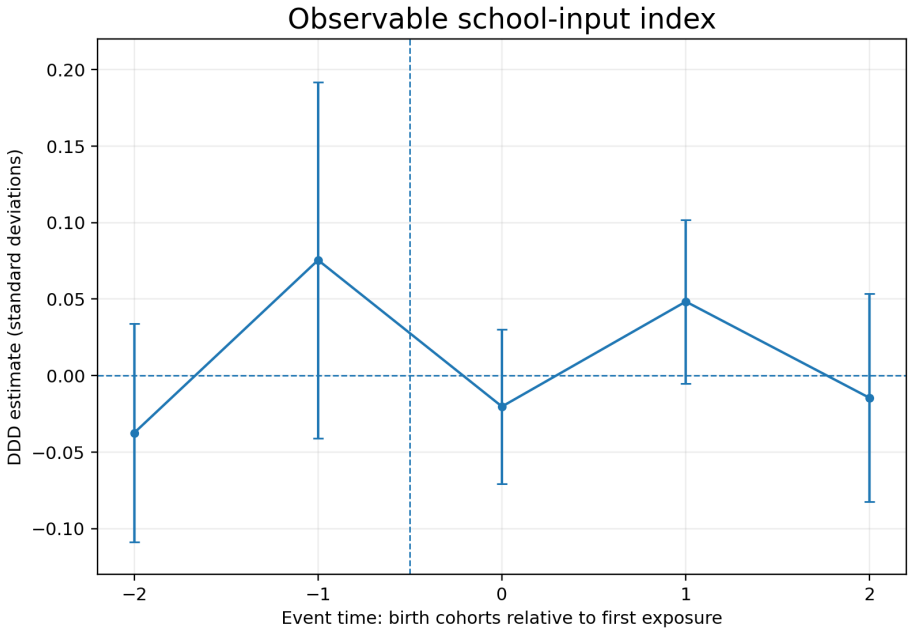


Figure A2: Observable school-input event study

Notes: Points are event-time HighBeef-minus-LowBeef DDD estimates for the observable school-input index. Bars are simultaneous 95 percent state-clustered Mammen wild-bootstrap confidence bands based on 999 repetitions.

Table A6: Webb-weight wild-bootstrap inference for the main learning estimates

Outcome	Period	HighBeef DD	LowBeef DD	DDD: High − Low	N (DDD)
Mathematics	Main estimate	−0.027 (0.022) [−0.076, 0.020]	0.048 (0.023) [−0.013, 0.092]	−0.075 (0.014) [−0.103, −0.043]	608,440
Mathematics	Pre-treatment placebo	−0.024 (0.024) [−0.077, 0.026]	−0.018 (0.018) [−0.065, 0.016]	−0.006 (0.010) [−0.028, 0.015]	608,440
Reading	Main estimate	−0.025 (0.017) [−0.063, 0.014]	0.040 (0.015) [0.005, 0.075]	−0.065 (0.015) [−0.095, −0.028]	609,658
Reading	Pre-treatment placebo	0.006 (0.012) [−0.023, 0.033]	−0.008 (0.021) [−0.059, 0.040]	0.014 (0.013) [−0.017, 0.042]	609,658

Notes: DDD point estimates match Table 1; the HighBeef and LowBeef component point estimates match Appendix Table A8. Parentheses report state-clustered recentered-influence-function standard errors. Brackets report pointwise 95 percent test-inverted wild-bootstrap confidence intervals obtained from `boottest`, using state clustering and bootstrap clustering, Webb's six-point weights, 9,999 repetitions and symmetric two-sided inference. The Webb procedure does not define a unique "Webb standard error," so the conventional clustered RIF standard error is reported rather than mislabeled as a bootstrap standard error. A separate higher-replication audit yields the same substantive conclusions.

