OXFORD JOURNAL
OF STUDENT
SCHOLARSHIP

The Opportunity Cost of Corruption in Public Resource Allocation: Evidence from India's Education Sector, 2016-2022

Vivan Balhara
vivanbalhara2009@gmail.com

ABSTRACT

India allocates close to 3 percent of GDP annually to education, yet foundational learning deficits remain widespread: the 2022 ASER survey found 42.8 percent of rural Class 5 students could not read a Class 2-level text. This paper asks whether state-level prosecutorial impunity under the Prevention of Corruption Act is associated with that gap, using a panel of 15 major Indian states over 2016-2022. Impunity is defined as one minus the ratio of PCA convictions to registered cases, from NCRB Crime in India reports. Two-way fixed effects regressions, controlling for per capita income and pupil-teacher ratio, estimate that a one-unit increase in impunity is associated with 6.3 additional points of reading deficiency. This estimate is not statistically significant ($p = 0.107$). The enrolment coefficient has the expected sign but is also insignificant. The dropout coefficient has the wrong sign, which the paper attributes to COVID-year reporting zeros and floor effects. As an illustration only, applying the reading coefficient to UDISE+ enrolment in three high-impunity states gives an order-of-magnitude figure of roughly 1 million students. This is an associational counterfactual rather than a causal prediction, and it inherits the coefficient's uncertainty. The results are best read as directional evidence from a small panel with a single, imperfect corruption proxy.

Keywords: corruption; education outcomes; fixed effects; NCRB; India; impunity; public resource allocation

1. INTRODUCTION

In 2022, Pratham's Annual Status of Education Report found that 42.8 percent of children enrolled in Class 5 across rural India could not read a text designed for Class 2, down from 50.5 percent in 2018, a decline driven substantially by pandemic-related school closures and the disruption to learning they caused. India's school education system enrols more than 265 million children from the pre-primary through higher secondary stages, of which elementary education (Classes 1 to 8) alone accounts for close to 189 million, and the country draws on combined central and state public expenditure equivalent to between 2.7 and 2.9 percent of GDP each year on education, according to successive editions of the

Economic Survey of India. That level of spending is not trivial, which makes the persistence of poor learning outcomes harder to explain away.

A natural question follows: why does substantial public spending on education fail to produce proportionate improvements in learning? The macroeconomic literature identifies several candidate answers: teacher supply constraints, household income, social inequality, and linguistic diversity. This paper investigates one that receives considerable qualitative attention but limited panel scrutiny in the Indian context: corruption in public resource allocation.

The intuition is relatively direct. Elementary education in India is primarily state-funded and state-administered. Resources flow from central government transfers and state budgets through district and block administrative offices to individual schools, and at each node of that chain there is opportunity for diversion. Teachers who are absent but remain on payroll consume salary expenditure without delivering instruction; infrastructure contracts that are inflated or left unexecuted absorb capital without improving classrooms; and underprovided textbooks and mid-day meals amount to budget allocations that reached an administrator rather than a student. Where the institutional machinery for detecting and penalising these diversions is weak, they tend to persist, and where it functions, they are at least partially constrained.

This paper uses a state-level corruption impunity score, constructed from Prevention of Corruption Act conviction and registration data in NCRB Crime in India reports, to assess whether within-state variation in prosecutorial accountability over 2016-2022 predicts elementary education outcomes. Three outcome variables are examined: reading deficiency rates from ASER surveys, school dropout rates, and gross enrolment ratios from UDISE+ Flash Statistics.

The results below are reported as they came out of the regressions, without dressing up the interpretation. The two-way fixed effects estimator, while appropriate for this setting, is conservative by construction: it absorbs all time-invariant cross-state variation through state fixed effects, leaving only within-state variation in impunity over time to explain within-state variation in outcomes. With fifteen states and seven years, degrees of freedom are limited. The headline reading deficiency result, a coefficient of 6.3 with $p = 0.107$, does not clear the conventional 5 percent threshold. The dropout coefficient has the wrong sign, a consequence of COVID-related data artifacts examined in detail below. The GER result is directionally correct but also statistically insignificant.

What comes out of this exercise is three things, best read together rather than in isolation. There is a reproducible panel dataset that combines administrative sources not previously used together in this way. There is regression evidence that is directionally consistent, alongside a cross-sectional pattern stark enough to be informative even without a causal interpretation. And there is an opportunity cost framework that turns the coefficient into a headcount of students, with the caveats stated plainly rather than left in a footnote.

The rest of the paper proceeds as follows. Section 2 reviews the relevant literature. Section 3 develops the theoretical mechanisms. Section 4 describes the data. Section 5 presents the methodology. Section 6

reports results. Section 7 estimates the opportunity cost. Section 8 presents robustness checks. Section 9 states limitations. Section 10 concludes.

2. LITERATURE REVIEW

The relationship between corruption and economic performance has been studied rigorously since Mauro's (1995) cross-country analysis, which showed that higher perceived corruption is associated with lower investment ratios and slower GDP growth. The paper used ethno-linguistic fractionalization as an instrument for perceived corruption, a creative if controversial identification strategy, and its directional finding has held across a range of subsequent methodological approaches. For the purposes of this paper, the relevant extension runs from corruption's aggregate effect on growth to its specific effect on public service delivery outcomes.

Gupta, Davoodi, and Tiongson (2000), in an IMF Working Paper examining a broad cross-country sample, found that higher corruption is associated not only with lower public expenditure on education as a share of GDP but also with worse outcomes conditional on that spending: higher dropout rates, lower literacy, worse health indicators. The key contribution of their analysis is the distinction between quantity and quality of public expenditure: corrupt governments do tend to spend less on education, but even the resources they do allocate deliver lower educational returns. This distinction matters for the Indian case because India's elementary education expenditure is not obviously deficient in quantity; the question is what those resources actually purchase at the school level.

The most direct empirical predecessor of this paper is Reinikka and Svensson's (2004) study of Uganda's primary school capitation grant programme. Using a newspaper public information campaign as quasi-experimental variation, they demonstrated that schools in districts with better access to information about their grant entitlements captured a substantially larger share of those grants. Their headline finding, that Ugandan schools received on average only 13 percent of the capitation grant to which they were entitled, became one of the most cited pieces of evidence for the magnitude of corruption's effect on public education funding. The design is not replicable in the Indian context studied here, since no analogous information shock occurs within the 2016-2022 study window, but the theoretical mechanism it validates is directly applicable: funds allocated at the centre do not reliably reach the schools they are intended for.