Table A7: Webb-weight wild-bootstrap inference for the dynamic DDD estimates

Outcome	$T-2$	$T-1$	$T0$	$T+1$	$T+2$
	−0.024	0.012	0.020	−0.094	−0.152
	(0.026)	(0.019)	(0.012)	(0.012)	(0.025)
Mathematics	[−0.084, 0.031]	[−0.028, 0.055]	[−0.003, 0.049]	[−0.119, −0.067]	[−0.206, −0.095]
	0.023	0.006	0.002	−0.090	−0.107
	(0.013)	(0.028)	(0.013)	(0.016)	(0.024)
Reading	[−0.011, 0.051]	[−0.057, 0.069]	[−0.027, 0.033]	[−0.123, −0.054]	[−0.156, −0.050]
Joint Webb tests (p-values)					
Mathematics: pre-treatment coefficients jointly zero, $T-2=T-1=0$					0.762
Mathematics: post-treatment coefficients jointly zero, $T0=T+1=T+2=0$					<0.001
Mathematics: all five displayed event coefficients jointly zero					<0.001
Reading: pre-treatment coefficients jointly zero, $T-2=T-1=0$					0.309
Reading: post-treatment coefficients jointly zero, $T0=T+1=T+2=0$					0.003
Reading: all five displayed event coefficients jointly zero					0.011

Notes: Point estimates are identical to Figure 2. Parentheses report state-clustered RIF standard errors; brackets report pointwise 95 percent test-inverted Webb wild-bootstrap confidence intervals. Joint tests come from a stacked RIF regression that preserves within-state covariance across event times. These intervals are pointwise, not simultaneous; Figure 2 retains the baseline simultaneous Mammen bands.

Table A8: HighBeef and LowBeef components underlying the main learning DDD estimates

Outcome	Period	HighBeef DD	LowBeef DD	DDD: High − Low	N (HighBeef)	N (LowBeef)
Mathematics	Main estimate	−0.027 (0.019) [−0.068, 0.014]	0.048 (0.021) [0.008, 0.089]	−0.075 (0.012) [−0.102, −0.048]	291,871	316,569
Mathematics	Pre-treatment placebo	−0.024 (0.015) [−0.067, 0.019]	−0.018 (0.015) [−0.051, 0.016]	−0.006 (0.010) [−0.025, 0.013]	291,871	316,569
Reading	Main estimate	−0.025 (0.017) [−0.059, 0.008]	0.040 (0.015) [0.008, 0.071]	−0.065 (0.015) [−0.095, −0.035]	292,484	317,174
Reading	Pre-treatment placebo	0.006 (0.011) [−0.017, 0.030]	−0.008 (0.016) [−0.047, 0.031]	0.014 (0.013) [−0.012, 0.040]	292,484	317,174

Notes: This table reports the HighBeef and LowBeef component DDs underlying the DDD-only main Table 1. The main estimate averages event times 0, 1, and 2; the pre-treatment placebo averages event times −2 and −1. Each coefficient cell reports the estimate, state-clustered Mammen wild-bootstrap standard error in parentheses, and pointwise 95 percent wild-bootstrap confidence interval in brackets. Inference uses 999 repetitions and 19 state clusters. The component estimates show that the main DDD reflects a widening relative learning gap: LowBeef scores rise while the HighBeef declines are modest and individually imprecise.

Table A9: ASER reading and mathematics assessment levels

Level	Reading	Mathematics
1	Cannot read anything	Cannot do basic arithmetic
2	Identifies letters	Recognizes numbers 1–9
3	Reads words	Recognizes numbers 11–99
4	Reads a grade-1-level sentence	Performs two-digit subtraction
5	Reads a grade-2-level paragraph	Performs division

Notes: ASER assesses children using ordered reading and mathematics tasks. The main analysis converts these five skill levels to age-standardized z-scores, so mathematics and reading outcomes are expressed in standard-deviation units relative to children of the same age.

Table A10: Summary statistics for the main analysis sample

Variable	Mean	SD	Minimum	Maximum	*N*
Panel A. Children, learning outcomes, and exposure					
Mathematics z-score	0.049	0.971	−2.507	4.158	608,440
Reading z-score	0.028	0.969	−2.702	4.149	609,658
Child age (years)	8.62	2.28	5	12	715,901
Female child	0.483	0.500	0	1	715,901
High beef-exposure district	0.478	0.499	0	1	715,901
Panel B. School participation and maternal schooling					
Currently enrolled	0.983	0.128	0	1	715,901
Dropped out	0.007	0.081	0	1	715,901
Never enrolled	0.010	0.100	0	1	715,901
Mother ever attended school	0.625	0.484	0	1	519,868
Panel C. Household educational investments and school inputs					
Home reading-environment index (SD)	−0.028	0.972	−0.611	4.148	701,289
Any paid private tutoring	0.272	0.445	0	1	653,261
Private tutoring expenditure (IHS, SD)	−0.024	0.985	−0.622	2.997	489,498
Observable school-input index (SD)	−0.015	1.008	−4.325	6.775	686,259

Notes: The sample is the main analysis sample used in Table 1: children ages 5–12 in Jharkhand, Gujarat, and never-treated states, born in 2000–2013. Always-treated states are excluded, and each retained state has both high- and low-exposure districts. Statistics are unweighted, and *N* is variable-specific. Variables reported here are descriptive; their inclusion does not imply their use as regression covariates. The main specifications include no covariates. Mathematics and reading scores are standardized by child age. The home reading-environment, private-tutoring expenditure, and observable school-input indices follow the definitions used in the corresponding analyses. A district is classified as high beef exposure when its likely beef-consumer share exceeds the unweighted district mean within its state.

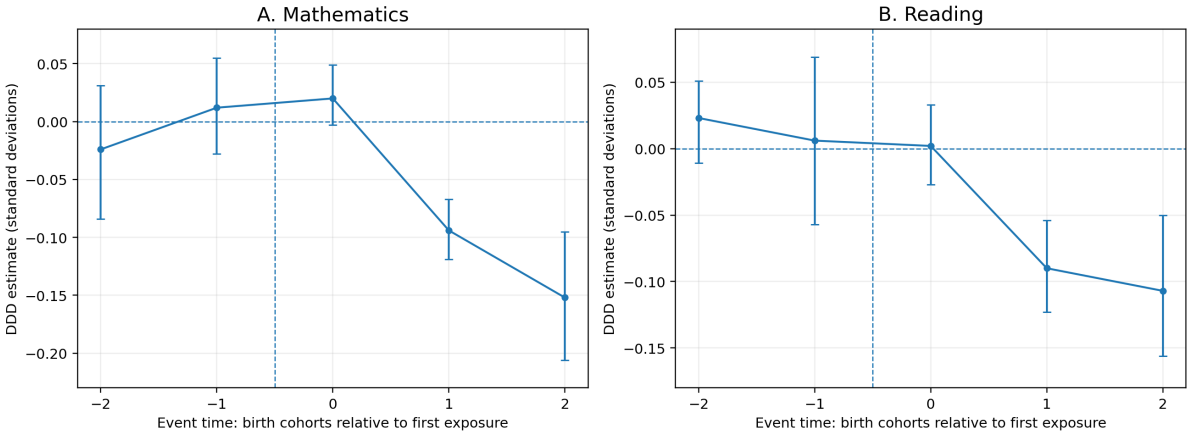


Figure A3: Dynamic effects under Webb-weight wild-bootstrap inference

Notes: Points are the baseline event-time HighBeef-minus-LowBeef DDD estimates. Bars are pointwise 95 percent test-inverted confidence intervals from `boottest`, using state clustering and bootstrap clustering, Webb's six-point weights, 9,999 repetitions. The vertical dashed line separates pre-treatment event times from the first exposed cohort. These intervals are not simultaneous; the main Figure 2 reports the baseline simultaneous Mammen wild-bootstrap bands.

Table A11: Cohort-representation diagnostic

Period	DDD estimate	Bootstrap SE	95% CI	Cohort cells	States
Post-treatment average ($T0$ to $T+2$)	−0.041	0.024	[−0.163, 0.082]	6,506	19
Pre-treatment placebo ($T-2$ to $T-1$)	0.016	0.018	[−0.077, 0.108]	6,506	19

Notes: The outcome is the log number of observed ASER children in each district-by-birth-cohort cell. The sample and design follow Table 1: children ages 5–12 in Jharkhand, Gujarat, and never-treated states; always-treated states are excluded. Jharkhand and Gujarat are first treated in the 2005 and 2011 birth cohorts. Districts are classified as high or low beef exposure relative to the unweighted district mean within each state. Entries report high-minus-low exposure DDD estimates. The post-treatment estimate averages event times 0, 1, and 2; the placebo averages event times −2 and −1. Standard errors and confidence intervals use a state-clustered Mammen wild bootstrap with 999 repetitions. The analysis contains one observation per nonempty district-by-birth-cohort cell; it therefore tests for relative changes in observed cohort representation, not population cohort size.