Teacher absenteeism is a related but distinct channel, with its own empirical literature. Chaudhury, Hammer, Kremer, Muralidharan, and Rogers (2006) conducted unannounced visits to primary schools across six developing countries and found an average teacher absence rate of approximately 19 percent across the sample, with India recording one of the higher country-level rates at roughly one-quarter of government primary teachers absent on a given day. The India-specific evidence was extended and deepened by Muralidharan, Das, Holla, and Mohpal (2017), who used a later nationally representative panel and estimated the salary cost of unauthorized teacher absence in India at approximately US$1.5 billion annually at 2010 exchange rates. They document that high-absence schools are disproportionately

concentrated in districts with weak administrative accountability structures. Their work provides the empirical foundation for the teacher absenteeism mechanism this paper invokes in Section 3.

For India specifically, the existing quantitative literature on corruption and education is largely composed of single-state studies or relies on perceived corruption indices with insufficient panel coverage. The Comptroller and Auditor General produces education department audit reports that document specific irregularities: ghost beneficiaries in mid-day meal schemes, construction payments for work not completed, and teacher recruitment irregularities. These reports are state-specific and inconsistently formatted across states and years, which is why the CAG audit trail could not be operationalized as a panel variable here. They serve in this paper as qualitative case evidence for the mechanisms invoked theoretically.

The contribution of this paper to the existing literature is narrow but specific. Cross-country studies lose the institutional specificity of India's federal structure: the variation in how state governments administer anti-corruption institutions, how actively they register and prosecute PCA cases, and how that variation has evolved over time. Single-state studies capture institutional detail but cannot isolate the effect of corruption from the many other ways Indian states differ from each other. A within-India panel exploiting variation across 15 states over 7 years, with state and year fixed effects, addresses the cross-state omitted variable problem that cross-sectional analysis cannot handle and generalises in a way that single-state work cannot. That is this paper's methodological contribution, and a modest one at that.

3. THEORETICAL FRAMEWORK

Three mechanisms plausibly link prosecutorial impunity under the Prevention of Corruption Act to elementary education outcomes. They operate at different points in the resource allocation chain and are not mutually exclusive.

Teacher absenteeism. Teacher salaries account for the large majority of recurring elementary education expenditure in India. Ghost teachers, meaning those on payroll who do not exist, and chronically absent teachers amount to a diversion of salary funds that produces zero instructional output. Maintaining this arrangement over time requires either active connivance by administrative supervisors or at minimum their unwillingness to investigate. In states where the institutional apparatus for detecting and penalizing corruption is weak, as proxied here by consistent failure to secure PCA convictions, the expected cost of maintaining such arrangements is low. One would therefore expect states with higher prosecutorial impunity to have higher effective teacher absence, lower instructional time, and as a consequence, worse learning outcomes. The reading deficiency outcome variable captures this most directly: it measures whether students have achieved the most basic threshold of literacy despite several years of enrollment.

Fund diversion at the transfer chain. A substantial portion of elementary education funding flows through scheme-based transfers: the Samagra Shiksha Abhiyan, the mid-day meal programme, free textbook distribution, and school infrastructure grants under the Rashtriya Madhyamik Shiksha Abhiyan. Each

transfer involves procurement decisions and funds moving through multiple administrative layers. Corruption at the block or district level can manifest as inflated procurement bills, payment for goods not delivered, or straightforward misappropriation of scheme funds. The educational impact depends on what is diverted: diversion of mid-day meal funds affects both nutrition and the incentive to attend school, diversion from infrastructure grants leaves schools without classrooms, toilets, or drinking water, and diversion of textbook funds simply deprives students of materials. These mechanisms affect both enrollment (captured in GER) and learning (captured in reading deficiency). The 2004 capitation grant evidence from Uganda (Reinikka and Svensson, 2004) provides the most rigorous cross-national evidence for how large these diversions can be; CAG audit reports for Indian states document that similar patterns occur domestically, even if not in a form that can be systematically quantified across states.

Procurement rent extraction in school construction. Capital expenditure in elementary education, covering classrooms, boundary walls, and separate toilet facilities for girls, is administered through state public works departments or dedicated construction agencies. Construction contracts are subject to the familiar problems of procurement corruption: bid rigging, contractor collusion, payments for substandard work, and certification of incomplete projects. The connection to educational outcomes is somewhat indirect: poor infrastructure does not by itself prevent teaching, but it affects student attendance, particularly girls' attendance given evidence that adequate toilet provision significantly increases female enrollment and retention in Indian schools. Gross Enrolment Ratio is most likely sensitive to this channel.

What makes all three mechanisms jointly susceptible to a single institutional measure is that each requires the persistent failure of investigative and prosecutorial accountability. A functional enforcement apparatus that actively registers PCA cases and secures convictions raises the expected cost of all three practices. The NCRB impunity score captures the long-run failure of this apparatus at the state level. In states where cases are registered but convictions rarely follow, the deterrent effect of anti-corruption law is minimal regardless of how many cases enter the pipeline.

4. DATA AND VARIABLES

4.1 SAMPLE AND STUDY PERIOD

The study covers 15 major Indian states over 2016-2022, yielding a maximum of 105 state-year observations. That maximum is not achieved in practice; missing and excluded observations documented in Section 4.2 reduce the usable panel. Summing Census 2011 population figures for the 15 sampled states against the all-India total of approximately 1.21 billion (Office of the Registrar General & Census Commissioner, 2011) gives a combined share of roughly 86 percent, calculated using Telangana's post-bifurcation population as derived from undivided Andhra Pradesh's 2011 figures, consistent with the approximation already flagged for Telangana's urbanisation share in Section 9. The sample spans a wide range of institutional quality, geographic size, and economic development, making it a reasonable if not exhaustive basis for within-India inference.

No prior classification of states into corruption tiers is imposed; states are compared only on their measured impunity scores. One state requires an immediate note. Madhya Pradesh's mean impunity score of 0.466 is the lowest in the sample, but this reflects anomalously high Anti-Corruption Bureau conviction rates in 2017-2019. These are likely an artifact of courts clearing a pre-existing backlog rather than a sign of institutionally robust enforcement. This anomaly is discussed as a limitation in Section 9 and flagged in all tables.

The study period begins in 2016 for a specific data reason. NCRB Crime in India Chapter 8C, which presents state-wise Prevention of Corruption Act conviction and registration data, adopts a consistent and comparable table format from 2016 onward. The 2014 and 2015 reports used a different chapter structure that makes longitudinal comparison problematic. 2016 is therefore the first year of clean, usable data.