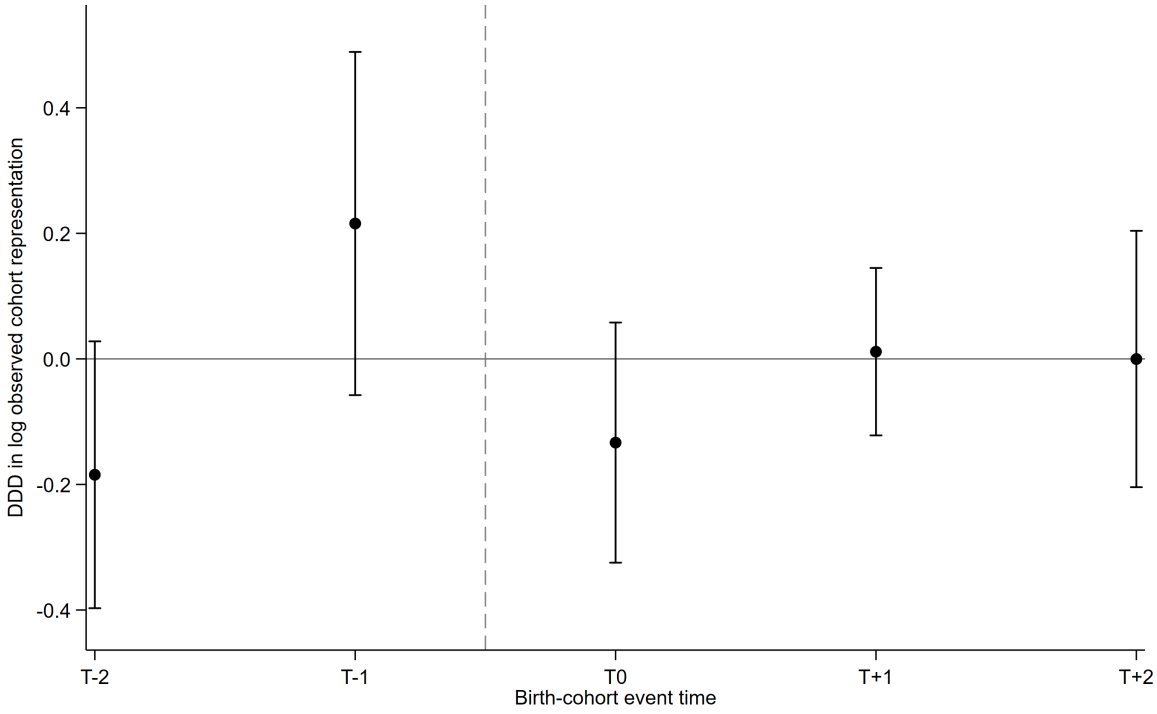


Figure A4: Dynamic cohort-representation diagnostic

Notes: Points are event-time high-minus-low exposure DDD estimates with the log number of observed ASER children in each district-by-birth-cohort cell as the outcome. Bars are simultaneous 95 percent state-clustered Mammen wild-bootstrap confidence bands based on 999 repetitions. The vertical dashed line separates pre-treatment event times from the first treated birth cohort. All five simultaneous confidence bands include zero. The outcome measures representation in the ASER sample, not population cohort size.

Table A12: Maternal-schooling composition diagnostic

Estimand	DDD estimate	Bootstrap SE	95% CI	*N*	States
Post-treatment average ($T0$ to $T+2$)	−0.010	0.015	[−0.059, 0.039]	519,868	19
Pre-treatment placebo ($T-2$ to $T-1$)	−0.009	0.006	[−0.029, 0.012]	519,868	19
Post-treatment minus placebo average	−0.001	0.014	[−0.061, 0.058]	519,868	19

Notes: The outcome equals one when the child's mother ever attended school. The sample and design follow Table 1, restricted to observations with maternal-schooling information. Entries report high-minus-low exposure DDD estimates. The post-treatment estimate averages event times 0, 1, and 2; the placebo averages event times −2 and −1. The final row is a paired influence-function contrast between these averages and accounts for their covariance. Because the averages contain different event-time cohorts, this contrast is not a conventional before-and-after comparison for the same cohort. Standard errors and confidence intervals use a state-clustered Mammen wild bootstrap with 999 repetitions. The main specifications include no covariates or survey weights.

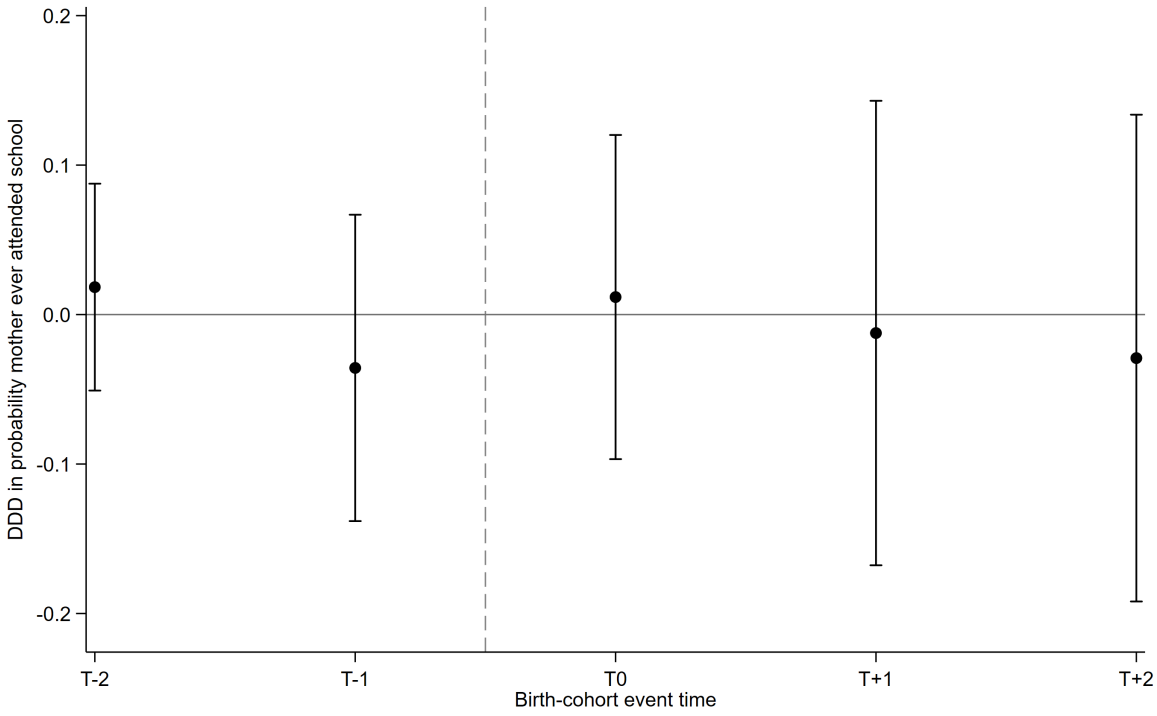


Figure A5: Dynamic maternal-schooling composition diagnostic

Notes: Points are event-time high-minus-low exposure DDD estimates for an indicator that the child's mother ever attended school. Bars are simultaneous 95 percent state-clustered Mammen wild-bootstrap confidence bands based on 999 repetitions. The vertical dashed line separates pre-treatment event times from the first treated birth cohort. All five simultaneous confidence bands include zero.