4.2 CORRUPTION MEASURE : NCRB IMPUNITY SCORE

The primary corruption measure is constructed from NCRB Crime in India, Volume II, Tables 8C.1 and 8C.4, for 2016-2022. The impunity score for state i in year t is:

$$ncrb(i,t) = 1 - [Convictions(i,t) / Cases\ Registered(i,t)]$$

A score approaching 1.0 indicates near-complete prosecutorial failure: cases are registered but almost none result in conviction. A score approaching 0 indicates convictions keep pace with registrations. The theoretical range is [0, 1] for positive case volumes; negative values occur when courts convict on backlog from prior years, and these years are excluded from the panel (see below).

Table 1 presents the full state-year impunity matrix.

Table 1: NCRB Impunity Scores by State and Year, 2016-2022

State	2016	2017	2018	2019	2020	2021	2022	Mean
Uttar Pradesh	0.833	0.948	0.857	0.918	0.984	0.866	0.832	0.891
Bihar	0.890	0.750	0.899	0.944	1.000	0.975	0.954	0.916
Madhya Pradesh	0.659	0.129	0.125	0.233	0.854	0.636	0.626	0.466
Rajasthan	0.806	0.748	0.659	0.658	0.846	0.890	0.787	0.770
Odisha	0.854	0.759	0.759	0.810	0.935	0.860	0.686	0.809
Jharkhand	1.000	0.977	0.789	0.701	0.762	0.982	0.776	0.855
West Bengal	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000
Chhattisgarh	0.459	—	—	—	—	1.000	0.692	0.717*

Assam	0.870	0.926	0.960	0.889	0.875	0.875	0.965	0.909
Telangana	0.753	0.455	0.698	0.944	0.917	0.928	0.925	0.803
Kerala	—	0.803	0.803	0.587	—	0.484	0.523	0.640*
Tamil Nadu	0.559	0.611	0.583	0.778	0.796	0.853	0.671	0.693
Gujarat	0.872	0.547	0.721	0.922	0.889	0.838	0.801	0.799
Maharashtra	0.910	0.936	0.940	0.938	0.985	0.977	0.957	0.949
Karnataka	—	0.820	0.892	0.868	0.865	0.894	0.846	0.864*

Source: NCRB Crime in India, Vol. II, Tables 8C.1 and 8C.4, 2016-2022. Notes: — indicates year excluded due to data anomaly (see text). * Mean computed over available years only.

Three states required year-specific exclusions. Karnataka's 2016 value is excluded because only 25 PCA cases were registered that year, against 401 in 2014 and 289 in 2017, indicating a source data error in the published report. Kerala's 2016 and 2020 values are excluded: the 2016 figure reflects anomalously high case registrations (430 in a single year) that distort the score, while the 2020 figure is missing from the published NCRB report with no explanation provided. Chhattisgarh's 2017-2020 values are excluded because court convictions exceeded registered cases in those years, producing negative nominal impunity values that reflect courts clearing pre-2017 backlogs rather than current enforcement activity.

West Bengal requires a separate note, because a score of 1.000 that holds unchanged across all seven years might look like a data error rather than a real finding. The West Bengal Lokayukta Act, 2003, established the office in name, but the West Bengal Lokayukta (Amendment) Act, 2018, narrowed its practical reach further: it excluded the Chief Minister from the institution's investigative purview and required prior state government approval before any complaint against a public servant could be investigated, restrictions that limit the office's independence relative to Lokayuktas in states such as Kerala or Karnataka. Against this institutional backdrop, annual PCA case registrations in West Bengal were minimal, between 8 and 25 cases per year across the full study window, with zero convictions recorded in each year. The 1.000 score accurately captures a state of sustained prosecutorial inaction and is retained in the panel as an extreme case of institutional weakness. Because West Bengal's impunity score never changes, it contributes no within-state variation to the fixed effects estimator and its inclusion does not affect the regression coefficients; it does, however, anchor the cross-sectional descriptive comparisons.

4.3 EDUCATION OUTCOME VARIABLES

Three outcome variables are used, each capturing a distinct dimension of elementary education performance.

Reading deficiency is defined as the percentage of Class 5 students unable to read a Class 2-level text. Data come from ASER (Annual Status of Education Report) surveys published by Pratham. ASER uses household-based sampling methodology, which means it captures out-of-school children as well as enrolled ones, an advantage over school-based surveys that systematically miss the most disadvantaged

students. Within the 2016-2022 study period, ASER conducted full national rural surveys in 2018 and 2022 only. The reading variable therefore has at most two observations per state, a severe constraint on statistical power discussed explicitly in the limitations.

School dropout is the percentage of enrolled elementary students (Classes 1-8) who leave before completing the elementary cycle, sourced from UDISE+ Flash Statistics for 2018-19 through 2022-23. UDISE+ provides annual state-level data with consistent variable definitions across the study period. A critical data quality problem affects 2020-21: dropout rates for that year are zero or near-zero for most states because schools were closed and administrative data collection was disrupted under COVID. These zeros are administrative placeholders rather than genuine measurements and create a downward artifact in the within-state correlation between impunity and dropout. The problem is tested in robustness checks and discussed as a limitation.

Gross Enrolment Ratio at the elementary level is also drawn from UDISE+ Flash Statistics, 2018-19 to 2022-23. GER measures total enrollment as a proportion of the age-eligible population. Five states in the sample report mean GER exceeding 100, namely West Bengal, Assam, Telangana, Karnataka, and Maharashtra, reflecting over-age student retention and cross-border enrollment flows. These values are genuine features of the administrative data and are not data errors, but they complicate regression interpretation and are discussed as a limitation. Net Enrolment Ratio, which is bounded between 0 and 100 and available in UDISE+, would be a cleaner alternative in future extensions of this work.

4.4 CONTROL VARIABLES

Per capita NSDP at constant 2011-12 prices is drawn from the RBI Handbook of Statistics on Indian Economy. This controls for income effects on education outcomes: wealthier states can fund better teachers, more materials, and better infrastructure independent of whatever corruption-related dynamics are also at work.

Urbanisation is measured as the proportion of state population residing in urban areas, using Census of India 2011 figures (Office of the Registrar General & Census Commissioner, 2011) applied statically across all years. This variable is time-invariant within the dataset. In the two-way fixed effects specification, it is automatically absorbed by state fixed effects and does not enter the estimating equation directly; it functions as a control only in pooled OLS.

Pupil-teacher ratio at the primary level (Classes 1-5) comes from UDISE+ Flash Statistics. Including PTR ensures that the NCRB coefficient does not pick up the effect of teacher shortages on learning outcomes, a channel distinct from corruption, namely insufficient teacher supply, as opposed to misallocation of the teachers who are nominally present.

4.5 VARIABLES EVALUATED AND DROPPED

Several variables considered in the original design were evaluated and dropped before estimation.

CAG irregularity data: Comptroller and Auditor General compliance audit reports document specific financial irregularities in state education departments but are not published in a consistent numerical format across all 15 states in every year. They cannot be transformed into a panel variable without imposing arbitrary discretion on what counts as an "irregularity unit." CAG findings are used qualitatively in Section 3 but not quantitatively.

RTI rejection rates: Right to Information compliance data for state School Education Departments is maintained by State Information Commissions rather than the Central Information Commission. Coverage and format vary significantly across states, and several states, including Tamil Nadu, Andhra Pradesh, and Jharkhand, had overlapping or fragmented SIC jurisdictions during the study period. A consistent cross-state panel series was not constructable.

Lokayukta complaint backlog: Annual pending complaint and resolution rate data is published by Lokayukta institutions in inconsistent formats. Kerala, Karnataka, and Madhya Pradesh have reasonably active and documented Lokayukta offices. Bihar and Uttar Pradesh lacked functional Lokayukta institutions or did not publish systematic data during large portions of the study period, while West Bengal's Lokayukta, though nominally established, operated under statutory restrictions, including the 2018 amendment narrowing its jurisdiction over complaints against senior officials, that left it without a meaningful enforcement record to draw on.

Transparency International sub-national scores: TI's state-level corruption surveys are not annual instruments. At most one or two observations per state exist within the 2016-2022 window, making panel inclusion impossible.

The failure to build a composite corruption index from multiple measures reflects the data constraints documented above rather than an oversight in the research design. NCRB alone is therefore used as a single proxy, and its measurement error has two components that bias the estimate in different ways. The first is timing noise: convictions in a given year often relate to cases registered years earlier (Section 9, conviction lag), which adds error plausibly unrelated to true corruption. Error of this kind attenuates the coefficient toward zero. The second is systematic: states with more active anti-corruption bureaus register more cases, which raises measured impunity where enforcement is actually stronger (Section 9, registration paradox). This error is correlated with the institutional quality the score is meant to capture, so the standard attenuation result does not apply to it. Within a state, a rise in registrations biases the coefficient toward zero if it reflects greater enforcement effort, but not if it reflects more underlying corruption, and the data cannot distinguish the two. Because the two components cannot be separated, this paper does not claim that the reported coefficient understates the true relationship.

5. METHODOLOGY

5.1 BASELINE OLS SPECIFICATION

As a baseline, a pooled OLS regression is estimated separately for each outcome variable:

$$Y(i,t) = a + b * ncrb(i,t) + c * ndsp(i,t) + d * ptr(i,t) + e * urban(i) + u(i,t)$$

where Y(i,t) is the outcome for state i in year t, ncrb(i,t) is the impunity score, ndsp(i,t) is the log of per capita NSDP, ptr(i,t) is the primary pupil-teacher ratio, and urban(i) is the time-invariant 2011 urbanisation share. OLS pools all variation, both within-state over time and across states, into a single coefficient. The limitation is well-known: states differ on many dimensions that correlate with both corruption and education outcomes. Richer states systematically tend toward lower impunity scores and better outcomes. The OLS coefficient for ncrb therefore conflates the income effect with the corruption effect and is reported only as a comparator to the fixed effects estimates.

5.2 TWO WAY FIXED EFFECTS SPECIFICATIONS

The main specification is a two-way fixed effects model:

$$Y(i,t) = a(i) + l(t) + b * ncrb(i,t) + c * ndsp(i,t) + d * ptr(i,t) + e(i,t)$$

State fixed effects a(i) absorb all time-invariant differences between states: historical institutional quality, geographic endowments, cultural factors, and any stable features of each state's political economy that might simultaneously predict both corruption and education outcomes. Year fixed effects l(t) absorb shocks common to all states in a given year, such as national policy changes, macroeconomic cycles, and central government education scheme modifications. What remains in the error term is idiosyncratic variation, and what the coefficient b identifies is the within-state relationship: when a state's impunity score rises or falls over 2016-2022, does its education outcome move in the predicted direction, after removing the national year trend?

This estimator is conservative precisely because it removes so much variation. For states with relatively stable impunity scores across the seven years, which describes most states in the sample since institutional change is slow, the FE estimator has limited variation to work with. West Bengal's invariant score of 1.000 contributes nothing to within-state identification. This is the trade-off of the design: it sacrifices statistical power in exchange for protection against cross-state omitted variable bias.

Estimation is carried out in R 4.6.0 using the plm package for panel models. The key command is:

```
plm(outcome ~ ncrb + ndsp + ptr, data=pdata, model='within', effect='twoways')
```

5.3 STANDARD ERRORS

Standard errors are clustered at the state level using the HC1 heteroskedasticity-consistent estimator, implemented via coeftest() with vcovHC(type='HC1', cluster='group'). Clustering at the state level accounts for serial correlation within states across years, which is to be expected when the same institutional and political features affect outcomes in successive years. Cameron and Miller (2015) note that cluster-robust inference with fewer than 30 clusters may be liberal, meaning confidence intervals may

be too narrow, and 15 states falls well below that threshold. The standard errors reported here should therefore be treated as approximate rather than exact; if anything, they likely understate the true uncertainty.

5.4 WHY DIFFERENCE-IN-DIFFERENCES WAS NOT PURSUED

Difference-in-Differences was the planned identification strategy, intended to exploit Lokayukta institutional reform timing as quasi-experimental variation. The design requires a well-defined reform event, such as passage of legislation, restoration of powers, or appointment of an active Lokayukta, occurring at different times in different states so that treated states can be compared against controls. No qualifying events occurred within the 2016-2022 study window. Karnataka, the most active Lokayukta state in India, had its institution's powers significantly curtailed in 2016 and only partially restored by 2019. Madhya Pradesh had a Lokayukta but its operational effectiveness fluctuated without a clean break. Bihar lacked a functional Lokayukta institution throughout, and West Bengal's Lokayukta operated under jurisdictional restrictions, formalized further by the 2018 amendment discussed in Section 4.2, that kept it from generating a usable reform event. The DiD design is kept in this section because it was the original plan, and the reason it could not be executed says something substantive about the institutional landscape of the period: no state in the sample underwent a clean, dateable reform of its anti-corruption apparatus during 2016-2022.

6. RESULTS

6.1 DESCRIPTIVE EVIDENCE

The cross-sectional pattern in the data deserves attention before regression results are presented. Table 2 reports mean impunity scores and mean education outcomes for each state.