Table A13: Household-wealth diagnostic

Estimand	DDD estimate	Bootstrap SE	95% CI	*N*	States
Post-treatment average ($T0$ to $T+2$)	−0.049	0.023	[−0.100, 0.002]	665,407	19
Pre-treatment placebo ($T-2$ to $T-1$)	−0.053	0.009	[−0.074, −0.033]	665,407	19
Post-treatment minus placebo average	0.005	0.025	[−0.046, 0.055]	665,407	19

Notes: The outcome is the household wealth index recorded in ASER. The sample and design follow Table 1, restricted to observations with nonmissing wealth information. Entries report high-minus-low exposure DDD estimates. The post-treatment estimate averages event times 0, 1, and 2; the placebo averages event times −2 and −1. The final row is a paired influence-function contrast between these averages and accounts for their covariance. Because the averages contain different event-time cohorts, this contrast is not a conventional before-and-after comparison for the same cohort. Standard errors and confidence intervals use a state-clustered Mammen wild bootstrap with 999 repetitions. The significant placebo differential prevents a causal interpretation of the wealth estimates.

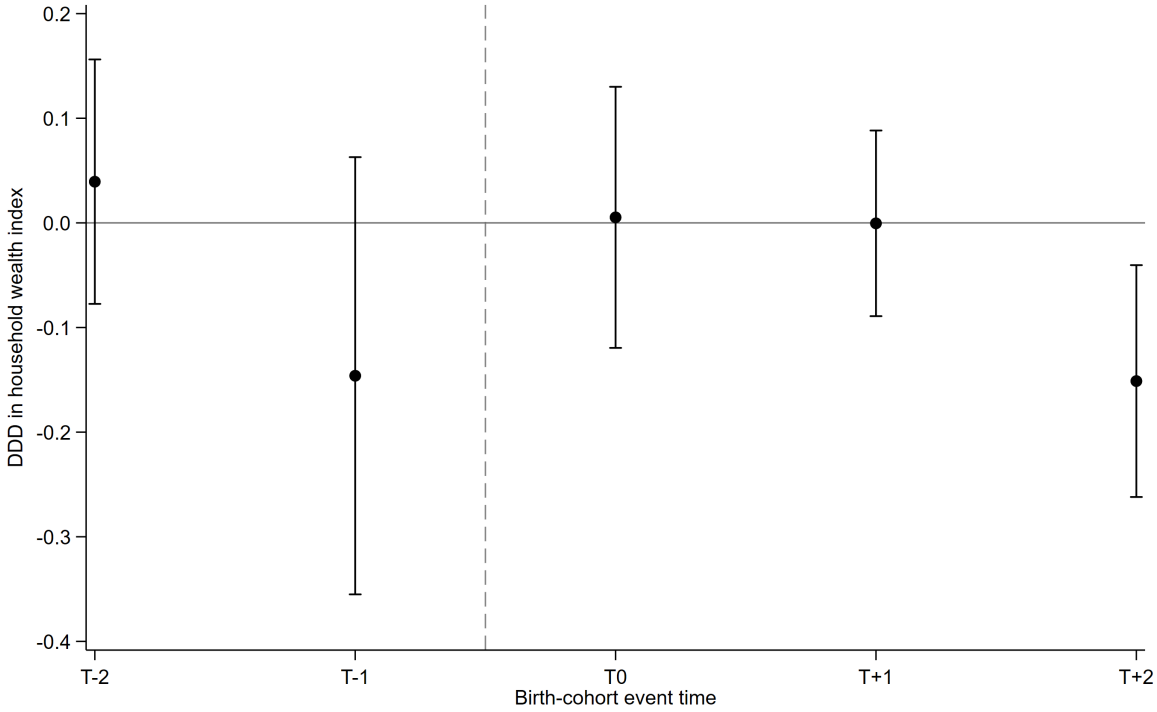


Figure A6: Dynamic household-wealth diagnostic

Notes: Points are event-time high-minus-low exposure DDD estimates for the household wealth index. Bars are simultaneous 95 percent state-clustered Mammen wild-bootstrap confidence bands based on 999 repetitions. The vertical dashed line separates pre-treatment event times from the first treated birth cohort. The $T + 2$ interval excludes zero, but the significant average pre-treatment differential and uneven event-time pattern prevent a causal interpretation.