Table 2: State-Level Averages, 2016-2022

State	Mean NCRB	Mean Dropout (%)	Mean GER
Uttar Pradesh	0.891	7.22	92.72
Jharkhand	0.855	5.70	95.00
Assam	0.909	5.04	107.64
Bihar	0.916	4.16	90.08
West Bengal	1.000	4.16	104.88
Rajasthan	0.770	3.42	100.12
Madhya Pradesh	0.466*	3.06	89.08
Chhattisgarh	0.717	2.28	94.78
Maharashtra	0.949†	1.42	102.60

Tamil Nadu	0.693	1.04	98.12
Gujarat	0.799	0.86	93.74
Odisha	0.809	0.84	94.94
Telangana	0.803	0.80	107.14
Kerala	0.640	0.66	99.04
Karnataka	0.864†	0.46	105.92

Source: NCRB Crime in India 2016-2022; UDISE+ Flash Statistics 2018-19 to 2022-23. * Madhya Pradesh NCRB score reflects ACB conviction artifact; see text. † Maharashtra registers more PCA cases than any other state in the sample, which mechanically raises the denominator of the impunity ratio even where conviction activity is comparatively normal (Section 9, registration paradox). Karnataka's score (0.864) may reflect a similar dynamic given its active anti-corruption bureau. This study does not have Karnataka's registration volumes to confirm it, so the comparison is suggestive.

The dropout gradient is the most immediate finding from Table 2. Uttar Pradesh records a mean dropout rate of 7.22 percent against a mean impunity score of 0.891. Kerala records 0.66 percent dropout against an impunity score of 0.640. Jharkhand (5.70 percent, 0.855) and Bihar (4.16 percent, 0.916) sit at the high end, while Karnataka (0.46 percent, 0.864) and Tamil Nadu (1.04 percent, 0.693) sit at the low end. The rank correlation is positive and visually apparent, though Maharashtra (0.949 impunity, 1.42 percent dropout) is a notable exception: a high-impunity state with relatively low dropout, suggesting that income effects and Maharashtra's urbanisation level partially offset the corruption channel for this outcome.

GER shows a noisier pattern. Several high-impunity states record GER above 100 (West Bengal, Assam), which reflects over-age student retention rather than genuine universal enrollment. Madhya Pradesh, with the lowest impunity score, has one of the lower GER values in the sample. The cross-sectional correlation between impunity and GER is weaker than for dropout, consistent with GER's sensitivity to demographic and infrastructure factors beyond the corruption channel.

6.2 REGRESSION RESULTS

Table 3 presents the two-way fixed effects regression results for all three outcome variables. The coefficient reported is for the NCRB impunity score; control variable coefficients are estimated but not reproduced in the main table.

Table 3: Two-Way Fixed Effects Estimates, NCRB Impunity Score Coefficients

Outcome Variable	Coefficient	Std. Error	p-value	Direction
Reading deficiency (%)	+6.301	3.413	0.107	Correct (positive)
Gross Enrolment Ratio	-4.131	2.855	0.155	Correct (negative)
Dropout rate (%)	-3.542	1.819	0.058	Wrong (negative)

Notes: All regressions include state and year fixed effects and control for log per capita NSDP and primary pupil-teacher ratio. Standard errors clustered at state level (HC1). N = 71 for dropout and GER models (out of a balanced maximum of 75, with

Chhattisgarh contributing 2 of 5 years and Kerala 4 of 5 years due to NCRB exclusions in those years); N = 29 for the reading model (out of a balanced maximum of 30, with Chhattisgarh contributing 1 of 2 years for the same reason).

6.3 READING DEFICIENCY

The reading deficiency coefficient is 6.301 percentage points (p = 0.107). A one-unit increase in the NCRB impunity score, moving from the theoretical limit of full conviction to complete prosecutorial failure, is associated with 6.3 additional percentage points of Class 5 students unable to read at grade level, controlling for per capita income, pupil-teacher ratio, and state and year fixed effects. The direction is consistent with the teacher absenteeism and fund diversion mechanisms in Section 3: states where anti-corruption enforcement produces fewer convictions tend, within their own history over 2016-2022, to show worse learning outcomes.

The coefficient does not achieve statistical significance at the 5 percent or 10 percent threshold. The low precision has a structural explanation beyond small N. The reading model uses ASER data from 2018 and 2022 only, producing 29 usable observations once Chhattisgarh's missing 2018 NCRB value is accounted for. Two-way fixed effects then absorb 15 state means and 2 year means, leaving only the change from 2018 to 2022 within each state, after removing the national trend, as the identifying variation. Estimating a coefficient from this residual with any precision requires that the within-state change in impunity between 2018 and 2022 correlates strongly with the within-state change in reading deficiency after the year fixed effect is removed. For states where impunity barely moved between 2018 and 2022, the reading model cannot detect a relationship even if one exists. The coefficient should be read as an indication of direction and order of magnitude rather than a precise estimate.

6.4 GROSS ENROLMENT RATIO

The GER coefficient is -4.131 (p = 0.155). The negative sign means higher impunity is associated with lower gross enrolment, which is directionally correct. The coefficient is not statistically significant. The GER outcome is further complicated by the presence of states where GER exceeds 100, which creates heterogeneity in the outcome's interpretation that the linear FE model cannot cleanly handle. The GER result provides weak corroborating evidence for the direction of the hypothesis but is not the primary empirical contribution of this paper.

6.5 DROPOUT RATE

The dropout rate coefficient is -3.542 (p = 0.058), the most problematic result, since the sign is wrong: a negative coefficient implies higher impunity predicts lower dropout, the opposite of the theoretical prediction.

The primary explanation is the COVID data artifact described in Section 4.3. For 2020-21, most states reported zero or near-zero dropout because schools were closed and administrative recording was suspended. High-corruption states, including Bihar (0 percent in 2020), Odisha (0 percent in 2020 and 2021), and UP (0 percent in 2020), are among the most affected. The presence of these zeros creates a

negative within-state association: in 2020, when impunity was at or near its highest for several states, dropout was artificially recorded at zero. The FE model interprets this pattern as: when impunity was high, dropout was low, an artifact of data collection failure rather than a substantive relationship.

This interpretation is examined in Section 8.1. Excluding 2020 and 2021 from the sample and reestimating produces a coefficient of -5.644 ($p = 0.146$), which is more negative and more imprecise rather than correct. The persistence of the wrong sign after COVID-year exclusion points to a secondary issue: floor effects in the dropout series for several otherwise well-performing states (Karnataka mean 0.46 percent, Gujarat mean 0.86 percent) leave very little within-state variation for identification. Dropout is concluded to be a weak outcome variable for two-way FE estimation in this setting.

7. OPPORTUNITY COST ESTIMATION

7.1 METHODOLOGY

The reading deficiency coefficient is used to estimate the number of additional elementary students who would read at grade level if high-impunity states reduced their impunity scores to the level of Kerala, the lowest-impunity state in the sample once Madhya Pradesh's conviction artifact is set aside (mean ncrb = 0.640). This is an associational counterfactual, not a causal prediction. The exercise asks: based on the within-state correlation estimated in the regression, what does the coefficient imply for the scale of the learning gap attributable to the impunity differential?

Four steps are applied for each of three target states: Uttar Pradesh, Bihar, and Rajasthan.

1. Compute the impunity gap between the target state and Kerala.

2. Multiply the gap by the FE coefficient (6.301) to obtain the predicted reduction in reading deficiency in percentage points.

3. Apply this percentage point reduction to the state's elementary enrollment, taken directly from Table 5.1 of the UDISE+ Flash Statistics 2021-22 report (Classes 1 to 8, all school management types), to obtain a student headcount.

4. State the assumptions and interpret the result as an order of magnitude rather than a point estimate.

7.2 STATE-LEVEL CALCULATIONS

Uttar Pradesh

Impunity gap: 0.891 - 0.640 = 0.251

Predicted reduction in reading deficiency: 6.301 x 0.251 = 1.58 percentage points

Elementary enrollment (UDISE+ Flash Statistics 2021-22, Table 5.1): 35,772,729 students

Implied student headcount: 0.015816 x 35,772,729 = approximately 566,000 students

Bihar

Impunity gap: 0.916 - 0.640 = 0.276

Predicted reduction in reading deficiency: 6.301 x 0.276 = 1.74 percentage points

Elementary enrollment (UDISE+ Flash Statistics 2021-22, Table 5.1): 21,433,075 students

Implied student headcount: 0.017391 x 21,433,075 = approximately 373,000 students

Rajasthan

Impunity gap: 0.770 - 0.640 = 0.130

Predicted reduction in reading deficiency: 6.301 x 0.130 = 0.82 percentage points

Elementary enrollment (UDISE+ Flash Statistics 2021-22, Table 5.1): 12,564,969 students

Implied student headcount: 0.008191 x 12,564,969 = approximately 103,000 students

Combined total across three states: approximately 1,041,000 additional students at grade-level reading.

7.3 INTERPRETATION AND CAVEATS

Several qualifications are essential. The regression coefficient carries substantial uncertainty ($p = 0.107$); the headcount therefore reflects that same uncertainty and should be interpreted as an order of magnitude, not a point estimate. The counterfactual assumes the impunity-outcome relationship is locally linear and approximately stable, neither of which is directly testable from this data. Kerala is used as the comparator because it has the lowest impunity score not driven by a known data artifact, but it is not a theoretical minimum:reducing impunity beyond Kerala's level would presumably yield further gains if the relationship holds.

The exercise does not translate these headcounts into rupee terms in the main text, because doing so requires wage return assumptions, employment projections, and discount rates that this study cannot independently estimate. The labour economics literature consistently finds positive returns to basic literacy, and over a million additional literate students would imply non-trivial productivity gains over their working lives, but the precision of any rupee figure would be driven by assumptions rather than evidence. The headcount alone conveys the human scale of what the regression coefficient implies.

The purpose of this exercise is to give empirical content to a coefficient that otherwise sits as an abstract statistical quantity. Taken at face value, the within-state coefficient applied to verified enrollment figures

implies an order of magnitude of around a million children across three states. Given p = 0.107, that figure illustrates what the estimate would mean in human terms if it holds. It is not a finding in its own right.

8. ROBUSTNESS CHECKS

Three robustness checks are reported, targeted at the main threats to the baseline results.

8.1 COVID YEAR EXCLUSION

The dropout model is most exposed to the COVID artifact in 2020-21 data. Excluding 2020 and 2021 and reestimating the fixed effects model produces a coefficient of -5.644 (p = 0.146) for dropout. The sign is unchanged and the magnitude is larger. This confirms that the COVID artifact is not the sole driver of the wrong sign: even without those years, within-state variation in dropout does not correlate positively with within-state variation in impunity. The floor effects discussed in Section 6.5 appear to be the residual problem, since states like Karnataka with consistently near-zero dropout leave no upward movement to be explained. The reading and GER results are not affected by COVID-year exclusion because ASER data for 2020 and 2021 does not exist in any form, and the GER COVID anomaly was handled separately.

8.2 LEAVE-ONE-OUT ANALYSIS

To assess whether any single state drives the reading deficiency result, the FE model is re-estimated fifteen times, each time dropping one state and recording the ncrb coefficient. Stability of the coefficient across these sub-samples would indicate that the result is not dependent on any one influential observation. West Bengal, which contributes no within-state variation (constant score of 1.000 across all seven years), should have no effect on the reading coefficient when dropped; this serves as an internal validity check. Madhya Pradesh is the most likely candidate for individual influence on the reading result, since its impunity score rises from 0.125 in 2018 to 0.626 in 2022, the largest within-state change in the reading sample, and this change likely reflects a conviction backlog artifact rather than a change in enforcement (see Section 9). Full leave-one-out results are available from the replication code in Appendix A.

8.3 LAGGED CORRUPTION

A one-year lag of the impunity score (ncrb(i, t-1)) replaces the contemporaneous score in an alternative specification. It is reasonable to think that education outcomes in year t are influenced by the corruption environment of the preceding year, since teacher absenteeism patterns and fund diversion arrangements develop over time. If the lagged coefficient is directionally similar to the main result, this provides partial protection against short-run reverse causality. The specification loses one year of data per state (2016

observations are used as lags for 2017 outcomes but cannot themselves have a lagged predictor). Results are available from the replication code.

9. LIMITATIONS

Each of the following limitations applies to the analysis and must be weighed by any reader evaluating its conclusions.

NCRB registration paradox. States with more active anti-corruption bureaus register more PCA cases, which mechanically inflates the denominator of the impunity score. Maharashtra registers the most PCA cases nationally each year and records a mean impunity score of 0.949, higher than Bihar and comparable to Jharkhand. This likely reflects genuine institutional activity in which cases are filed faster than courts can process them, rather than corruption of the same character as in states that register almost no cases. The impunity score conflates these two institutional situations, a problem that a composite index with multiple measures would partially address. Because this distortion is tied to enforcement activity rather than randomly distributed, it is a non-classical measurement error. It is the main reason this paper does not rely on the attenuation argument (see Section 4.5).

Conviction lag. Courts convict on cases registered in prior years, sometimes many years prior. Using same-year registrations and convictions introduces noise in the impunity measure, particularly for states with small annual case volumes where the ratio is sensitive to the timing of individual verdicts. Chhattisgarh's anomalous 2017-2020 scores directly reflect this problem.

Chhattisgarh thin data. Only three years are usable for Chhattisgarh (2016, 2021, 2022). Its mean impunity score of 0.717 should be treated with caution. A leave-one-out check excluding Chhattisgarh is recommended alongside the main results. **Madhya Pradesh conviction artifact.** Madhya Pradesh's 2017-2019 impunity scores (0.129, 0.125, 0.233) sit far below its 2016 and 2020-2022 values (0.626 to 0.854). This pattern is consistent with a burst of backlog convictions rather than a change in enforcement. Because these swings are large and within-state, Madhya Pradesh contributes disproportionate identifying variation to the fixed effects estimates, particularly the reading model, where its score moves from 0.125 in 2018 to 0.626 in 2022.

ASER not published annually. The reading deficiency model uses ASER data from 2018 and 2022 only. The 2016, 2017, 2019, 2020, and 2021 cells are blank and are not imputed. The FE model for reading effectively identifies from a single within-state change across a four-year gap, after removing the year fixed effect. This is the most significant structural limitation on the reading result.

COVID dropout zeros. Dropout data for 2020-21 are zero or near-zero for most states due to school closures and administrative disruption. These are not genuine zero dropout rates. The robustness check excluding these years does not resolve the wrong-sign result for dropout, indicating that floor effects in the dropout series are a secondary contributor.

GER above 100. Five states report mean GER exceeding 100. These values are not data errors but complicate regression interpretation. Net Enrolment Ratio would be a cleaner outcome variable for future work.

Reverse causality. Poor education outcomes may contribute to corruption by producing populations with weaker capacity to monitor and sanction public officials. The FE design partially addresses this by exploiting within-state variation over time, but full resolution requires an instrumental variable that is not developed here.

Static urbanisation. Census 2011 urbanisation figures are applied across all seven study years. India's urbanisation shifted during this period. In OLS specifications, this introduces measurement error; in the FE specification, the time-invariant urbanisation variable is absorbed by state fixed effects and does not affect the main estimates.

Telangana urbanisation approximation. Telangana's urbanisation figure (38.9%) is approximated from Andhra Pradesh Census 2011 data adjusted for the districts constituting Telangana at the time of bifurcation in 2014. This is standard practice for post-bifurcation analysis but introduces some imprecision.

Composite index not constructable. The original design intended a composite corruption index combining NCRB, Lokayukta, CAG, RTI, and TI data. Four of five intended components were unavailable as consistent panel variables. Relying on a single proxy exposes the estimate to both random and systematic measurement error (Section 4.5). The random component biases toward zero, but the systematic component has no determinate direction. The weak significance of the results should therefore be read as reflecting limited power and an imperfect measure, not as a lower bound on the true relationship. Difference-in-Differences not executed. The planned quasi-experimental design was infeasible within the study window. This is stated as a limitation rather than omitted from the paper.

Small N. Fifteen states is a small cross-section for a two-way fixed effects estimator. Approximately 20 degrees of freedom, 14 from state fixed effects and 6 from year fixed effects, are consumed relative to pooled OLS, from a panel that begins at 105 observations and loses more to missing data. Inferences from this sample should not be generalised without further validation on a larger state panel.

10. CONCLUSION

This paper examined whether prosecutorial impunity under the Prevention of Corruption Act predicts elementary education outcomes across 15 Indian states from 2016 to 2022. The evidence points in the expected direction, but the signal is weak and one result has the wrong sign entirely. Both findings are reported honestly rather than one of them being buried.

The reading deficiency result, 6.3 percentage points of additional Class 5 students unable to read at grade level per unit increase in the impunity score, is the most informative finding. The direction is consistent

with the theoretical channels of teacher absenteeism and fund diversion. The coefficient is not statistically significant at conventional thresholds. The opportunity cost exercise gives this coefficient empirical content: applied to verified UDISE+ 2021-22 elementary enrollment figures for Uttar Pradesh, Bihar, and Rajasthan, it implies an order of magnitude of just over 1 million additional students who could be at grade-level reading if those states matched Kerala's prosecutorial accountability. That figure is the associational implication of the regression estimate, translated into headcount terms, and not a causal prediction.

The dropout result is the least informative. A wrong-sign coefficient that persists after excluding COVID-contaminated years reflects a data quality problem that UDISE+ dropout series shares in the 2018-2023 window: floor effects in well-performing states and administrative zeros in COVID years leave insufficient within-state variation for the FE estimator to recover a positive coefficient even if a genuine positive relationship exists. Reporting this honestly is a service to researchers who might otherwise treat the dropout series as a reliable panel outcome variable in similar designs.

The descriptive evidence is, in some respects, the most legible element of the paper. The rank ordering of states by dropout rate corresponds closely enough to the rank ordering by impunity score, with UP and Jharkhand at the top and Kerala and Karnataka at the bottom, to be informative on its own, even without a causal interpretation. Cross-sectional patterns are not evidence of causation, but they establish that the hypothesis is at minimum consistent with the data rather than contradicted by it.

Three policy implications follow from the evidence, though each is stated at the level of confidence the data actually support. The most direct is institutional strengthening of state Lokayuktas. States like Bihar, which lack a functional anti-corruption oversight institution, and West Bengal, where the Lokayukta exists in statute but operates under significant jurisdictional restrictions, sit at the high end of the impunity distribution, and the study period contains no Lokayukta reform events simply because none occurred. Whether reform of these institutions would causally reduce corruption is a question that needs a longer time series and an identification strategy this paper could not execute, but the pattern in the data still makes it a reasonable policy priority. A second lever is extending the GeM (Government e-Marketplace) procurement platform to cover school construction and materials supply contracts, which would shrink the discretionary administrative space where procurement corruption operates and address the fund diversion channel directly. Third, moving ASER to annual rather than biennial surveys would substantially improve the statistical power of any reading deficiency model built on within-state variation, and would make the design used here considerably more useful for researchers who take it further.

Research that is honest about what it cannot show is more useful than research that projects false certainty, especially when the stakes involve the educational trajectories of hundreds of millions of children, and although the findings reported here are modest, the pattern underlying them looks large enough to be worth taking seriously.

REFERENCES

Bardhan, P. (1997). Corruption and development: A review of issues. Journal of Economic Literature, 35(3), 1320-1346.

Cameron, A. C., & Miller, D. L. (2015). A practitioner's guide to cluster-robust inference. Journal of Human Resources, 50(2), 317-372. https://doi.org/10.3368/jhr.50.2.317

Chaudhury, N., Hammer, J., Kremer, M., Muralidharan, K., & Rogers, F. H. (2006). Missing in action: Teacher and health worker absence in developing countries. Journal of Economic Perspectives, 20(1), 91-116. https://doi.org/10.1257/089533006776526058

Gupta, S., Davoodi, H., & Tiongson, E. (2000). Corruption and the provision of health care and education services (IMF Working Paper No. WP/00/116). International Monetary Fund.

Mauro, P. (1995). Corruption and growth. Quarterly Journal of Economics, 110(3), 681-712. https://doi.org/10.2307/2946696

Ministry of Finance, Government of India. (2017-2023). Economic Survey of India (Annual editions; education expenditure as a percentage of GDP, Social Services chapter). Department of Economic Affairs. https://www.indiabudget.gov.in

Muralidharan, K., Das, J., Holla, A., & Mohpal, A. (2017). The fiscal cost of weak governance: Evidence from teacher absence in India. Journal of Public Economics, 145, 116-135. https://doi.org/10.1016/j.jpubeco.2016.11.005

National Crime Records Bureau. (2016-2022). Crime in India, Volume II (Annual reports for data years 2016-2022). Ministry of Home Affairs, Government of India. https://ncrb.gov.in

Office of the Registrar General & Census Commissioner, India. (2011). Census of India 2011: Primary census abstract data highlights (DDW-PCA0000). Ministry of Home Affairs, Government of India. https://censusindia.gov.in

Pratham Education Foundation. (2018). Annual status of education report (rural) 2018. ASER Centre.

Pratham Education Foundation. (2022). Annual status of education report (rural) 2022. ASER Centre.

Reinikka, R., & Svensson, J. (2004). Local capture: Evidence from a central government transfer program in Uganda. Quarterly Journal of Economics, 119(2), 679-705. https://doi.org/10.1162/0033553041382120

Reserve Bank of India. (2016-2022). Handbook of statistics on Indian economy (Annual editions, state-wise NSDP tables). Reserve Bank of India. https://rbi.org.in

Unified District Information System for Education Plus. (2018-2022). Flash statistics (Annual reports, 2018-19 to 2022-23). Ministry of Education, Government of India. https://udiseplus.gov.in

APPENDIX A : REPLICATION CODE

The following R code replicates the main two-way fixed effects estimates. Data must be loaded from my_data.csv using the column names specified in the text. The working directory should be set to the folder containing the data file. The script renames the pupil-teacher ratio column to a clean identifier on load and enters per capita NSDP in log form to match the specification described in Section 5.1, where ndsp(i,t) is defined as the log of per capita NSDP.

```
library(plm)
library(lmtest)
library(sandwich)

# Load data and create panel frame
my_data <- read.csv("my_data.csv")

# Clean the pupil-teacher ratio column name (R mangles spaces and brackets
# from the original spreadsheet header into dots on import)
names(my_data)[names(my_data) == "ptr..Primary.only."] <- "ptr"
pdata <- pdata.frame(my_data, index = c("State", "Year"))

# Main FE model: Reading deficiency
fe_reading <- plm(reading ~ ncrb + log(ndsp) + ptr,
  data = pdata, model = "within", effect = "twoways")
coeftest(fe_reading, vcov = vcovHC(fe_reading, type = "HC1", cluster = "group"))

# Main FE model: Dropout rate
fe_dropout <- plm(dropout ~ ncrb + log(ndsp) + ptr,
  data = pdata, model = "within", effect = "twoways")
coeftest(fe_dropout, vcov = vcovHC(fe_dropout, type = "HC1", cluster = "group"))

# Main FE model: Gross Enrolment Ratio
fe_ger <- plm(ger ~ ncrb + log(ndsp) + ptr,
  data = pdata, model = "within", effect = "twoways")
coeftest(fe_ger, vcov = vcovHC(fe_ger, type = "HC1", cluster = "group"))

# Robustness: COVID year exclusion (dropout)
data_nc <- my_data[!(my_data$Year %in% c(2020, 2021)), ]
pdata_nc <- pdata.frame(data_nc, index = c("State", "Year"))
fe_drop_nc <- plm(dropout ~ ncrb + log(ndsp) + ptr,
  data = pdata_nc, model = "within", effect = "twoways")
coeftest(fe_drop_nc, vcov = vcovHC(fe_drop_nc, type = "HC1", cluster = "group"))

# Robustness: Leave-one-out for reading
states <- unique(my_data$State)
loo_coefficients <- sapply(states, function(s) {
```

```
  temp <- pdata.frame(my_data[my_data$State != s, ], index = c("State", "Year"))
  m <- plm(reading ~ ncrb + log(ndsp) + ptr,
    data = temp, model = "within", effect = "twoways")
  coef(m)["ncrb"]
})
print(data.frame(State_dropped = states, ncrb_coefficient = round(loo_coefficients, 3)))

# Robustness: Lagged impunity score
# plm:: is specified explicitly because base R's lag() and dplyr's lag()
# are not panel-aware; only plm's method respects the State-Year index
# structure of the pdata.frame, so an explicit namespace call protects
# this line if dplyr or another masking package is loaded elsewhere in
# the same session.
fe_reading_lag <- plm(reading ~ plm::lag(ncrb, 1) + log(ndsp) + ptr,
  data = pdata, model = "within", effect = "twoways")
coeftest(fe_reading_lag, vcov = vcovHC(fe_reading_lag, type = "HC1", cluster = "group"))
```

APPENDIX B : NOTES ON FIGURE CONSTRUCTION

Figure 1 should be a scatter plot of mean NCRB impunity score (x-axis, range 0.45 to 1.00) against mean dropout rate (y-axis, range 0 to 8), with state abbreviations as point labels. A least-squares fitted line through all points will show the positive cross-sectional association. Madhya Pradesh should be labelled explicitly as the outlier above the fitted line, with a note explaining the ACB artifact. West Bengal sits precisely at (1.000, 4.16) and should also be labelled as the extreme case.

Figure 2 should be a scatter plot of mean NCRB impunity score (x-axis) against mean GER (y-axis, range 85 to 110), with state labels. Points with GER above 100 should be marked with a hollow circle to signal the measurement concern. The fitted line will show a weakly negative association, though the relationship is noisier than for dropout.

Figure 3 should use the 29 state-year observations available for the reading model (15 states, ASER 2018 and 2022). The x-axis is the NCRB impunity score for the corresponding year; the y-axis is the state reading deficiency rate. Points should be coloured by year (one colour for 2018, another for 2022). A single fitted line through all 29 pooled observations will show the positive raw association before fixed effects are applied.

Figure 4 should be a horizontal bar chart, ordered from lowest to highest mean NCRB score, showing all 15 states. This figure makes the institutional variation in the sample immediately visible and allows readers to situate any state they are familiar with within the broader distribution. Note that Karnataka (0.864) and Maharashtra (0.949) appear in the upper half of the chart despite both states having comparatively active anti-corruption bureaus, which illustrates the registration paradox discussed in Section 9.